

Sorting on Survival Gains under Incomplete Health Insurance

Dominique Araya*, Wei G. Xiong†

June 8, 2026

Abstract

This paper studies whether patients sort into health care providers based on expected health gains when access to care entails financial exposure. We develop and implement a marginal treatment effect framework for a binary mortality outcome with both treatment and selection heterogeneity. Using individual-level hospital discharge data from Chile and a continuous instrument based on variations in relative private-public surgery prices, we estimate how the mortality effects of private hospital care vary with patients' unobserved resistance to seeking private care. We find large mortality reductions for clinically severe patients, but these high-gain patients are often the most reluctant to access private care. This negative sorting is concentrated among low-income patients, suggesting that incomplete insurance undermines allocative efficiency by deterring patients with the largest potential survival gains from using private care.

Keywords: *Hospital Choice; Marginal Treatment Effects; Roy selection*

JEL Codes: *I13; C21; C26*

*The Ohio State University, USA, email: araya.46@osu.edu.

†Universidad de los Andes, Chile, email: wxiong@uandes.cl.

1 Introduction

In many government-sponsored health systems, care is free at the point of use but constrained by limited capacity, often resulting in long waiting times. Some patients therefore seek faster access through the private sector, either by purchasing private insurance or paying out of pocket.¹ Whether and how public insurance should subsidize private care thus depends on a central economic question: does cost exposure improve the allocation of medical resources through patient choice? If patients who stand to benefit most from timely care are also the most willing to bear its financial cost, then cost-sharing can improve targeting by directing private capacity toward high-gain patients. If instead financial risk deters those patients from accessing private care, cost exposure may generate misallocation and avoidable health losses. The welfare consequences are further complicated if the gain from timelier private treatment is attenuated or amplified by differences in clinical quality across sectors.

We formalize this trade-off in a generalized [Roy \(1951\)](#) model in which patients differ in both the survival gains from private care and the financial costs they face. The direction of sorting is theoretically ambiguous. When financial costs correlate modestly with severity, patients with larger expected gains are more likely to select into private care, generating positive sorting on gains. In this case, cost exposure serves a useful targeting function, aligning private decisions with social objectives. When financial risk rises sufficiently with severity, however, it can deter precisely those patients who stand to benefit most, leading to negative sorting. Whether cost exposure improves or undermines allocative efficiency therefore depends on how patients' survival gains covary with their net cost of choosing private care, a source of heterogeneity that is privately observed by patients but not by the researcher.

We study this question in Chile's mixed public-private health system, where the private sector is widely perceived as high-quality and offers faster access, but requires substantial out-of-pocket payments for publicly insured patients. Within the public system, care is generally free at the point of use, but public hospitals operate under substantial capacity constraints, especially for surgeries, with above-average waiting times compared with other OECD countries.² Delays to treatment can have severe consequences. For example, in hip fracture, longer waits for surgery are associated with higher mortality and more complications; in emergency general surgery, delays are associated with worse postoperative outcomes; and in cancer, even relatively short delays to treatment, including surgery, are associated with higher mortality across several major indications ([Pincus et al., 2017](#); [Meschino et al.,](#)

¹Examples include European countries such as the United Kingdom ([Propper, 2000](#)) and Italy ([Fabbri and Monfardini, 2009](#)), as well as Australia, New Zealand, Brazil and Mexico.

²See OECD, [Health at a Glance 2023](#).

2020; Hanna et al., 2020).

In this setting, we empirically estimate the mortality difference between public and private hospitals for publicly insured patients, using individual-level hospital discharge data linked to death records. We develop and implement a marginal treatment effect (MTE) framework (Heckman and Honoré, 1990; Heckman and Vytlačil, 2005) for a binary survival outcome that accommodates both selection heterogeneity and treatment heterogeneity, recovering how mortality differences vary with unobserved resistance to private care. We document negative sorting with respect to both the level of mortality risk and the magnitude of survival gains. Patients with higher mortality risk are more resistant to private care, even though they stand to gain the most in survival from private admission. This pattern is concentrated among low-income patients, consistent with financial exposure distorting access to high-value care. These findings suggest that incomplete insurance can undermine allocative efficiency by deterring high-gain patients from accessing private providers. We also present additional evidence suggesting that the estimated gains from private care are unlikely to be driven primarily by superior clinical quality and instead more plausibly reflect differential access to timely treatment.

Our econometric model closely follows the recent work of Acerenza et al. (2023), which adapts the MTE framework to nonlinear outcome settings. In the first-stage choice model, individuals choose between staying in the public system and opting for a private hospital. Under general assumptions of the local average treatment effect (LATE) (Vytlačil, 2002), the propensity score, conditional on observables and the excluded instrument, represents the cumulative density function of the unobserved resistance. In the second-stage treatment effect model, we estimate two marginal treatment response (MTR) curves based on the first-stage propensity score. These curves represent the potential survival probabilities in public and private hospitals, conditional on the propensity to choose private care and on observable characteristics. The MTE is then defined as the difference between these two potential outcomes.

To address endogeneity, we use a continuous choice shifter as an instrumental variable: the interaction between the relative generosity of the public voucher – measured as the ratio of the public reimbursement to the private-market price for a given surgery – and the number of private providers performing that surgery. This interaction proxies for the effective choice set of private hospitals willing to accept publicly insured patients under the voucher scheme. The key excluded variation comes from private-market prices, which are plausibly exogenous to the demand and unobserved mortality risk of publicly insured patients because they are negotiated in the privately insured market. The instrument thus shifts the likelihood of private hospital use without directly affecting patients’ survival probabilities.

Our methodology has clear advantages over the conventional two-stage least squares (2SLS) estimator. First, 2SLS identifies a local average treatment effect tied to a particular compliance margin, but it does not recover how treatment effects vary across individuals or how selection responds to expected gains; as a result, it provides limited guidance for welfare analysis when selection operates on gains rather than levels (Mogstad et al., 2021). Second, in the presence of substantial treatment-effect heterogeneity, a single 2SLS estimand averages over units with very different responses, potentially masking economically meaningful effects for key subpopulations. Third, death is a rare event in our data, with a mean mortality close to zero, so linear approximations to a binary outcome can be poorly behaved and difficult to interpret. By contrast, our MTE-based approach allows treatment effects to vary with unobserved resistance to private care, recovers the distribution of causal effects, and directly characterizes selection on survival gains.

The policy implications of our findings are particularly relevant in countries with capacity constraints in government-sponsored health care. Evidence on the direction of sorting on gains helps evaluate whether and how public subsidies should expand access to private providers for publicly insured patients. If treatment in private hospitals offers a genuine survival advantage, restricting access for financial reasons may lead to avoidable mortality and under-use of life-saving capacity. By identifying the magnitude and variability of survival gains, this paper provides the empirical basis for assessing whether policies that lower financial barriers can improve overall survival.

Related Literature. Our paper contributes to several strands of literature. The first is the literature estimating the causal impact of medical care on mortality in the presence of patient self-selection. This literature largely consists of two approaches. One exploits quasi-experimental variation to identify hospital effects on survival, including ambulance referral patterns (Almond et al., 2010), neonatal thresholds (Doyle et al., 2015) and rainfall (Adhvaryu and Nyshadham, 2015), and consistently finds that causal estimates exceed naïve comparisons – implying positive sorting on gains. Another approach uses distance as an instrument, combined with structural methods (e.g., Geweke et al. (2003); Einav et al. (2025)) or machine learning techniques (Olenski and Sacher, 2024) to recover quality estimates while explicitly modeling selection. In these settings, however, patients typically do not face the financial cost of care, and selection is driven largely by clinical assignment rather than individual trade-offs. We contribute to this literature by deploying an intermediate approach that requires fewer structural assumptions to reveal the direction of selection in the context of active patient choice under incomplete insurance.

The second contribution concerns the classic debate over whether cost exposure improves

the targeting of publicly funded benefits. A longstanding view in public economics is that public provision of low-quality free health care serves an important redistributive role, allowing higher-income households to opt for private care with shorter waiting times (Besley and Coate, 1992; Hoel and Sæther, 2003). In parallel, the adverse-selection view of insurance suggests that, if riskier individuals are more likely to purchase private coverage, self-selection may help allocate high-risk patients to faster care (Akerlof, 1970). More generally, the self-targeting literature argues that monetary or non-monetary costs can improve program targeting by screening in individuals with higher expected benefit (Nichols and Zeckhauser, 1982). However, empirical work shows that such screening can improve targeting in some settings, but also that application costs and administrative burdens may screen out vulnerable recipients, reducing take-up and welfare (Alatas et al., 2016; Dupas et al., 2016; Deshpande and Li, 2019; Finkelstein and Notowidigdo, 2019). More recently, Shepard and Wagner (2025) show that “ordeal” targeting is unlikely to perform well in selection markets in the presence of adverse selection. We contribute to this literature by studying cost exposure in health care, where costs not only screen for willingness to pay but also for patients with the largest expected survival gains.

Because our empirical setting features both a gap in out-of-pocket costs and a potential gap in quality between public and private hospitals, this article also relates to the literature that concerns the implications of shifting care delivery between public and private health systems. Existing work has examined this question primarily through institutional reforms that reassign patients across sectors. In the United States, studies of the Military Health System and the Veterans Affairs system document mixed evidence on differences in quality and outcomes between public and private providers (Frakes et al., 2023; Chan et al., 2023). Outside the U.S., related questions have been explored in European health systems, often in the context of regulated competition or outsourcing of care (see Kruse et al. (2018) for a review), as well as in Latin America, where public financing frequently coexists with private provision (Bronsoler et al., 2024). While this literature focuses on average quality differences across sectors, we contribute by emphasizing heterogeneity in survival gains and the role of patient selection. Rather than asking whether private care is uniformly better or worse, we study which patients benefit from private care and whether institutional arrangements – such as incomplete insurance coverage – facilitate or hinder the efficient allocation of patients across public and private providers.

This paper also makes a methodological contribution by extending the marginal treatment effect (MTE) framework to the binary outcome setting, where linear approximations are inappropriate. While the MTE approach has been widely applied to study heterogeneous treatment effects, most existing applications focus on continuous outcomes and rely on linear

or quasi-linear specifications (Carneiro et al., 2011; Cornelissen et al., 2018; Mountjoy, 2022; Borghesan and Vasey, 2024). The closest related application in context is Kowalski (2022), which analyzes selection heterogeneity in a clinical trial environment, where the instrument variable is binary rather than continuous. Relative to this literature, our contribution is to demonstrate how the MTE framework can be used to study selection on gains in a binary mortality outcome, a commonly observable outcome in empirical health economics but one that has received limited attention in the MTE literature.


The paper proceeds as follows. Section 2 describes the institutional setting and presents descriptive statistics using our data. Section 3 presents our econometric model. Section 4.1 provides details on our instrumental variable and estimating procedure. Section 5 shows the results, and section 6 discusses the mechanism. Section 7 concludes.

2 Institutional setting and data

2.1 Setting


Chile’s mixed health care system is characterized by a largely segmented public and private sector, with separate payers and providers. Figure 1 summarizes the main institutional features. By law, workers must contribute 7% of taxable income either to the public insurance fund or toward the purchase of private insurance. In practice, this mandate functions as an income-based tax credit, up to a cap, for private insurance. Figure 2 illustrates the income composition of public insurance beneficiaries and private insurance enrollees.

Figure 1: The mixed healthcare system in Chile



Our analysis focuses on beneficiaries of the public insurance system, the National Health Fund (Fondo Nacional de Salud, FONASA), which covers roughly 75% of the population.

Figure 2: Income distribution by health insurance groups



FONASA beneficiaries are classified into four income groups, labeled A through D, in ascending order of household per capita income. As shown in Figure 2, these categories are strongly associated with income, although they do not correspond exactly to fixed income thresholds because eligibility depends in part on the number of dependents.

Within the public system, care is generally free at the point of use, but hospital choice is limited. Patients are typically assigned to providers through a referral system based on place of residence, and public hospitals operate under substantial capacity constraints. These constraints are reflected in long waiting times for specialist consultations and elective procedures: for example, median waiting times are at least three months for digestive surgery and roughly six months for orthopedics (Benítez, 2026).

To expand provider choice, FONASA operates a voucher scheme known as the Free Choice Model (Modalidad de Libre Elección, MLE), under which beneficiaries may purchase covered services from any participating provider, public or private. The public insurer administratively sets the voucher value for each service and subsidizes 30% of that value.³ The voucher list is updated annually and covers a broad range of services, including outpatient consultations, tests, specialty procedures, and surgeries requiring hospitalization.

³The subsidy is 50% for childbirth-related procedures.

Private hospitals may decide whether to participate in the MLE system and, if so, which vouchers to accept. For most private providers, however, the main source of revenue comes from privately insured patients, who pay fee-for-service prices negotiated between private insurers and hospitals. If a publicly insured patient seeks treatment at a private hospital that does not accept the relevant MLE voucher, the patient must pay the full list price out of pocket, which is often substantially higher than the voucher reimbursement. About 10% of Fonasa beneficiaries voluntarily purchase commercial complementary insurance that partially or fully reimburses these out-of-pocket expenses. Supplementary coverage increases with income: approximately 20% of patients in group D hold such insurance, compared with 12% in group C, 6% in group B, and 3% in group A.⁴

Although private hospitals primarily serve privately insured patients – who account for only about 15% of the population and are, on average, wealthier and healthier (Sapelli and Vial, 2003) – Figure 3 shows that private-sector capacity exceeds this share for both surgical beds and intensive care units (IMCU and ICU). This pattern suggests that private hospitals are less capacity-constrained than public hospitals, particularly for surgical and high-intensity care. While our data do not directly measure waiting times, these capacity differences are likely to translate into shorter delays for surgical procedures and faster access to critical care in the private sector. By contrast, the public sector’s more limited capacity relative to its patient population may generate congestion and treatment delays, especially for patients requiring timely intervention.


2.2 Data and descriptive statistics

Our analysis primarily draws on hospitalization discharge records from all Chilean hospitals, public and private, between 2012 and 2017. For each hospitalization, the data report patient demographics, admission and discharge dates, the treating hospital, type of health insurance, primary diagnosis, primary surgery (if any), admission channel, and discharge status. We augment these records with several complementary data sources. First, we link discharge records to official death records to measure post-admission mortality. Second, we combine the annual MLE voucher schedules with private insurers’ payments to private hospitals to construct the relative voucher-to-private price for each service. Third, we use Diagnosis-Related Group (DRG) records from public hospitalizations to construct disease-level severity and mortality-risk measures for the publicly insured population.⁵ Finally, to compare patient self-selection with centralized assignment, we use transfer records from the Centralized Bed

⁴Commercial insurance policies are offered by private firms, including banks and international insurance groups. See the regulator’s [report](#) based on the 2015 national socioeconomic survey.

⁵We construct these disease-level risk measures using public-hospital DRG records from 2017-2018.

Figure 3: Share of hospital beds by sector



Management Unit, which allocates beds to publicly insured patients when their designated public hospital is at capacity.

Table 1 reports hospital choices by insurance type. Although private insurance covers both public and private hospitals, 93% of privately insured patients are treated in private hospitals. This pattern suggests a strong revealed preference for private care and is consistent with substantial perceived differences in quality across sectors.

Table 1: Overview of hospital discharges (2012-2017)

Insurance	Hospitals			Total
	Public	Private	Special-force	
Public	6,102,320 (85.2%)	981,089 (13.7%)	81,250 (1.1%)	7,164,659 (100%)
Private	111,827 (5.9%)	1,763,630 (92.6%)	30,143 (1.6%)	1,905,600 (100%)
Special-force	186,942 (21.2%)	321,240 (36.5%)	373,131 (42.3%)	881,313 (100%)
Total	6,401,089 (64.3%)	3,065,959 (30.8%)	484,524 (4.9%)	9,951,572 (100%)

Notes: “Special-force” refers to the separate system for military and police, and an integrated occupational care system.

We focus on approximately 7 million hospitalizations of publicly insured patients, 85% of

which occur in public hospitals and 14% in private hospitals. Table 2 compares observable characteristics across these two groups. The descriptive evidence points to both differential treatment and patient selection. On the treatment side, admissions to private hospitals are markedly more likely to involve surgery and modestly more likely to involve ICU or IMCU care, consistent with greater effective capacity in the private sector. On the selection side, private-hospital admissions involve diagnoses with lower in-hospital mortality, lower DRG weights, and lower severity and mortality-risk scores, suggesting that patients who choose private care are, on average, less severely ill on observables. Differences in demographics are modest, whereas income differences are large: private-hospital patients are substantially more concentrated in the highest-income group and underrepresented in the lowest-income group. Taken together, these patterns suggest that naïve outcome comparisons across sectors will reflect both differences in treatment opportunities and non-random sorting of patients.

Table 2: Characteristics of publicly insured patients by hospital sector

	Public hospitals	Private hospitals	Difference (Private-Public)
<i>Admission characteristics</i>			
Involve surgery	0.35	0.72	0.37***
Discharge from ICU/IMCU	0.03	0.04	0.01***
<i>Patient characteristics</i>			
Mean age	41.03	41.83	0.80***
Female	0.62	0.67	0.05***
Santiago resident	0.31	0.43	0.12***
Income group			
A	0.33	0.03	-0.30***
B	0.42	0.30	-0.13***
C	0.11	0.15	0.04***
D	0.14	0.52	0.39***
<i>Diagnostic disease severity</i>			
Expected in-hospital mortality	0.03	0.02	-0.01***
Average DRG weights	0.82	0.77	-0.06***
Severity index (1-3)	1.42	1.19	-0.23***
Mortality-risk index (1-3)	1.30	1.06	-0.24***
<i>Local condition (municipality)</i>			
Public bed occupancy	0.63	0.68	0.05***
Distance to nearest private hospital	39.47	21.62	-17.85***
Observations	6,102,320	981,089	7,083,409

Notes: ***statistical significance at a 1% alpha level. A disease is defined as a category in the International Classification of Diseases (ICD).

3 Econometric model

3.1 Conceptual framework: a generalized Roy model

We use a generalized Roy 1951 model to illustrate why the direction of sorting on gains is ambiguous when patients face financial risk in accessing private care. Conceptually, we decompose the benefit of private care into two components. The first is mortality-relevant and may arise from greater life-saving capacity, better life-saving technology, or shorter waiting times. The second is mortality-irrelevant and captures what we refer to as “consumer experience,” including better facilities, more personalized attention, reputation, or prestige.

Let $B(\theta)$ denote consumers’ willingness to pay for the survival benefit of private care relative to public care for a patient of type θ , where θ indexes underlying mortality risk. For expositional simplicity, we abstract from differences in clinical quality across sectors and assume that $B(\theta)$ is weakly positive, reflecting the private sector’s capacity advantage. As we show in the appendix, allowing for clinical-quality differences affects the magnitude of gains but not the direction of sorting.

We further assume that willingness to pay for the mortality-irrelevant component of private care does not vary with mortality risk, and denote it by X . A patient chooses private care whenever


$$OOP(\theta) < B(\theta) + X$$

where $OOP(\theta)$ is the out-of-pocket cost of treatment in private hospitals. Under incomplete insurance, $OOP(\theta)$ is increasing in θ , since patients with higher mortality risk typically require more intensive and therefore more costly treatment.

Figure 4 shows that the choice of private care depends on the shapes of the out-of-pocket cost function relative to the benefit function. The two panels illustrate alternative configurations of these objects. When the benefit $B(\theta)$ increases with θ and dominates financial risk, patients with higher expected gains are more likely to select into private care, generating positive sorting on gains. In contrast, when financial risk increases sufficiently with θ , it can offset the survival benefit, leading high-gain patients to avoid private care and producing negative sorting. These cases highlight that the direction of sorting depends on the relative slopes of survival gains and financial costs across patient types.

It should be noted that a uniform difference in clinical quality between the two sectors shifts $B(\theta)$ up or down in parallel, affecting the share of patients who opt for private care but not the direction of sorting. We illustrate these scenarios in the Appendix A.

Figure 4: Patient self-selection into private care



3.2 Empirical strategy

We estimate a generalized Roy model based on a latent variable binary choice model for selection into private hospital and a potential outcome model for mortality (Heckman and Vytlacil, 1999, 2005). We present our methodology using the notations by Mogstad and Torgovitsky (2024). In our context, “choice propensity” indicates a patient’s likelihood of selecting private care, “treatment” refers to private care, and “treatment effect” describes the impact of private care on mortality. Treatment effects are allowed to vary with unobservable “resistance to treatment”, which we interpret as each individual’s net cost of private care in our generalized Roy model. The MTE estimator measures how mortality changes for patients shifted into or out of private care by a marginal change in their choice probability.

3.2.1 First stage: binary choice

We are interested in estimating the effect of receiving care in a private hospital (treatment $D_i = 1$) on individual mortality outcome Y_i . Selection into private care follows a latent variable choice model

$$D_i = \mathbf{1}\{V_i \leq \nu(Z_i, X_i)\} \quad (1)$$

where X_i are covariates, Z_i are exogenous instruments, and V_i is a random variable of unobservable “resistance to private care”.

In our conceptual Roy model, V_i is an individual’s *net cost* of choosing private care, equivalent to $OOP(\theta) - [B(\theta) + X]$. Note that it is a function of one’s unobservable mortality risk θ , which naturally correlates with mortality outcome. Moreover, as illustrated in Figure

4, not only are gains in survival probability heterogeneous for medical reasons, but out-of-pocket costs are also heterogeneous due to unobservable commercial insurance and family wealth (the latter determines the shadow price of hospital expenses).

The MTE framework is well-suited to identify the unobserved heterogeneity in treatment effects. Under the general LATE assumptions (Vytlacil, 2002), instruments Z_i allows for identifying the distribution of V_i conditional on X_i via the propensity score

$$p(z, x) = \mathbb{P}[V_i \leq \nu(z, x) | X_i = x] \equiv F_{V|X}(\nu(z, x) | x) \quad (2)$$

Assuming that $V_i | X$ is continuously and uniformly distributed, we can define a new random variable $U_i \equiv F_{V|X}(V_i | X)$, the cumulative distribution of the unobserved resistance to treatment. Then, the individuals treated are those for whom their probability to choose the private system is above the U-quantile of the distribution of V_i .

$$D_i = \mathbf{1}[U_i < p(Z_i, X_i)] \quad (3)$$

3.2.2 Second stage: binary outcome

Let potential outcomes be $Y_i(1)$ (if treated in a private hospital) and $Y_i(0)$ (if treated in a public hospital), the MTE is defined as the ATE among subpopulations with the same propensity to use private care:

$$MTE(p, x) \equiv \mathbb{E}[Y_i(1) - Y_i(0) | U_i = p, X_i = x] \quad (4)$$

In MTE applications in which the outcome is a continuous variable such as wage or test score, the econometrician observes the outcome under one treatment such that $Y_i = D_i Y_i(1) + (1 - D_i) Y_i(0)$. In contrast, we observe a binary survival outcome at the individual level. We therefore define a latent index model for the potential outcomes

$$Y_i^*(d) = X_i \beta_d + \epsilon_d, \quad d \in \{0, 1\} \quad (5)$$

while the observed outcomes are

$$Y_i = D_i Y_i(1) + (1 - D_i) Y_i(0) = \mathbf{1}[D_i Y_i^*(1) + (1 - D_i) Y_i^*(0) \geq 0] \quad (6)$$

Following Acerenza et al. (2023), we use the stratified approach and define our separable

marginal treatment response (MTR) functions, which can be written as

$$\begin{aligned} MTR(1|p, x) &\equiv \mathbb{E}[Y_i(1) \mid U_i = p, X_i = x] = \mathbb{E}[Y_i^*(1) \leq 0 \mid U_i = p, X_i = x] \\ &= \frac{\partial \mathbb{E}[Y_i = 0, D_i = 1 \mid U_i = p, X_i = x]}{\partial p} \end{aligned} \quad (7a)$$

$$\begin{aligned} MTR(0|p, x) &\equiv \mathbb{E}[Y_i(0) \mid U_i = p, X_i = x] = \mathbb{E}[Y_i^*(0) \leq 0 \mid U_i = p, X_i = x] \\ &= -\frac{\partial \mathbb{E}[Y_i = 0, D_i = 0 \mid U_i = p, X_i = x]}{\partial p} \end{aligned} \quad (7b)$$

where $MTR(d|p, x)$ gives the probability that a patient with propensity score p of choosing private care would have died if they were admitted to a private hospital ($d = 1$) or a public hospital ($d = 0$). We provide detailed mathematical steps in the Appendix B.

As noted by Kowalski (2023), the shape of $MTR(0)$ captures the selection on levels, i.e., the underlying severity or risk. An upward-sloping $MTR(0)$ curve indicates that risky individuals with higher mortality risk in public hospitals exhibit greater resistance to private care; a downward-sloping $MTR(0)$ curve indicates the opposite.

We then define the MTE function as

$$MTE(p, x) = MTR(1|p, x) - MTR(0|p, x) \quad (8)$$

where positive (negative) values of the MTE function indicate that private hospitals increase (decrease) the survival probability compared to public hospitals.

3.3 Identifying assumptions

Identification of the MTE estimator relies on a set of assumptions standard in the generalized Roy and instrumental variables literature (Vytlacil, 2002; Heckman and Vytlacil, 2005).

- A.1. **Instrument Relevance:** The propensity score $p(z, x)$ is a nondegenerate function of z and is continuous in Z , with support given by the interval $[\underline{p}, \bar{p}]$.
- A.2. **Full Exogeneity:** The potential outcomes $Y^*(0)$, $Y^*(1)$ and the unobserved resistance term U are jointly independent of Z conditional on X .
- A.3. **Continuity:** The distribution of U is continuous with respect to the Lebesgue measure, conditional on X .
- A.4. **Overlap:** Both treatment states occur with positive probability $P(D = d) \in (0, 1)$ for $d \in \{0, 1\}$

A.5. Regularity of potential outcomes: The latent indices $Y^*(d)$ that govern the potential outcome are continuous with respect to the Lebesgue measure.

Assumptions A.1 and A.2 correspond to the standard instrumental variables conditions of relevance and conditional exogeneity. In particular, Assumption A.2 implies first-stage independence: conditional on observables X , the instrument affects outcomes only through its impact on treatment status and is independent of unobserved determinants of both treatment and survival. Together with Assumptions A.3-A.5, these conditions are equivalent to the LATE assumptions, under which treatment choice can be represented as a latent threshold-crossing model (Vytlacil, 2002). This set of assumptions permits identification of marginal treatment effects by tracing how outcomes vary with the latent resistance to treatment.

4 Estimation

4.1 Instrumental variable

We employ a continuous instrumental variable to identify treatment effects in the presence of unobserved heterogeneity in both hospital choice and mortality. The instrument exploits supply-side variation in private hospitals' willingness to accept public vouchers for surgical procedures, driven by private market conditions and plausibly orthogonal to unobserved mortality risk among publicly insured patients.

The choice shifter captures the choice set of private hospitals that are willing to accept the public voucher for a given surgery for a publicly insured patient. It is constructed by combining three components: (i) the regulated public voucher value for each surgical procedure; (ii) the median price paid by private insurers for the same surgery in the private market; and (iii) the number of private hospitals that perform the surgery in the local market. Let R_{st} denote the ratio of the publicly set voucher reimbursement to the median private-market price paid by private insurers for surgery s in year t , and let N_{st} denote the number of private hospitals performing surgery s in year t . Our excluded instrument is their interaction, $Z_{st} = R_{st} \times N_{st}$. The economic logic is that a publicly insured patient's choice set among private hospitals depends jointly on how competitive the voucher reimbursement is relative to prevailing private prices – captured by R_{st} – and on how many alternative private providers exist for that surgery – captured by N_{st} . Neither component alone is sufficient: a favorable price ratio generates little additional access when there are few private providers, and a large number of providers shifts patients toward the private sector only when reimbursement is competitive. It is their interaction that captures the effective supply

of private hospital care available to publicly insured patients.

A notable advantage of an instrument that varies at the surgery-year level rather than the individual level is that it is orthogonal to all individual-level unobservables by construction. Formally, let ε_i denote unobserved individual heterogeneity in mortality risk. Since Z_{st} varies only across surgery and year cells and is determined by hospital supply conditions, we have $Cov(Z_{st}, \varepsilon_i | s, t) = 0$ by construction – no individual characteristic can influence a surgery-year-level price ratio or provider count. The exclusion restriction therefore only requires that surgery-year-level supply conditions affect mortality solely through private admission, conditional on fixed effects, which is a strictly weaker requirement than ruling out correlation between the instrument and individual unobservables. This stands in contrast to individual-level instruments – such as distance to the nearest private hospital – where the exclusion restriction must additionally rule out direct effects operating through individual-specific channels.

Crucially, we include the main effects R_{st} and N_{st} directly in the control vector X_i , so that identification comes exclusively from the interaction $Z_{st} = R_{st} \times N_{st}$. These controls absorb variation related to procedure complexity and baseline risk, ensuring that identification is driven by private market conditions that shift provider participation but do not directly affect patient mortality. The key excluded variation derives from private-market prices, which are negotiated exclusively between private insurers and hospitals and are therefore determined by conditions in the privately insured market, not by the demand, health status, or risk profile of publicly insured patients.


Figure 5 shows the distribution of the instrumental variable in the analytical sample. Panel (a) presents the distribution of the relative price ratio, while panel (b) displays the resulting distribution of the continuous instrumental variable. We hypothesize that this interaction expands the effective choice set of private hospitals available to publicly insured patients and increases the propensity to seek private care.

4.2 Sample Selection

Our sample selection procedure serves two purposes. First, we restrict attention to the subpopulation for whom mortality is a clinically meaningful and empirically informative outcome. Second, we construct an analytic sample in which hospital choice is plausibly responsive to our instrumental variable. Figure 6 summarizes the selection steps.

We begin with the universe of hospital discharges for publicly insured patients between 2012 and 2017. To focus on settings in which mortality is informative, we retain only each patient’s last admission for a given diagnosis and exclude diagnoses with near-zero

Figure 5: Relative price by the number of private providers as a choice shifter



in-hospital mortality. Restricting the sample to clinically serious conditions is important because it reduces the influence of purely risk-based sorting, as in Table 2, where patients selecting private hospitals are, on average, less severe.

We then impose additional restrictions to obtain a sample for which our instrumental variable is relevant. Since our instrument is defined at the level of surgery, we retain only admissions involving an identifiable primary surgery and beneficiaries who are eligible to purchase public vouchers. Because private hospitals are heavily concentrated in the Santiago metropolitan area, we further restrict the sample to Santiago residents in order to limit geographic variation in access to private providers that is unrelated to voucher incentives. We also retain only surgeries observed in both sectors and across higher- and lower-income groups, ensuring meaningful within-surgery variation in hospital choice. Finally, we exclude surgeries with very small cell sizes and a small number of procedures with implausibly large year-to-year changes in private prices, which are likely to reflect coding or measurement error.

Taken together, these restrictions yield a final analytic sample of 55,369 admissions. This final sample is therefore designed to satisfy both substantive and econometric objectives: mortality is a meaningful outcome, and the instrument has sufficient power to shift hospital choice along a relevant margin.

4.3 Model specification

As is standard in the empirical MTE literature, we assume additive separability between observed covariates X_i and unobserved resistance U_i (Mogstad and Torgovitsky, 2024). Identification relies on the exogeneity of Z_{st} conditional on X_i , which includes four sets of covariates.


First, we include the main effects R_{st} and N_{st} directly in the control vector X_i , so that

Table 3: Descriptive statistics for the analytic sample

	Full sample	Public hospitals	Private hospitals	Difference (Private-Public)
<i>Outcomes</i>				
28-day mortality	0.02	0.03	0.02	-0.01***
90-day mortality	0.05	0.05	0.03	-0.02***
180-day mortality	0.07	0.08	0.05	-0.03***
<i>Patient characteristics</i>				
Mean age	50.8	50.0	52.5	2.5***
Female	0.57	0.56	0.57	0.01
Indigenous	0.27	0.30	0.20	-0.10***
Readmission, same diagnosis	0.11	0.12	0.08	-0.05***
Emergency admission	0.32	0.38	0.19	-0.19***
Income group				
B	0.41	0.49	0.25	-0.24***
C	0.12	0.15	0.07	-0.08***
D	0.46	0.37	0.68	0.31***
<i>Disease characteristics</i>				
Expected in-hospital mortality	0.05	0.05	0.05	-0.00***
Average DRG weights	1.18	1.19	1.14	-0.05***
Severity index (1-3)	1.54	1.55	1.51	-0.04***
Mortality-risk index (1-3)	1.45	1.46	1.42	-0.04***
<i>Local condition and private market variables</i>				
Public bed occupancy	0.70	0.69	0.71	0.02***
Public/private price ratio	0.74	0.74	0.74	-0.00
Number of private providers	16.3	16.1	16.7	0.61***
Observations	55,369	38,116	17,253	55,369

Notes: This table reports descriptive statistics for the final analytic sample after imposing the sample restrictions described in Figure 6. The difference column reports the mean difference between patients admitted to private and public hospitals. Statistical significance is indicated by *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

Figure 6: Sample Selection



identification comes exclusively from their interaction, $Z_{st} = R_{st} \times N_{st}$. Conditional on the main effects, the exclusion restriction requires that this interaction affect outcomes only through hospital choice.

The second set includes surgery, year, and municipality fixed effects to absorb three important sources of confounding that could otherwise compromise the instrument's exogeneity. Surgery fixed effects control for systematic differences across procedures in regulated voucher values, private prices, clinical complexity, and baseline mortality risk. This is especially important because our instrument is partly constructed from surgery-specific relative prices. Year fixed effects absorb aggregate time variation in hospital markets, reimbursement schedules, technology, and macro health-system conditions that may simultaneously affect

private-hospital use and survival. Municipality fixed effects absorb local determinants of hospital choice and mortality, including geographic access to providers, commuting costs, local socioeconomic conditions, and shocks to the supply or congestion of the public health system.

The third set includes individual demographics and admitting characteristics: age, age squared and their interaction with sex, income group, an indicator for readmission with the same diagnosis, and an indicator for admission through the emergency department. These variables capture observable heterogeneity in both the demand for private care and baseline mortality risk. In particular, controlling for income group accounts for persistent differences across income categories in both access to private care and underlying health status, while readmission and emergency admission capture dimensions of illness severity and urgency that may affect both hospital choice and mortality.

The fourth set includes observed patient severity based on primary diagnosis, interacting with income group. At the diagnosis level, we construct the expected in-hospital mortality, DRG weights, severity, and mortality-risk measures for the publicly insured population. These controls capture systematic differences across conditions in clinical risk and treatment intensity. Interacting them with income group allows the relationship between diagnosis severity and hospital choice to vary with patients' financial resources, which is important in our setting because incomplete insurance may affect selection differently for richer and poorer patients facing the same condition.

Estimation of the two marginal treatment response functions in equation (7) then requires additional functional-form assumptions. The procedure we implement closely follows [Acerenza et al. \(2023\)](#).

1. **Estimate the propensity score.** We begin by estimating the first-stage choice model for private hospital admission,

$$D_i = \mathbf{1}[U_i < p(Z_i, X_i)],$$

where $p(Z_i, X_i) = \Pr(D_i = 1 \mid Z_i, X_i)$ is the propensity score. We estimate $p(Z_i, X_i)$ using a Logit specification and obtain the fitted propensity score $\hat{p}_i$ for each observation.

2. **Estimate the conditional distribution regressions.** For each $d \in \{0, 1\}$ estimate the conditional expectation $E[Y_i = 0, D_i = d \mid U_i = p, X_i = x]$ semi-parametrically given $p(z, x)$. To do so, we assume

$$\begin{aligned} \Gamma(p, x; y, d) &= E[Y_i = 0, D_i = d \mid U_i = p, X_i = x] \\ &= \Lambda(\beta_0(0, d) + X' \beta_x(0, d) + B_d(P)) \end{aligned} \tag{9}$$

where $Y_i = 0$ denotes death, Λ is a link function and B is a B-spline function.

3. **Impose the link function and estimate by maximum likelihood.** We assume that $\Lambda(\cdot)$ is the logistic cumulative distribution function, so that

$$\Gamma_d(p, x) = \frac{\exp(\beta_{0,d} + X'\beta_{x,d} + B_d(p))}{1 + \exp(\beta_{0,d} + X'\beta_{x,d} + B_d(p))}.$$

In estimation, we replace P_i by the fitted propensity score $\hat{p}_i$, and estimate $\beta_{0,d}$, $\beta_{x,d}$, and $B_d(\cdot)$ by maximum likelihood. In all specifications, we approximate $B_d(\cdot)$ with a quadratic B-spline (degree 2) with 7 degrees of freedom (5 internal knots).⁶

4. **Differentiate with respect to the propensity score.** For each $d \in \{0, 1\}$, we differentiate the fitted conditional expectation with respect to p :

$$\gamma_d(p, x) \equiv \frac{\partial E[Y_i = 0, D_i = d \mid U_i = p, X_i = x]}{\partial p}.$$

Under the logistic specification, this derivative is

$$\gamma_d(p, x) = B'_d(p) \Gamma_d(p, x) (1 - \Gamma_d(p, x)),$$

where $B'_d(p)$ denotes the derivative of the spline basis with respect to p . We compute $B'_d(p)$ using the `gensplines` command in Stata and evaluate $\hat{\gamma}_d(p, x)$ at the fitted values.

5. **Recover the marginal treatment response functions.** Using equation (7), we recover the two marginal treatment response functions as

$$\widehat{MTR}(d \mid p, x) = \frac{\sum_i^I 1(D = d, P = p, X = x) \hat{\gamma}_d}{\sum_i^I 1(D = d, P = p, X = x)}$$

By construction, $\widehat{MTR}(d \mid p, x)$ gives the estimated probability of death under treatment state d for an individual with observed characteristics x and latent resistance $U_i = p$. We additionally impose the natural restriction that the estimated MTR functions lie in the unit interval.

6. **Construct the marginal treatment effect function.** We then estimate the marginal treatment effect as the difference between the two marginal treatment response func-

⁶The B-spline specification was chosen for the full-sample model using Bayesian Information Criteria and held for all other estimates.

tions:

$$\widehat{MTE}(p, x) = \widehat{MTR}(1 | p, x) - \widehat{MTR}(0 | p, x).$$

Positive (negative) values of $\widehat{MTE}(p, x)$ imply that private hospital admission increases (decreases) mortality relative to public hospital admission for patients with resistance level p .


7. **Conduct inference.** We obtain confidence intervals using 999 block bootstrap replications clustered at the municipality-of-residence level.

4.4 Verification of identifying assumptions

To assess relevance and monotonicity (assumption A.1, also known as the uniformity condition), we examine the first-stage relationship across observable subgroups. Figure 7 reports the estimated effect of the instrument on private hospital admission separately by age, income group, diagnosis severity, and other pre-determined patient characteristics. The estimated first-stage coefficient is positive and statistically significant in all subgroups. This pattern supports the relevance condition by showing that the instrument consistently shifts the propensity to choose private care, and it is also consistent with monotonicity, as we find no evidence of subgroup-level reversals in the direction of the first-stage. While monotonicity is ultimately an assumption on unobservables and therefore cannot be tested directly, the stability of the positive first-stage relationship across observables provides supportive evidence for a complier interpretation in our final analytic sample.

Full exogeneity (Assumption A.2) requires that, conditional on observables, the instrument be independent of both the unobserved resistance to treatment and the potential mortality outcomes, except through its effect on hospital choice. In our setting, a violation would arise if the interaction between the public-to-private price ratio and the availability of private hospitals were correlated with unobserved determinants of patients' net cost of private care. We view this channel as unlikely. The key source of excluded variation comes from private-market prices, which are negotiated exclusively between private insurers and private hospitals and are therefore determined by conditions in the privately insured market rather than by the demand or health risk of publicly insured patients. As a result, changes in private prices should be unrelated to individual-level shocks affecting the unobserved resistance of publicly insured patients to private care. Moreover, the costs relevant for publicly insured patients are not the private prices themselves, but rather the copayment schedule set administratively by the government, together with other individual-specific barriers such as access frictions, liquidity constraints, and unobserved severity. Conditional on surgery, municipality-by-year, and patient and diagnosis characteristics, the excluded variation in the

Figure 7: Monotonicity of the instrumental variable by observable characteristics



Notes: ER, Emergency Room.

instrument is therefore plausibly orthogonal to both unobserved mortality risk and unobserved resistance to private care.

We further assess Assumption A.2 with a falsification exercise using patients in income group A, who are ineligible to purchase MLE vouchers and are therefore excluded from the main analysis. If the instrument were correlated with unobserved determinants of mortality or hospital choice beyond the voucher channel, it should retain predictive power in this population. Instead, we find a null first stage, with a Kleibergen-Paap F-statistic of 0.1, and no reduced-form effect of the instrument on mortality. This absence of predictive power in a population that cannot respond to voucher incentives supports the interpretation that the instrument affects outcomes only by shifting access to private care through the voucher system.

Figure 8 provides evidence in support of the overlap assumption (A.4). Both public and private hospital admissions are observed throughout the relevant region of the propensity-score distribution, indicating substantial common support in the final analytic sample. Because identification of the MTE relies on common support in the propensity score, we trim

the upper and lower 10% of the distribution, where overlap is thin, to avoid basing the estimates on poorly supported tail regions.

At the same time, the figure reveals economically meaningful sorting along this dimension. As discussed in the generalized Roy model in Section 3, an important component of resistance to private care in our setting is the financial burden created by incomplete insurance coverage. Because households differ in their ability to absorb medical expenditures, the shadow cost of out-of-pocket spending varies systematically with income and credit constraints. Consistent with this mechanism, panels (b) and (c) of Figure 8 show that the upper end of the propensity-score distribution is disproportionately populated by high-income patients (group D), whereas the lower-propensity region is concentrated among lower-income patients (groups B and C). Accordingly, in the low-income subsample we further restrict attention to the region $p \leq 0.7$, where overlap remains adequate for estimation.

Assumption A.3 requires that the latent resistance to treatment, U , be continuously distributed conditional on X . While this is not directly testable, it is plausible in our setting because U captures a composite of several continuously varying determinants of private hospital choice, including out-of-pocket exposure, liquidity constraints, supplementary insurance, travel frictions, and unobserved health risk. Consistent with this interpretation, the estimated propensity score exhibits a smooth distribution with no evidence of mass points in the relevant support.

Finally, Assumption A.5 requires that the latent potential-outcome indices be continuous conditional on observables. As with the continuity of U , this condition is not directly testable, but it is natural in our setting. The latent mortality index reflects the combined influence of observed covariates and unobserved determinants of health status, treatment intensity, and hospital performance, all of which are plausibly continuous. Under our parametric specification of the binary outcome model, this assumption ensures that the marginal treatment response functions are well defined and vary smoothly with the propensity score.


5 Empirical results

5.1 Two-Stage Least Squares


The conventional Two-Stage Least Squares (2SLS) recovers LATE for the margin of patients whose hospital choice is shifted by the instrument. It is not our preferred model for two reasons. First, 2SLS recovers only a single LATE and therefore cannot capture unobserved heterogeneity in treatment effects, which is central to our policy question and to a Roy-selection interpretation of hospital choice. Second, because death is relatively rare in our

Figure 8: Estimated Propensity Score by Sector


(a) Full sample



(b) High-income patients



(c) Low-income patients



Notes: the vertical lines in panel (b) indicate the cutoffs that divide resistance U_i into deciles.

sample, a linear approximation to the binary outcome is not well-suited to the data.

The 2SLS estimates are nonetheless informative. In Table 4, first-stage Kleibergen-Paap F -statistics exceed 28 in all specifications, confirming that the instrument is relevant. Across all horizons and specifications, the point estimates are negative, suggesting that private hospital care reduces mortality. Comparing columns (1) and (2), the inclusion of surgery, year, and municipality fixed effects reduces the magnitude of the estimates, consistent with these controls absorbing important systematic differences across procedures, locations, and time. The estimates in columns (2)-(4) are then relatively stable across specifications. At the same time, the standard errors remain large relative to the point estimates, so that most estimates are statistically indistinguishable from zero.

This is not evidence of no effect. Rather, the combination of a strong first stage and imprecise second-stage estimates is a hallmark of substantial unobserved heterogeneity in treatment effects: when the mortality reduction from private care is large for some patients and close to zero for others, a single linear IV estimand averages over these heterogeneous responses, inflating variance and masking the policy-relevant variation. This pattern directly motivates the MTE approach, which recovers the full distribution of treatment effects rather than forcing them into a single average.

Table 4: Two-Stage Least Squares Estimates

Independent variable: 1(private)	(1)	(2)	(3)	(4)
<i>Outcome: 28-day mortality ($N = 54,584$, $mean(std.) = 0.024(0.154)$)</i>				
Logit without IV	-0.011 (0.002)	-0.007 (0.002)	-0.004 (0.002)	-0.005 (0.002)
2SLS with IV	-0.082 (0.029)	-0.051 (0.057)	-0.060 (0.067)	-0.074 (0.067)
First-stage F -statistic	97.4	37.6	35.5	32.5
<i>Outcome: 90-day mortality ($N = 52,577$, $mean(std.) = 0.049(0.216)$)</i>				
Logit without IV	-0.022 (0.003)	-0.011 (0.003)	-0.009 (0.003)	-0.010 (0.003)
2SLS with IV	-0.266 (0.050)	-0.089 (0.077)	-0.107 (0.090)	-0.123 (0.088)
First-stage F -statistic	92.4	34.8	33.6	30.7
<i>Outcome: 180-day mortality ($N = 49,555$, $mean(std.) = 0.069(0.253)$)</i>				
Logit without IV	-0.033 (0.004)	-0.016 (0.003)	-0.014 (0.003)	-0.015 (0.003)
2SLS with IV	-0.408 (0.068)	-0.134 (0.090)	-0.164 (0.106)	-0.179 (0.108)
First-stage F -statistic	90.0	36.9	31.3	28.5
Instrument main effects	✓	✓	✓	✓
Surgery, year and municipality FE		✓	✓	✓
Demographic and admission characteristics			✓	✓
Disease severity controls				✓

Notes: Marginal effects reported for Logit. Standard errors clustered at the municipality level (52 clusters) in parentheses. First-stage F -statistics are Kleibergen-Paap rk Wald F -statistics.

5.2 Marginal Treatment Effects

Figure 9 presents the estimated MTR and MTE curves at the 28-, 90- and 180-day horizons. The results reveal three patterns. First, $MTR(0 | p, x)$ – the potential mortality under public hospital admission – rises steeply with unobserved resistance to private care. This implies strong *selection on levels*: patients who are most resistant to private care are also those facing the highest mortality risk if treated in the public sector. In other words, unobserved resistance is positively associated with underlying clinical severity. Second, $MTR(1 | p, x)$ – potential mortality under private admission – is comparatively flat and low across the support of the propensity score. The private sector thus provides a survival floor that the public sector cannot match for severely ill patients, who likely face congestion and delays. Third, the implied $MTE(p, x) = MTR(1 | p, x) - MTR(0 | p, x)$ is negative over most of the support and becomes increasingly negative as resistance rises. This is the signature of *negative sorting on gains*: the patients who would benefit most from private admission – those at the high end of the resistance distribution – are precisely those least likely to access it. The pattern is qualitatively identical at all three horizons and becomes stronger at longer horizons.

Thus, the two dimensions of selection reinforce each other. High-resistance patients are not only sicker on average, as seen in the upward-sloping $MTR(0|p, x)$, but also derive larger mortality reductions from private care, as seen in the downward-sloping MTE curve. These estimates suggest that incomplete insurance distorts the allocation of care. Rather than sorting private treatment toward the patients with the highest expected mortality reductions, financial exposure appears to deter exactly those patients from accessing private hospitals. This combination of upward-sloping $MTR(0|p, x)$ and downward-sloping $MTE(p, x)$ is the central empirical signature of negative sorting in our setting, and it motivates the heterogeneity analysis that follows, where we examine whether these patterns are concentrated among patients facing tighter financial constraints.

The estimated mortality levels are high: predicted mortality reaches nearly 40% at the 28-day horizon and close to 50% at the 90-day horizon. These magnitudes are plausible for three reasons. First, our sample is restricted by construction to diagnoses with above 1% in-hospital mortality. Within this high-acuity sample, mortality varies enormously across procedure types – neurosurgical or oncological surgeries carry mortality rates of 30-40% – and these high-risk conditions receive non-trivial weight in the average MTR , pulling the estimated levels upward. Second, our estimates are based on a trimmed propensity score support of $p \in [0.1, 0.9]$, excluding patients who are either almost certain to choose public care ($p < 0.1$) or almost certain to choose private care ($p > 0.9$). The complier margin in the interior is likely to have greater clinical severity than that of patients near the

boundaries of the excluded tails. Third, our specification includes rich fixed effects, and the nonlinear estimation necessarily relies on cells with within-cell variation in mortality. Cells with no deaths (i.e. perfect classifiers) are effectively excluded, which shifts the estimation toward higher-risk parts of the sample. Using the average of the estimated MTEs as an approximation to the sample LATE, we obtain mean (standard deviation) effects of -0.089 (0.222), -0.158 (0.259), and -0.201 (0.256) for the 28-, 90-, and 180-day horizons, respectively. These averages are broadly consistent with the 2SLS estimates reported in column (4) of Table 4, while being modestly larger in absolute value. For these reasons, the MTR and MTE curves shown in Figure 9 are most informative in terms of their shape rather than their exact magnitudes. The economically relevant magnitude of private care’s treatment effect requires integrating these marginal estimates into *policy-relevant treatment effects* (PRTE) using appropriate weights.

5.3 Selection patterns by income

Figure 10 shows that the negative selection on severity documented in the full sample is concentrated among lower-income patients, even though we lack statistical power to make conclusive statements regarding sorting on gains. For this group – panels (b),(d) and (f) – $MTR(0|p, x)$ rises steeply with resistance, with negative $MTE(p, x)$ concentrated at the high-resistance tail. Thus, severe, lower-income patients are the most reluctant to use private care. By contrast, among higher-income patients – panels (a),(c) and (e), $MTR(0|p, x)$ are relatively flat, and $MTE(p, x)$ are negative throughout most of the support. These results suggest that incomplete insurance distorts access to private care primarily among financially constrained patients.


Taken together, these results sharpen the interpretation of the full-sample estimates. The overall negative sorting pattern is not simply a general feature of hospital choice, but arises specifically because lower-income patients with high expected gains are deterred from accessing private care. This is precisely the mechanism emphasized in the generalized Roy model: when the shadow cost of out-of-pocket medical spending is high, financial exposure can overturn the targeting benefits of self-selection and prevent high-gain patients from using the sector with greater effective capacity.

6 Mechanism and policy implications


The evidence of negative sorting on gains raises a natural follow-up question: what is the source of the private-sector advantage that high-resistance patients fail to access? Although

Figure 9: Marginal Treatment Responses and Marginal Treatment Effects


(a) 28-day mortality since admission



(b) 90-day mortality since admission




(c) 180-day mortality since admission




Notes: The figure plots the estimated $MTR(1 | p, x)$, $MTR(0 | p, x)$, and $MTE(p, x)$ against the resistance U . Shaded bands are 90% pointwise confidence intervals from 999 block-bootstrap replications clustered at the municipality level. The horizontal axis covers the region of common support in the full analytic sample.

Figure 10: Selection pattern by income level


(a) High-income, 28-day mortality




(b) Low-income, 28-day mortality




(c) High-income, 90-day mortality




(d) Low-income, 90-day mortality



(e) High-income, 180-day mortality



(f) Low-income patients, 180-day mortality



private hospitals' greater capacity – most notably, their ability to provide more timely access to surgery and intensive care – may be associated with lower mortality, this correlation does not by itself identify the underlying mechanism. Public and private hospitals also differ along other dimensions that are difficult to observe, including physicians' and nurses' clinical skills, technology adoption, organizational practices, and hygienic standards. Distinguishing between these mechanisms is not necessary for recovering the direction of patient sorting, but it matters for policy interpretation. If the advantage of private care is primarily capacity-based, then the relevant policy problem is one of access to timely treatment. If instead it

reflects superior clinical quality, the case for subsidizing private care is less about alleviating congestion and more about reallocating patients toward intrinsically better providers.

To shed light on this distinction, we examine centrally coordinated patient transfers. Publicly insured patients may be transferred by the Ministry of Health’s Centralized Bed Management Unit (Unidad de Gestión Centralizada de Camas, UGCC) when their designated public hospital lacks an available bed at the required level of care. The UGCC first searches for beds in public hospitals within the patient’s region and turns to private hospitals only when public capacity is exhausted. Crucially, this process differs from patient self-selection in two respects: transferred patients face no out-of-pocket cost, and hospital assignment is determined by real-time bed availability rather than patient choice.⁷ In this setting, financial barriers are removed and waiting-time differences are largely neutralized, so outcome differences should more closely reflect clinical performance rather than selection through willingness or ability to pay.

Conditional on obtaining an ICU or IMCU bed, Table 5 shows that mortality among centrally transferred patients is higher in private hospitals than in public hospitals – the opposite of the pattern observed under patient self-selection. This comparison is not fully causal and may still reflect residual differences in case severity, especially because roughly two-thirds of transferred patients are assigned to private hospitals. Nevertheless, the sign reversal is informative. Once treatment is centrally allocated and the financial cost of private care is eliminated, private hospitals do not appear to deliver uniformly better outcomes than public hospitals. Taken together, these results suggest that the survival advantage identified in our main analysis is more likely to reflect differential access to scarce capacity than superior clinical quality per se.

This interpretation also sharpens the connection to the generalized Roy model. In our main results, patients with the largest expected mortality reductions from private care are often those with the highest resistance to private admission, implying negative sorting on gains. The UGCC evidence suggests that this pattern is not driven by private hospitals having intrinsically higher clinical quality across the board, but rather by financially constrained patients being unable to access the sector with greater effective capacity. In other words, the relevant margin is not simply provider quality, but whether high-gain patients can reach timely surgical and intensive care.

The policy implication is therefore not that public insurance should indiscriminately expand private provision. Rather, if the private-sector advantage is primarily one of capacity, then the welfare gains from subsidizing private care come from improving targeting: reducing

⁷Clinicians at the designated hospital assess the patient’s medical needs, determine the required level of care, and submit the transfer request. Patients play no role in the assignment process.

financial frictions can allow patients with the largest expected survival gains to sort into the sector with available capacity. In this sense, relaxing cost exposure can restore the targeting role of patient choice by converting negative sorting on gains into positive sorting on gains. Subsidies to private care are thus most valuable not when they shift all patients toward private providers, but when they enable high-gain patients to access timely treatment that they would otherwise forgo.

Table 5: Mortality among centralized transferred patients (publicly insured and originally admitted to public hospitals, years 2010-2019)

Bed type	Destination hospital type	N	Mortality	Difference from public hospitals
Basic	Public hospitals	6,348	0.008	—
	Private hospitals	3,684	0.002	-0.003 (0.003)
	Special force hospitals	518	0.002	-0.001 (0.002)
ICU/IMCU	Public hospitals	11,680	0.092	—
	Private hospitals	32,071	0.107	0.019 (0.007)**
	Special force hospitals	1,553	0.089	0.004 (0.010)

** statistically significant at the 5% level.

Notes: The mortality difference is estimated with Logit regression. All difference estimates are conditional on year, the original region of transfer, whether the transfer crosses regions, the requested bed type and the medical cause. All standard errors are clustered at the original hospital level.

7 Conclusion

Private hospital care generates large mortality reductions for some publicly insured patients, but those patients are often the least likely to access it. This negative sorting on gains is concentrated among lower-income patients, consistent with incomplete insurance distorting access to high-value care. Our evidence further suggests that the private-sector advantage primarily reflects greater effective capacity, especially faster access to surgery and intensive care, rather than uniformly superior clinical quality.

The policy implication is straightforward. Subsidies to private care are most valuable when they relax financial barriers for patients with the largest expected survival gains. In this setting, the main welfare gain from public support is not broader use of private care per se, but better targeting of timely treatment to the patients who benefit most.

References

- Acerenza, Santiago, Vitor Possebom, and Pedro HC Sant’Anna (2023) “Was Javert Right To Be Suspicious? Marginal Treatment Effects with Censored Data.”
- Adhvaryu, Achyuta and Anant Nyshadham (2015) “Return to Treatment in the Formal Health Care Sector: Evidence from Tanzania,” *American Economic Journal: Economic Policy*, Vol. 7, pp. 29–57, DOI: <http://dx.doi.org/10.1257/pol.20120262>.
- Akerlof, George A. (1970) “The Market for ”Lemons”: Quality Uncertainty and the Market Mechanism,” *The Quarterly Journal of Economics*, Vol. 84, pp. 488–500.
- Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, Ririn Purnamasari, and Matthew Wai-Poi (2016) “Self-Targeting: Evidence from a Field Experiment in Indonesia,” *Journal of Political Economy*, Vol. 124, pp. 371–427, DOI: <http://dx.doi.org/10.1086/685299>.
- Almond, Douglas, Jr. Doyle, Joseph J., Amanda E. Kowalski, and Heidi Williams (2010) “Estimating Marginal Returns to Medical Care: Evidence from At-risk Newborns,” *The Quarterly Journal of Economics*, Vol. 125, pp. 591–634, DOI: <http://dx.doi.org/10.1162/qjec.2010.125.2.591>.
- Benitez, Alejandra (2026) “Waiting Time and Demand for Elective Surgery: Evidence from Chile,” Work in progress.
- Besley, Timothy and Stephen Coate (1992) “Workfare versus Welfare: Incentive Arguments for Work Requirements in Poverty-Alleviation Programs,” *The American Economic Review*, Vol. 82, pp. 249–261.
- Borghesan, Emilio and Gabrielle Vasey (2024) “The Marginal Returns to Distance Education: Evidence from Mexico’s Telesecundarias,” *American Economic Journal: Applied Economics*, Vol. 16, pp. 253–285, DOI: <http://dx.doi.org/10.1257/app.20220065>.
- Bronsoler, Ari, Jonathan Gruber, and Enrique Seira (2024) “Private Sector Provision as an “Escape Valve”: The Mexico Diabetes Experiment,” *The Review of Economic Studies*, Vol. 92, pp. 129–161, DOI: <http://dx.doi.org/10.1093/restud/rdae024>.
- Carneiro, Pedro, James J. Heckman, and Edward J. Vytlačil (2011) “Estimating Marginal Returns to Education,” *American Economic Review*, Vol. 101, p. 2754–81, DOI: <http://dx.doi.org/10.1257/aer.101.6.2754>.
- Chan, David C., David Card, and Lowell Taylor (2023) “Is There a VA Advantage? Evidence from Dually Eligible Veterans,” *American Economic Review*, Vol. 113, p. 3003–43, DOI: <http://dx.doi.org/10.1257/aer.20211638>.

- Cornelissen, Thomas, Christian Dustmann, Anna Raute, and Uta Schönberg (2018) “Who Benefits from Universal Child Care? Estimating Marginal Returns to Early Child Care Attendance,” *Journal of Political Economy*, Vol. 126, pp. 2356–2409, DOI: <http://dx.doi.org/10.1086/699979>.
- Deshpande, Manasi and Yue Li (2019) “Who Is Screened Out? Application Costs and the Targeting of Disability Programs,” *American Economic Journal: Economic Policy*, Vol. 11, pp. 213–248, DOI: <http://dx.doi.org/10.1257/pol.20180076>.
- Doyle, J., J. Graves, J. Gruber, and S. Kleiner (2015) “Measuring Returns to Hospital Care: Evidence from Ambulance Referral Patterns,” *Journal of Political Economy*, Vol. 123, pp. 170–214.
- Dupas, Pascaline, Vivian Hoffmann, Michael Kremer, and Alix Peterson Zwane (2016) “Targeting health subsidies through a nonprice mechanism: A randomized controlled trial in Kenya,” *Science*, Vol. 353, pp. 889–895, DOI: <http://dx.doi.org/10.1126/science.aaf6288>.
- Einav, Liran, Amy Finkelstein, and Neale Mahoney (2025) “Producing Health: Measuring Value Added of Nursing Homes,” *Econometrica*, Vol. 93, pp. 1225–1264, DOI: <http://dx.doi.org/10.3982/ECTA21016>.
- Fabbri, Daniele and Chiara Monfardini (2009) “Rationing the public provision of healthcare in the presence of private supplements: Evidence from the Italian NHS,” *Journal of Health Economics*, Vol. 28, pp. 290–304, DOI: <http://dx.doi.org/10.1016/j.jhealeco.2008.11.004>.
- Finkelstein, Amy and Matthew J Notowidigdo (2019) “Take-Up and Targeting: Experimental Evidence from SNAP,” *The Quarterly Journal of Economics*, Vol. 134, pp. 1505–1556, DOI: <http://dx.doi.org/10.1093/qje/qjz013>.
- Frakes, Michael D., Jonathan Gruber, and Timothy Justicz (2023) “Public and Private Options in Practice: The Military Health System,” *American Economic Journal: Economic Policy*, Vol. 15, p. 37–74, DOI: <http://dx.doi.org/10.1257/pol.20210625>.
- Geweke, John, Gautam Gowrisankaran, and Robert J. Town (2003) “Bayesian Inference for Hospital Quality in a Selection Model,” *Econometrica*, Vol. 71, pp. 1215–1238, DOI: <http://dx.doi.org/10.1111/1468-0262.00444>.
- Hanna, Timothy P., William D. King, Stephanie Thibodeau, Michael Jalink, Gregory A. Paulin, Elizabeth Harvey-Jones, Dylan E. O’Sullivan, Christopher M. Booth, Richard Sullivan, and Ajay Aggarwal (2020) “Mortality Due to Cancer Treatment Delay: System-


- atic Review and Meta-Analysis,” *BMJ*, Vol. 371, p. m4087, DOI: <http://dx.doi.org/10.1136/bmj.m4087>.
- Heckman, James J. and Bo E. Honoré (1990) “The Empirical Content of the Roy Model,” *Econometrica*, Vol. 58, pp. 1121–1149.
- Heckman, James J. and Edward Vytlacil (2005) “Structural Equations, Treatment Effects, and Econometric Policy Evaluation,” *Econometrica*, Vol. 73, pp. 669–738, DOI: <http://dx.doi.org/j.1468-0262.2005.00594.x>.
- Heckman, James J. and Edward J. Vytlacil (1999) “Local Instrumental Variables and Latent Variable Models for Identifying and Bounding Treatment Effects,” *Proceedings of the National Academy of Sciences of the United States of America*, Vol. 96, pp. 4730–4734.
- Hoel, Michael and Erik Magnus Sæther (2003) “Public health care with waiting time: the role of supplementary private health care,” *Journal of Health Economics*, Vol. 22, pp. 599–616, DOI: [http://dx.doi.org/10.1016/S0167-6296\(03\)00007-9](http://dx.doi.org/10.1016/S0167-6296(03)00007-9).
- Kowalski, Amanda E (2022) “Behaviour within a Clinical Trial and Implications for Mammography Guidelines,” *The Review of Economic Studies*, Vol. 90, pp. 432–462, DOI: <http://dx.doi.org/10.1093/restud/rdac022>.
- (2023) “Reconciling seemingly contradictory results from the Oregon health insurance experiment and the Massachusetts health reform,” *Review of Economics and Statistics*, Vol. 105, pp. 646–664.
- Kruse, Florian M., Niek W. Stadhouders, Eddy M. Adang, Stef Groenewoud, and Patrick P.T. Jeurissen (2018) “Do private hospitals outperform public hospitals regarding efficiency, accessibility, and quality of care in the European Union? A literature review,” *The International Journal of Health Planning and Management*, Vol. 33, pp. e434–e453, DOI: <http://dx.doi.org/https://doi.org/10.1002/hpm.2502>.
- Meschino, Michael T., Andrew Giles, Rahima Nenshi, C. G. Ball, Guy Telford, Owen Anderson, Kevin Vogt, Neil Parry, S. Morad Hameed, Michael Simons, Andrew W. Kirkpatrick, John Kortbeek, Sarah Widder, and S. Morad Hameed (2020) “Operative timing is associated with increased morbidity and mortality in patients undergoing emergency general surgery: a multisite study of emergency general surgical services in a single academic network,” *Canadian Journal of Surgery*, Vol. 63, pp. E321–E328, DOI: <http://dx.doi.org/10.1503/cjs.012919>.
- Mogstad, Magne and Alexander Torgovitsky (2024) “Chapter 1 - Instrumental variables with unobserved heterogeneity in treatment effects,” Vol. 5 of *Handbook of Labor Economics*: Elsevier, pp. 1–114, DOI: <http://dx.doi.org/10.1016/bs.heslab.2024.11.003>.

- Mogstad, Magne, Alexander Torgovitsky, and Christopher R. Walters (2021) “The Causal Interpretation of Two-Stage Least Squares with Multiple Instrumental Variables,” *American Economic Review*, Vol. 111, p. 3663–98, DOI: <http://dx.doi.org/10.1257/aer.20190221>.
- Mountjoy, Jack (2022) “Community Colleges and Upward Mobility,” *American Economic Review*, Vol. 112, pp. 2580–2630, DOI: <http://dx.doi.org/10.1257/aer.20181756>.
- Nichols, Albert L. and Richard J. Zeckhauser (1982) “Targeting Transfers through Restrictions on Recipients,” *The American Economic Review*, Vol. 72, pp. 372–377.
- Olenski, Andrew and Szymon Sacher (2024) “Estimating Nursing Home Quality with Selection,” *The Review of Economics and Statistics*, pp. 1–31, DOI: http://dx.doi.org/10.1162/rest_a_01449.
- Pincus, Daniel, Bheeshma Ravi, David Wasserstein, Allan Huang, J. Michael Paterson, Avery B. Nathens, Hans J. Kreder, Robert Jenkinson, and Walter P. Wodchis (2017) “Association Between Wait Time and 30-Day Mortality in Adults Undergoing Hip Fracture Surgery,” *JAMA*, Vol. 318, pp. 1994–2003, DOI: <http://dx.doi.org/10.1001/jama.2017.17606>.
- Propper, Carol (2000) “The demand for private health care in the UK,” *Journal of Health Economics*, Vol. 19, pp. 855–876, DOI: [http://dx.doi.org/10.1016/S0167-6296\(00\)00045-X](http://dx.doi.org/10.1016/S0167-6296(00)00045-X).
- Roy, A. D. (1951) “Some Thoughts on the Distribution of Earnings,” *Oxford Economic Papers*, Vol. 3, pp. 135–146.
- Sapelli, Claudio and Bernardita Vial (2003) “Self-selection and moral hazard in Chilean health insurance,” *Journal of Health Economics*, Vol. 22, pp. 459–476, DOI: [http://dx.doi.org/10.1016/S0167-6296\(02\)00121-2](http://dx.doi.org/10.1016/S0167-6296(02)00121-2).
- Shepard, Mark and Myles Wagner (2025) “Do Ordeals Work for Selection Markets? Evidence from Health Insurance Auto-Enrollment,” *American Economic Review*, Vol. 115, pp. 772–822, DOI: <http://dx.doi.org/10.1257/aer.20231133>.
- Vytlacil, Edward (2002) “Independence, Monotonicity, and Latent Index Models: An Equivalence Result,” *Econometrica*, Vol. 70, pp. 331–341.


A Sorting with clinical quality differences

Figure A.1: Patient self-selection into private care (with clinical quality differences)


(a) Positive sorting on gains, lower clinical quality in the private sector




(b) Negative sorting on gains, lower clinical quality in the private sector



(c) Positive sorting on gains, higher clinical quality in the private sector



(d) Negative sorting on gains, higher clinical quality in the private sector



B Derivation of the MTR functions

Note that

$$\begin{aligned} E[Y = 0, D = 1 \mid X = x, P = p] &= P(Y = 0, D = 1 \mid X = x, P = p) \\ &= P(Y = 0 \mid D = 1, X = x, P = p)P(D = 1 \mid X = x, P = p) \end{aligned}$$

since $D = 1$, we have that $P(Z_i) \geq U_{di}$ and $U_{di} \sim Unif[0, 1]$ conditional on X and P . Therefore, the treated population are those with $U_{di} \in [0, p]$.

$$\begin{aligned} P(Y = 0, D = 1 \mid X = x, P = p) &= \int_0^p P(Y = 0 \mid X = x, U_d = u, D = 1) du \\ &= \int_0^p P(Y_1^* \leq 0 \mid X = x, U_d = u) du \end{aligned}$$

Apply Leibniz Rule, then

$$\frac{\partial E[Y = 0, D = 1 \mid X = x, P = p]}{\partial p} = P(Y_1^* \leq 0 \mid X = x, U_d = p)$$

Now, note that

$$\begin{aligned} E[Y = 0, D = 0 \mid X = x, P = p] &= P(Y = 0, D = 0 \mid X = x, P = p) \\ &= \int_p^1 P(Y = 0 \mid X = x, U_d = u, D = 0) du \\ &= \int_0^1 P(Y_0^* \leq 0 \mid X = x, U_d = u) du \\ &\quad - \int_0^p P(Y_0^* \leq 0 \mid X = x, U_d = u) du \end{aligned}$$

Apply Leibniz Rule

$$\frac{\partial E[Y = 0, D = 0 \mid X = x, P = p]}{\partial p} = -P(Y_0^* \leq 0 \mid X = x, U_d = p)$$