Easy Like Sunday Morning: Organizational Design and Physician Agency in Birth Timing*

Jimena S. Ferraro† Raul A. Sosa Franco M. Vazquez

Universidad de San Andrés

July 2026

Abstract

Cesarean rates differ sharply between private and public hospitals in Buenos Aires, even though payments do not vary with delivery mode in either sector. We argue this gap reflects organizational design rather than procedure prices. In the Buenos Aires metropolitan area, individualized private and team-based public maternity care coexist within the same labor market. Using ex ante preferences and medical records for low-risk first-time mothers, we document a preference margin and a timing margin. Stated preferences for cesarean predict scheduled procedures strongly in private hospitals and not at all in public ones. Elective cesareans in private hospitals cluster on boundary working days (the Monday or Friday adjacent to a weekend or holiday), and the intrapartum cesarean rate is about 13 percentage points higher on working than nonworking days in private hospitals, but indistinguishable from zero in public ones, a gradient that reflects selection induced by the timing margin. A simple model rationalizes both margins through the opportunity cost of physician time, internalized only when physicians are residual claimants of their own schedule. Organizational design, not procedure prices, is the primary lever behind the discretionary component of cesarean use.

JEL Codes: I11, I18, L23, J22, J44

Keywords: cesarean section, physician agency, organizational design, birth timing, maternity care, mixed health systems.

*This paper incorporates empirical work from an earlier unpublished working paper, "Cesarean Sections in Buenos Aires: Influence of Physician Convenience and Maternal Preferences", co-authored by Jimena S. Ferraro, Shagun Khare, and Alan Acosta. We thank Shagun and Alan for their contributions to the empirical findings.

†Corresponding author. Universidad de San Andrés. Email: ferraroj@udesa.edu.ar.

1 Introduction

In the Buenos Aires metropolitan area, cesareans account for more than 70 percent of births in private hospitals and around 43 percent in public hospitals (Figure 1).[1] This gap of close to 30 percentage points emerges despite the two sectors sharing the same metropolitan labor market, overlapping pools of obstetricians, and payment structures that do not differentiate cesarean from vaginal delivery within either sector. It is difficult to reconcile with the canonical economic explanations of excess cesarean use: price-driven supplier behavior and autonomous patient demand. Nor is this pattern unique to Argentina. Across low- and middle-income countries, cesareans are 1.6 times more frequent in private than in public facilities (Boerma et al., 2018). The gap's size and persistence, in Argentina and beyond, lead us to a possible additional margin in how obstetric care is organized.

The stakes are substantial. Cesarean delivery is life-saving when clinically indicated, but at the population level rates above roughly 10–15 percent are not associated with further reductions in maternal or neonatal mortality (World Health Organization, 2015). Unnecessary cesareans consume more resources than uncomplicated vaginal delivery (Gibbons et al., 2010), elevate short- and long-term risks for mothers and infants (Sandall et al., 2018), and reduce subsequent fertility (Halla et al., 2020). Latin America bears a disproportionate share of this excess, with cesarean rates among the highest in the world and a concentration of unnecessary procedures in the private sector (Belizán et al., 1999; Villar et al., 2006).

Economic research has identified two main channels through which nonmedical factors shape cesarean use. The first operates through prices. Obstetric care exemplifies physician agency: clinicians hold superior information about clinical risk and substantial control over treatment decisions (Arrow, 1963; Green, 1978; McGuire, 2000). Cesarean rates respond to procedure-specific fee differentials (Gruber et al., 1999; Foo et al., 2017) and to malpractice exposure, which induces defensive practices (Kessler and McClellan, 1996; Tussing and Wojtowycz, 1997; Currie and MacLeod, 2008; Frakes and Gruber, 2019). Yet price differentials are not necessary for supply-side incentives: de Elejalde and Giolito (2021) show that a Chilean copayment reduction paying private hospitals a flat fee per delivery still raised cesarean rates by 8.6 percentage points, and rationalize this through a demand-smoothing mechanism in which scheduled cesareans raise total deliveries per unit of capacity.

[1] Private-sector rates are from Librandi et al. (2023), who report 69.0 to 73.6 percent in IOMA-affiliated private facilities in the Province of Buenos Aires over 2017–2021. Public-sector rates are from the Sistema Informático Perinatal para la Gestión (SIP-G), Ministerio de Salud de la Nación Argentina, Indicadores básicos sobre atención perinatal a partir del SIP-G, Argentina, 2022. Within the same insured population, Ariovich et al. (2025) report 74.5 percent in private and 53.1 percent in public facilities over 2020–2023.

Figure 1. Cesarean section rates in Argentina

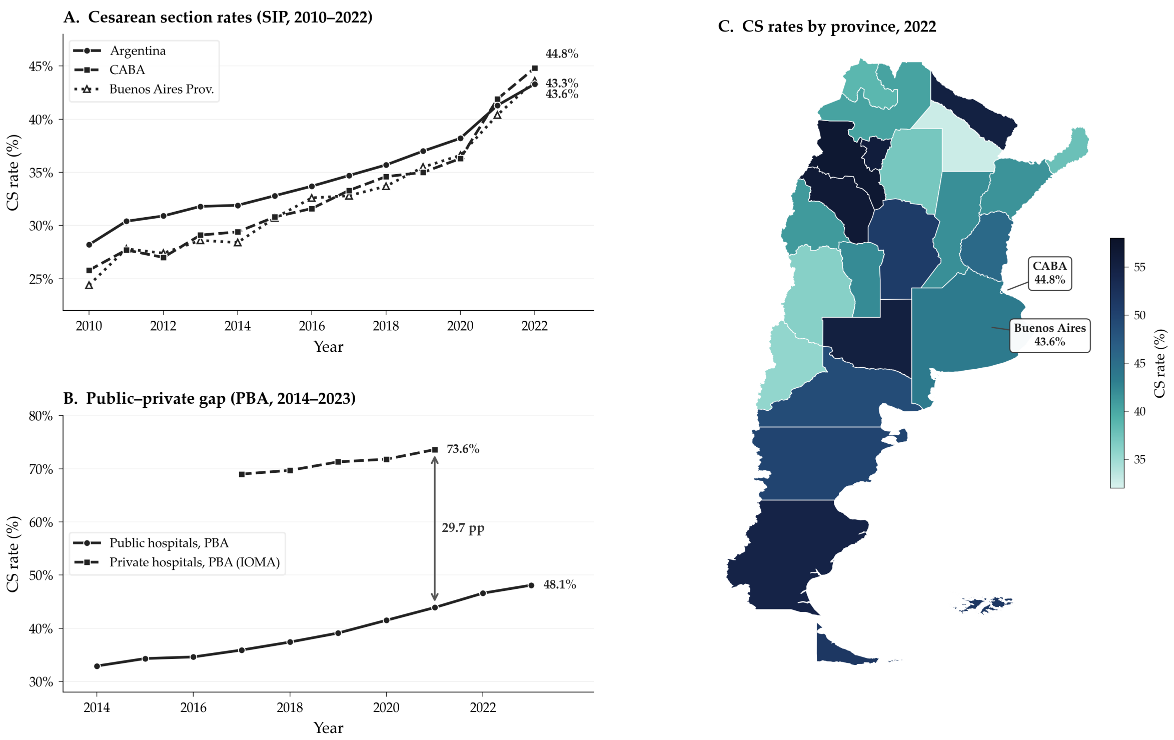

Notes: Panel A reports cesarean section rates from the Sistema Informático Perinatal para la Gestión (SIP-G) for Argentina, the City of Buenos Aires (CABA), and the Province of Buenos Aires (PBA), 2010–2022. Panel B compares public-hospital rates from SIP-G with private-hospital rates from the Instituto de Obra Médico Asistencial (IOMA) for PBA, 2014–2023. The vertical arrow indicates the gap of 29.7 percentage points in 2021. Panel C maps provincial cesarean rates from SIP-G for 2022.

The second channel is the value physicians place on schedule control. Because vaginal deliveries are typically longer and less predictable than cesareans, time-related incentives can tilt clinical choices toward cesarean delivery, particularly at evening hours and near the end of the work week (Brown, 1996; Spetz et al., 2001). Recent quasi-experimental work quantifies downstream consequences of timing-driven cesarean use for neonatal outcomes, subsequent fertility, and long-run child health (Costa-Ramón et al., 2018; Borra et al., 2019; Halla et al., 2020; Costa-Ramón et al., 2022).

A complementary literature focuses on patient preferences. Globally, only a minority of women prefer cesarean delivery (Mazzoni et al., 2011). Yet direct evidence on how preferences shape clinical decisions is limited: preferences are rarely observed before delivery. Most existing work therefore relies on indirect proxies: immigrant mothers' home-country cesarean norms (Grytten et al., 2013), manipulation of birth timing around salient dates (Dickert-Conlin and Chandra, 1999; Schulkind and Shapiro, 2014; Gans and Leigh, 2009), and cultural beliefs about auspicious days of delivery (Lo, 2003).

The most revealing evidence comes from Brazil, where public and private patients express broadly similar preferences for vaginal birth yet face dramatically different cesarean rates (Hopkins, 2000; Potter et al., 2001). Béhague et al. (2002) nuances the finding by showing that some lower-income women actively seek cesareans to avoid perceived poor-quality public labor care, underscoring that the public–private gap reflects institutional context rather than autonomous demand. What remains open is whether, and how, organizational context itself mediates the translation of preferences into clinical decisions.

In practice, prices and organization typically covary: fee-for-service reimbursement comes bundled with individualized continuity of care, while salaried pay comes bundled with team-based production. The organizational component of physician discretion has therefore rarely been isolated from the price component. This paper investigates that organizational component in the Buenos Aires metropolitan area, where two sharply different organizational forms coexist within a common labor market. Public maternity care is delivered by teams working fixed shifts: a pregnant woman is not followed by a specific physician during pregnancy, care at admission is provided by the team on duty, and clinicians are paid salaries tied to hours worked. Private maternity care, by contrast, is organized around a designated physician who, as the residual claimant of her own schedule, follows the woman through pregnancy and attends her delivery.

Three features of this setting isolate the organizational channel. First, within each sector, monetary payments do not differ by delivery mode: public salaries depend on hours worked, while private physicians are paid per delivery without any fee differential between cesarean and vaginal birth. This shuts down the canonical price-based channel of supplier-induced demand. Second, both sectors operate within the same metropolitan obstetric labor market and draw from a common pool of obstetricians. The cross-sector contrast is therefore between organizational forms acting on a shared physician population, not across separate labor markets or training systems. Third, the sectors differ sharply in how clinical authority is organized: continuous individualized care on one side and shift-based team production on the other. Holding procedure prices fixed and varying organizational form isolates the component of physician agency that depends on how clinical authority and schedule control are structured, rather than on how they are priced.

The closest prior studies are Triunfo and Rossi (2009), who documents a roughly doubled cesarean probability in private practice in Uruguay but conflates remuneration structure with organizational form, and de Elejalde and Giolito (2021), who hold prices flat in Chile but locate the incentive at the hospital level—capacity utilization—rather than in the organization of clinical authority. Neither observes ex ante preferences nor the scheduled/intrapartum distinction.

Our data advantage is that we observe, for the same low-risk first-time mothers, both the stated preference over delivery mode elicited before delivery and a medical record that distinguishes scheduled (elective) from intrapartum cesareans. The data come from a prospective cohort study of nulliparous women enrolled at five Buenos Aires hospitals (two public and three private) between August 2010 and December 2011 (Mazzoni et al., 2016), with an analytic sample of 379 deliveries. Two design features are critical. First, ex ante measurement of preferences rules out justification bias, the concern that reported preferences adjust to match realized outcomes (Bound, 1991). Second, the elective–intrapartum distinction separates the decision to schedule a procedure before labor, where timing is chosen, from the decision to escalate once labor has begun. Together, these features allow us to decompose discretionary cesarean use into a *preference margin* (whether physicians accommodate a stated preference for cesarean) and a *timing margin* (whether physicians shift deliveries away from weekends and public holidays). Previous work has examined these margins separately: studies of preferences use proxies but rarely observe timing, while studies of timing exploit calendar variation but rarely observe preferences.

We structure the empirical analysis with a simple model of physician decision-making. Each physician incurs a nonmonetary cost of performing a non-indicated cesarean, plus an additional cost of overriding a maternal preference for vaginal birth. The physician weighs these costs against the time saved by cesarean relative to vaginal delivery, with weekend and holiday time carrying a higher opportunity cost than working-day time. The key institutional primitive is how strongly physicians value time savings. In the private sector, the physician captures the full value of the time saved. In the public sector, fixed shifts and team-based production attenuate the private return to time saved on any single case.

The model yields three testable implications. First, within a sector, stated preferences translate into elective cesareans through the preference-responsive subset of physicians. The empirical asymmetry we document across sectors reflects both a higher intervention threshold in the private sector and the sorting of low-professionalism-cost physicians into it. Our estimates capture the combined effect of these two channels, which the design does not separate. Second, elective cesareans should bunch on boundary working days (the Monday or Friday adjacent to a weekend or holiday) more strongly in the private sector, where physicians who internalize the saved time stand to gain most from shifting an anticipated weekend or holiday onset onto a boundary day. Third, among women who enter labor, intrapartum cesareans should be more common on working days than on nonworking days in the private sector. This is a selection effect of the timing margin: discretionary cases with anticipated weekend or holiday onset are shifted to scheduled boundary-day cesareans rather than performed intrapartum, leaving the weekend intrapartum pool disproportionately concentrated in indicated escalations, while working-day labors retain their discretionary cesareans.

The empirical evidence supports all three predictions. On the preference margin, a stated preference for cesarean raises the probability of an elective cesarean by about 37 percentage points in private hospitals (from 11 to 48 percent), while having no detectable effect in public hospitals, where the corresponding probabilities are 5 and 6 percent. The private-sector differential survives a sensitivity analysis for selection on unobservables (Altonji et al., 2005; Oster, 2019): unobservables would have to explain more than the full set of observed covariates to overturn the estimate, and augmenting the model with a direct proxy for patient agency would increase the coefficient rather than attenuate it.

On the timing margin, elective cesareans in private hospitals concentrate on Mondays and Fridays and are almost absent from weekends and public holidays, while public hospitals show no comparable Mon/Fri concentration. Among women who enter labor, the intrapartum cesarean rate in private hospitals is about 13 percentage points higher on working days than on nonworking days (35.7 versus 22.4 percent), while the public sector shows a small gradient (28 versus 36 percent) that is statistically indistinguishable from zero. We interpret this calendar gradient as a selection effect of the timing margin: in the private sector, discretionary cesareans operate primarily through advance scheduling rather than through intrapartum conversions.

Our study makes three contributions. First, for the literature on physician agency, we isolate organizational design as a lever distinct from procedure prices. Prior work has identified the role of information asymmetry using patients who are themselves physicians (Johnson and Rehavi, 2016) and the role of reimbursement using exogenous variation in prices (Clemens and Gottlieb, 2014; Foo et al., 2017). Holding procedure prices fixed and varying the organization of clinical authority isolates the total effect of organizational form, a contribution that neither approach captures. The cross-sector difference in how physicians are paid, salaried shift work versus fee-per-delivery, is itself a component of organizational design rather than a separate price channel: it shapes whether a physician personally captures the time a faster delivery saves, which is the mechanism we study. This total effect operates through two complementary channels that are both endogenous to organizational design: the direct change in physician incentives and the sorting of physicians induced by differential rents. The reduced-form estimate is therefore the policy-relevant object even if the two channels cannot be separately identified.

Second, for the economics of cesarean delivery, we jointly test the preference and timing margins through the elective–intrapartum decomposition, guided by a model that links both to the organization of physician time. This complements evidence on hospital-level practice variation (Card et al., 2023), on the downstream consequences of timing manipulation for neonatal and long-run outcomes (Costa-Ramón et al., 2018; Halla et al., 2020; Costa-Ramón et al., 2022), and on public–private remuneration differences in Latin America (Triunfo and Rossi, 2009). Third, for the design of mixed public–private health systems, our findings suggest that interventions targeting procedure prices can miss the dominant distortion when the price channel is already inactive. Policies that dilute the extent to which any single physician is the residual claimant of her schedule, such as shared-call arrangements, group practices, and hospitalist-style coverage (Wachter and Goldman, 1996), are candidate levers to attenuate discretionary cesarean use without restricting clinically indicated surgery (Betrán et al., 2018).

The remainder of the paper proceeds as follows. Section 2 describes the Argentine health system and the organization of maternity care. Section 3 presents the model. Sections 4 and 5 describe the data and empirical strategy. Section 6 reports the results. Section 7 concludes.

2 Background

The cross-sectoral cesarean gap documented in the introduction reflects institutional differences in how maternity care is organized in Argentina. Subsection 2.1 describes the three-tier Argentine health system, which combines public, social-security, and private sectors. Subsection 2.2 then characterizes the organization of maternity care, focusing on the features of physician compensation, continuity of care, and clinical authority that feed into the model of physician discretion in Section 3.

2.1 Health system in Argentina

Argentina has a mixed health system composed of three sectors: the public sector, the social security sector, and the private sector. Public hospitals provide care free of charge and are accessible to the entire population, although in practice they are the main source of care for individuals lacking formal employment or private insurance. Public facilities are financed by the Ministry of Health and by reimbursements from social-security and private insurers when their beneficiaries use public services.

The social security sector is funded through social insurance. Employers and employees in the formal labor market are required to contribute to funds that cover workers and their families, as well as retirees. Because eligibility is tied to formal employment, social insurance is not universal. Social security funds act primarily as paying agents, purchasing health services largely from private providers[2]. The Superintendency of Health, an autonomous entity that reports to the Ministry of Health, regulates both social security funds and private insurers.

The private sector comprises private hospitals and other providers, along with

[2] A special case is the Programa de Atención Médico Integral (PAMI), which covers retired individuals and contracts mainly with private providers, although part of its demand is also directed to the public sector.

private insurance agencies and prepaid health plans. Prepaid plans are financed mostly through premiums paid by families and, in some cases, employers. In practice, private hospitals primarily serve individuals with social-security or private-insurance coverage, and are associated with higher socioeconomic status and more personalized care.

2.2 Maternity care

In public hospitals, maternity care is delivered by teams working fixed shifts. Pregnant women are not typically followed by a specific physician or midwife. During pregnancy, a woman may be seen by different clinicians across visits, and she generally cannot choose the provider responsible for follow-up examinations. When labor begins and she is admitted, care is provided by the physicians and midwives on duty at that time. If labor extends beyond a shift change, responsibility transfers to the incoming team. Decisions about how to proceed are made within the team. Compensation is also shift-based: clinicians are paid a fixed amount that depends on the hours worked, which limits any individual incentive to be present at the moment of delivery.

Private hospitals are organized around individualized continuity of care. Women choose a particular physician or midwife during prenatal care, and this clinician typically provides follow-up care and attends the delivery. The resulting continuity can strengthen trust and increase maternal confidence at birth because the responsible provider is known in advance, and patients often expect their chosen clinician to be present. It can also help clinicians anticipate upcoming demand based on their panel of patients. At the same time, individualized responsibility can increase mothers' reliance on a specific clinician and make delivery timing a binding constraint for providers: physicians may face periods of unexpectedly high workload, and coordinating availability is more difficult when the same clinician is expected to attend each delivery. Private physicians are also paid per delivery, which increases their dependence on patient demand and reinforces incentives to accommodate expectations about personalized care. Finally, individualized responsibility may raise perceived legal exposure (Jena et al., 2011).

These institutional differences imply distinct incentive environments. In private hospitals, continuity of care and individualized responsibility increase the extent to which clinicians internalize the time costs of unpredictable labor and can benefit from predictability in delivery timing. In public hospitals, team-based care and shift work limit individual control and attenuate the private return to convenience. Section 3 formalizes these ideas parsimoniously, letting the sector difference operate through the value physicians place on time saved while holding the cost of overriding a patient's stated preference common across sectors. Under this structure, the stronger preference responsiveness in private hospitals reflects the higher return to time and the sorting of physicians across sectors, and a distinct incentive to accommodate patient demand would reinforce this pattern rather than substitute for it.

3 Model

This section develops a model of delivery mode choice that links calendar-based convenience incentives to institutional structure and maternal preferences. The model clarifies when providers have incentives to shift deliveries away from weekends and public holidays, and why the strength of those incentives should differ across public and private hospitals in Buenos Aires. The model builds on three primitives: cesarean delivery is faster and more predictable than vaginal delivery, physicians bear professionalism costs for performing non-indicated procedures and for overriding maternal preferences, and institutions shape how easily time saved on a delivery converts into leisure or other valued activity. Monetary reimbursements are assumed not to vary with delivery mode within a sector, so the mechanism operates through time and nonclinical costs.[3] Appendix Table A1 summarizes the variables and parameters introduced in this section.

3.1 Setup and primitives

Consider a pregnancy that reaches delivery. Let d denote the *biological onset day*, defined as the day on which labor would begin absent discretionary intervention that changes timing. Let D denote the *delivery day*. Without discretionary intervention, $D = d$. With intervention, the physician may alter timing so that $D \neq d$. We assume that, when choosing a delivery plan, the physician has sufficient late-pregnancy information to anticipate whether onset is likely to occur during nonworking time, so that scheduling decisions are feasible on the margin of discretion.

To capture convenience incentives, we map the calendar into three categories $t \in \{0, 1, 2\}$. Category $t = 0$ denotes regular working days that are not adjacent to weekends or holidays, which in a standard week corresponds to Tuesday, Wednesday, and Thursday. Category $t = 1$ denotes boundary working days that are adjacent to weekends or holidays, which in a standard week corresponds to Friday and Monday[4].

[3] If cesarean delivery carried a direct monetary premium over vaginal delivery, the net gain from intervention would rise uniformly across physician types and calendar states, amplifying P1–P3 without altering their signs. Shutting down this channel by assumption isolates the organizational-design mechanism we want to identify, and the price differential is in any case absent in the Buenos Aires setting (Section 2). The assumption concerns nominal fees by mode; if the per-delivery payment requires the physician's personal attendance, predictable scheduling could additionally raise the chance of capturing it. This monetary channel would reinforce the time-cost incentive in the same direction and is itself part of the fee-per-delivery structure of private practice rather than a procedure-price differential.

[4] More generally this category can include working days immediately before or after a public holiday.

Category $t = 2$ denotes nonworking days, namely weekends and public holidays. From the physician's perspective these are the *inconvenient* days, since they carry the higher opportunity cost α. The working days ($t \in \{0, 1\}$) are correspondingly *convenient*. Let $t(\cdot)$ denote this classification rule, so that $t(d)$ is the category of the biological onset day and $t(D)$ is the category of the delivery day.

Physicians face a higher opportunity cost of working on nonworking days. We capture this with a multiplier $\alpha > 1$ applied to time spent on days with $t = 2$. Define

$$\omega_t \equiv \begin{cases} 1, & t \in \{0, 1\}, \\ \alpha, & t = 2, \end{cases}$$

so that an hour of work on a nonworking day carries weight α relative to an hour on a working day. This is consistent with Brown (1996), who documents that physician demand for leisure translates into fewer cesareans on weekends and higher cesarean rates during evenings and near the end of the work week.

Discretionary changes in timing are assumed to be local. If $t(d) = 2$, the physician can shift delivery to an adjacent boundary working day, so that $t(D) = 1$. This captures the idea that a weekend or holiday delivery can sometimes be anticipated by scheduling a procedure on the preceding working day or, in some cases, postponed to the following working day, but that clinically meaningful rescheduling is limited in scope.[5] When $t(d) \in \{0, 1\}$, no shift across calendar categories is advantageous: moving delivery within working-day categories leaves the time weight ω_t unchanged, while moving to a nonworking day strictly raises the time cost. So no shift is strictly beneficial, and we set $D = d$, resolving the physician's indifference across working-day dates in favor of the onset day. The delivery day coincides with the onset day, but the physician may still perform a cesarean delivery that is not medically indicated. In the data, such discretionary cesareans can appear either as an elective (scheduled) procedure performed before labor begins (with $D = d$) or as an intrapartum conversion performed after labor onset (also with $D = d$).

Patients are of two broad clinical types. A fraction $\xi \in (0, 1)$ of deliveries have a clear medical indication for cesarean section (CS), so vaginal delivery is not feasible. These cases are not the focus of the nonmedical decision problem because the physician has no discretion over the mode of delivery. The analysis therefore centers on the remaining share $1 - \xi$ of *discretionary* cases.

[5] Both boundary days provide feasible elective-shift targets when the physician anticipates a weekend or holiday onset at a prenatal visit. For Friday, the procedure is scheduled in advance of the nonworking period. For Monday, it is scheduled for the day after. The evidence shows clustering on both days for elective cesareans in the private sector, consistent with both shift mechanisms operating in practice. The Monday channel can additionally reflect conservative intrapartum management of a spontaneous weekend labor completed on a working day, a possibility we return to when interpreting the intrapartum gradient.

Within discretionary cases, mothers are heterogeneous in their preferred mode of delivery. Each mother is of type $\tau \in \{C, N\}$, where $\tau = C$ indicates a preference for CS and $\tau = N$ indicates a preference for vaginal delivery. Let $\theta \equiv \Pr(\tau = C)$ denote the share of mothers who prefer CS.

Delivery modes differ in time requirements. Let L_N denote the time required for a vaginal delivery and L_C the time required for a CS. We assume $L_N > L_C$, consistent with the fact that vaginal deliveries are typically longer and less predictable than surgery.

Physicians differ in their costs of performing a non-indicated CS. A physician incurs a *professionalism cost* $\phi \geq 0$ when she performs a non-medically indicated CS. This parameter summarizes adherence to clinical guidelines and the disutility from deviating from medical appropriateness. Let ϕ be distributed according to a continuous distribution F on $[0, \bar{\phi}]$. If the mother prefers vaginal delivery ($\tau = N$) and the physician performs a CS, the physician incurs an additional *override cost* $\gamma > 0$, capturing the friction of acting against the mother's preference and the effort required to justify the decision.

Institutions determine how strongly time savings enter physician utility. Let $s \in \{\text{priv}, \text{pub}\}$ index the sector. A physician in sector s values time at rate $\rho_s > 0$, and we assume $\rho_{\text{priv}} > \rho_{\text{pub}}$. This inequality reflects that private physicians are closer to residual claimants of their schedule, so an hour saved at delivery can be converted into leisure or other valued activity, while public physicians typically remain on duty for a fixed shift and therefore cannot fully appropriate time saved on any single patient.

3.2 Per-case utility and feasible actions

We assume that within a sector monetary payments do not differ across modes of delivery, so the physician's decision trades off time costs against nonclinical costs. The physician's per-case utility from action $a \in \{N, C\}$ performed on a day of type t is

$$U^s(a, t; \tau, \phi) = -\rho_s \, \omega_t \, L_a - \mathbb{1}\{a = C\}\Big(\phi + \mathbb{1}\{\tau = N\}\gamma\Big), \tag{1}$$

where L_N and L_C are the time requirements, ω_t is the nonworking-day multiplier, and $\mathbb{1}\{\cdot\}$ is the indicator function.[6]

Equation (1) embeds the two features that separate our framework from the demand-smoothing model of de Elejalde and Giolito (2021), the closest antecedent in which cesarean use rises without procedure-price differentials. First, the decision unit is the individual physician rather than the hospital: in their model, scheduled cesareans are valuable because random vaginal arrivals would otherwise idle capacity, so a higher

[6] Since the model restricts attention to discretionary cases, the professionalism cost ϕ is incurred every time a cesarean is performed under this utility function.

cesarean share raises total deliveries per unit of capacity; here, the value of schedulability accrues to the physician through $\rho_s \, \omega_t$, and organizational design determines how much of the saved time she captures. Second, the mother is an agent rather than a bystander: de Elejalde and Giolito (2021) assume women are indifferent between modes of delivery, whereas here the stated preference τ enters the physician's payoff directly through the override cost γ. The preference margin and the calendar margin derived below follow from these two ingredients and would not arise from hospital-level smoothing alone.

For a given biological onset day d, the physician compares feasible actions subject to the local shifting constraint. It is convenient to express choices as net utility gains relative to the baseline of vaginal delivery on the biological onset day, $U^s(N, t(d); \tau, \phi)$, which removes terms common to all actions for a given d.

When onset would occur on a working day, $t(d) \in \{0, 1\}$, shifting across calendar categories is not beneficial (as discussed in Section 3.1), so the physician chooses between vaginal delivery and CS on that day. Let subscript w denote working-day onset and subscript 2 denote nonworking-day onset. The net utility gain from choosing CS rather than vaginal delivery is

$$\begin{aligned} \Gamma^s_w(\tau, \phi) &\equiv U^s(C, t(d); \tau, \phi) - U^s(N, t(d); \tau, \phi) \\ &= \rho_s(L_N - L_C) - \phi - \mathbb{1}\{\tau = N\}\gamma. \end{aligned} \tag{2}$$

When onset would occur on a nonworking day, $t(d) = 2$, the physician can either deliver vaginally on that day or shift delivery to an adjacent boundary working day through a scheduled CS. Relative to the baseline of vaginal delivery on the nonworking day, the net gains from performing CS on the nonworking day and from shifting CS to a boundary working day are

$$\widetilde{\Gamma}^s_{C@2}(\tau, \phi) \equiv U^s(C, 2; \tau, \phi) - U^s(N, 2; \tau, \phi) = \rho_s \alpha(L_N - L_C) - \phi - \mathbb{1}\{\tau = N\}\gamma, \tag{3}$$

$$\widetilde{\Gamma}^s_{\text{Shift}}(\tau, \phi) \equiv U^s(C, 1; \tau, \phi) - U^s(N, 2; \tau, \phi) = \rho_s(\alpha L_N - L_C) - \phi - \mathbb{1}\{\tau = N\}\gamma. \tag{4}$$

Lemma 1. *For any $\rho_s > 0$, performing a non-indicated CS on a nonworking day is strictly dominated by shifting the CS to a boundary working day whenever shifting is feasible.*

Proof. Subtracting (3) from (4) yields

$$\widetilde{\Gamma}^s_{\text{Shift}}(\tau, \phi) - \widetilde{\Gamma}^s_{C@2}(\tau, \phi) = \rho_s(\alpha - 1)L_C,$$

which is strictly positive because $\rho_s > 0$, $\alpha > 1$, and $L_C > 0$. □

Lemma 1 implies that when $t(d) = 2$ the relevant comparison is between vaginal delivery on the nonworking day and a shifted CS on a boundary working day, with the net gain given by (4).

Now define the sector-specific time incentives

$$S_w^s \equiv \rho_s(L_N - L_C), \qquad S_2^s \equiv \rho_s(\alpha L_N - L_C).$$

Because $\alpha > 1$, we have

$$S_2^s - S_w^s = \rho_s(\alpha - 1)L_N > 0. \tag{5}$$

Thus, for a given physician and patient, the incentive to intervene is stronger when delivery would otherwise occur on a nonworking day. In practice, timing can also be influenced through induction of labor, another physician-controlled technology. We exclude it from the baseline action set because induction does not reliably control the delivery day: an induced labor still yields a variable-length vaginal delivery that may not complete on the intended working day and may itself end in an intrapartum cesarean, so it is not a clean substitute for a scheduled cesarean as a day-shifting technology. The empirical analysis evaluates robustness to excluding induced labors (Section 6.2).

3.3 Optimal policy and implications for observed outcomes

The physician chooses the delivery plan that maximizes per-case utility. Because payments are assumed not to vary with delivery mode within a sector, the discretionary decision reduces to a comparison between the time-related benefit of cesarean delivery and the nonclinical costs of performing a non-indicated procedure and, when relevant, overriding the mother's stated preference. Using the net gains defined in (2) and (4), intervention is optimal whenever the relevant net gain is positive.

It is convenient to collect the non-time components of the decision into a single term. Define the total non-time friction as

$$T(\tau, \phi) \equiv \phi + \mathbb{1}\{\tau = N\}\gamma.$$

On working-day onset, $t(d) \in \{0, 1\}$, shifting across calendar categories is not advantageous and a non-indicated cesarean is chosen if and only if $S_w^s > T(\tau, \phi)$. On nonworking-day onset, $t(d) = 2$, the relevant discretionary comparison is between vaginal delivery on the nonworking day and a scheduled cesarean shifted to a boundary working day, and a shifted cesarean is chosen if and only if $S_2^s > T(\tau, \phi)$. These rules admit a cutoff representation. Define

$$\hat{\phi}_w^s(\tau) = S_w^s - \mathbb{1}\{\tau = N\}\gamma, \qquad \hat{\phi}_2^s(\tau) = S_2^s - \mathbb{1}\{\tau = N\}\gamma,$$

so that intervention occurs if and only if $\phi < \hat{\phi}_w^s(\tau)$ on working-day onset and if and only if $\phi < \hat{\phi}_2^s(\tau)$ on nonworking-day onset. Each $\hat{\phi}_t^s(\tau)$ is the threshold marked by the dashed lines in Figure 2: a physician with ϕ below this threshold has a professionalism cost low enough that the time gain from intervening on a type-τ patient outweighs her disutility of performing a non-indicated cesarean, and therefore intervenes; a physician with ϕ above it declines.

Figure 2 summarizes these decision regions and highlights the role of institutional structure.

Figure 2. Policy regions by physician professionalism cost ϕ in public and private sectors

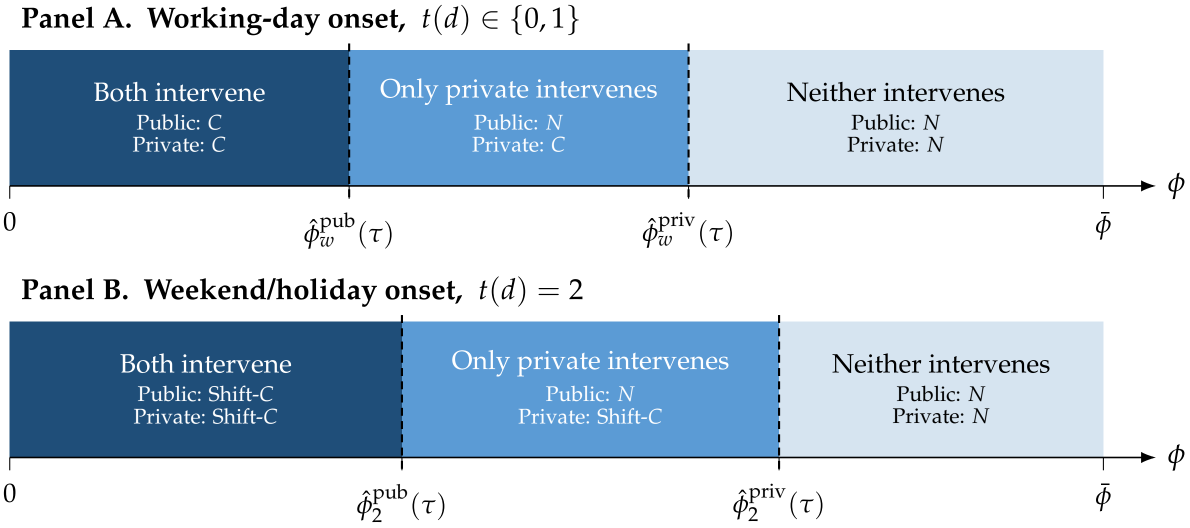

Notes: Each panel plots the physician's professionalism cost $\phi \in [0, \bar{\phi}]$ on the horizontal axis. Dashed lines indicate sector-specific cutoffs: $\hat{\phi}_w^s(\tau)$ for working-day onset (Panel A) and $\hat{\phi}_2^s(\tau)$ for weekend/holiday onset (Panel B). Dark shading marks the region where both sectors intervene, medium shading where only the private sector does, and light shading where neither does. Actions are N (vaginal delivery), C (cesarean on the onset day), and Shift-C (cesarean shifted to an adjacent boundary working day). Because $\rho_{\text{priv}} > \rho_{\text{pub}}$, the private cutoff lies to the right of the public cutoff in both panels, and the gap widens from Panel A to Panel B. The cutoffs are ex-ante thresholds applied at the planning stage, conditional on the anticipated onset category $t(d)$; ex-post deviations—for example, when spontaneous labor begins earlier than anticipated, before a boundary-day cesarean can be arranged—leave a residual of discretionary cases in the weekend intrapartum pool (see footnote to P3).

In Panel A (working-day onset), the private cutoff lies to the right of the public cutoff because $\rho_{\text{priv}} > \rho_{\text{pub}}$: a private provider finds a discretionary cesarean optimal for an intermediate range of physician types where a public provider does not. In Panel B (nonworking-day onset), both cutoffs lie further to the right because the higher opportunity cost of nonworking-day work can be avoided by scheduling a cesarean on an adjacent boundary working day, and the private cutoff shifts more than the public one because $\rho_{\text{priv}} > \rho_{\text{pub}}$ amplifies the weekend premium. This cutoff structure is the core mechanism behind the model's predictions for bunching of elective cesareans on boundary working days and for differential working-day versus nonworking-day patterns among women who enter labor.[7]

Lemma 2. *Under $\alpha > 1$ and $\rho_{\text{priv}} > \rho_{\text{pub}}$, the sector-specific net time incentives S_t^s satisfy:*

(i) $S_2^s > S_w^s$ *for each* $s \in \{\text{priv}, \text{pub}\}$.

(ii) $S_t^{\text{priv}} > S_t^{\text{pub}}$ *for each* $t \in \{w, 2\}$.

(iii) $S_2^{\text{priv}} - S_w^{\text{priv}} > S_2^{\text{pub}} - S_w^{\text{pub}}$.

Proof. (i) $S_2^s - S_w^s = \rho_s(\alpha - 1)L_N$, which is strictly positive. (ii) $S_w^{\text{priv}} - S_w^{\text{pub}} = (\rho_{\text{priv}} - \rho_{\text{pub}})(L_N - L_C)$ and $S_2^{\text{priv}} - S_2^{\text{pub}} = (\rho_{\text{priv}} - \rho_{\text{pub}})(\alpha L_N - L_C)$, both strictly positive. (iii) $(S_2^{\text{priv}} - S_w^{\text{priv}}) - (S_2^{\text{pub}} - S_w^{\text{pub}}) = (\rho_{\text{priv}} - \rho_{\text{pub}})(\alpha - 1)L_N$, strictly positive. □

These three properties translate into observable patterns. Part (i) is the source of the timing margin: since the time incentive is strictly higher on nonworking-day onset, a physician who would not perform a discretionary cesarean on a working-day onset may nonetheless schedule one when delivery would otherwise fall on a weekend or holiday. Part (ii) implies that the set of physicians who intervene is larger in the private sector: for any (τ, ϕ), intervention in the public sector implies intervention in the private sector, and physicians whose professionalism cost falls in the interval $[\hat{\phi}_t^{\text{pub}}(\tau), \hat{\phi}_t^{\text{priv}}(\tau)]$ intervene only in private. Part (iii) implies that the size of the timing margin itself is larger in the private sector: the difference $S_2^{\text{priv}} - S_w^{\text{priv}}$ exceeds $S_2^{\text{pub}} - S_w^{\text{pub}}$, so the gap in intervention propensity between nonworking-day and working-day onset is more pronounced in private practice.

Aggregating over physician heterogeneity clarifies how preferences and institutions interact. For a given sector s and calendar state $t \in \{w, 2\}$,

$$\Pr(\text{intervene} \mid s, t, \tau) = F(S_t^s - \mathbb{1}\{\tau = N\}\gamma)\,.$$

Averaging over maternal types yields

$$\Pr(\text{intervene} \mid s, t) = \theta\, F(S_t^s) + (1 - \theta)\, F(S_t^s - \gamma)\,. \tag{6}$$

Equation (6) highlights two channels through which institutions affect observed intervention rates. A higher ρ_s increases both S_w^s and S_2^s, shifting the arguments of $F(\cdot)$ upward and raising intervention for both preference types. The override cost γ lowers intervention among mothers who prefer vaginal delivery and implies that, in sectors where discretionary intervention occurs, measured preferences should be more predictive of realized delivery mode.

[7] The depicted ordering $\hat{\phi}_2^{\text{pub}} < \hat{\phi}_w^{\text{priv}}$ depends on the relative magnitudes of $\rho_{\text{priv}}/\rho_{\text{pub}}$ and α. The bunching and intrapartum predictions depend only on the private cutoff lying to the right of the public cutoff in each panel and on the cross-panel shift being larger in private.

This structure implies that elective cesareans should bunch on boundary working days (the Monday or Friday adjacent to a weekend or holiday) and that this bunching should be stronger in private hospitals. It also implies a selection effect in the composition of births that occur on nonworking days: discretionary cases with low ϕ are precisely those most likely to be shifted off weekends and holidays, so the remaining nonworking-day deliveries are disproportionately vaginal.

3.4 Testable implications

The decision rule derived above yields three testable implications that structure the empirical analysis in Section 5.

P1 (Preference margin). Within a sector, stated preferences for cesarean ($\tau = C$) raise the probability of an elective cesarean through the preference-responsive subset of physicians, those with $\phi \in [S_t^s - \gamma, S_t^s)$, who perform a non-indicated cesarean for a $\tau = C$ mother but not, after incurring the override cost γ, for a $\tau = N$ mother. Because the mother's preference is elicited before labor, a physician who honors it schedules the cesarean in advance rather than awaiting labor, so the preference margin operates through elective procedures.

P2 (Timing margin). Within working days, elective cesareans bunch on boundary days (the Monday or Friday adjacent to a weekend or holiday), and this bunching is more pronounced in the private sector. Lemma 1 implies that non-indicated cesareans scheduled on a nonworking day are dominated by the same cesarean shifted to an adjacent boundary working day, so a discretionary cesarean with anticipated nonworking-day onset is scheduled on those boundary days rather than on interior working days or on the nonworking day itself. Because $\rho_{\text{priv}} > \rho_{\text{pub}}$, the set of discretionary cases shifted off nonworking days is larger in the private sector, and the bunching on boundary days is therefore sharper.

P3 (Intrapartum calendar gradient via selection). Among women who enter labor, the intrapartum cesarean rate is higher on working days than on nonworking days in the private sector. The pattern reflects selection: by Lemma 1, discretionary cases with anticipated nonworking-day onset are shifted to scheduled boundary-day cesareans rather than performed intrapartum, so the nonworking-day intrapartum pool is concentrated in medically indicated escalations, the cases over which the physician has no discretion and which therefore fall outside the decision problem addressed by this framework.[8] The working-day intrapartum pool, by contrast, includes both indicated escalations and discretionary cesareans decided during labor on working-day onsets, cases for which no shift is needed because the anticipated onset already falls on a working day.

The three predictions have different theoretical structures. P2 and P3 follow from the inequality $\rho_{\text{priv}} > \rho_{\text{pub}}$ and Lemma 1: the timing channel varies across sectors by construction, because its strength scales with ρ_s. P1 operates through the override cost γ, which is invariant across sectors, so the prediction is stated within a sector. The empirical asymmetry we document across sectors is then rationalized jointly by the higher private-sector cutoffs and by sector selection, which we formalize in Section 3.5.

3.5 Two-stage model: endogenous sector choice

Up to this point we have held a physician's professionalism cost ϕ fixed and asked how she behaves in each sector. The implicit picture is that the same physician could practice in either sector and that only the incentive environment changes. In reality, physicians choose where to work, and that choice depends on the same trade-offs the model is designed to capture. If physicians with a low ϕ, those most willing to perform a cesarean for nonclinical reasons, find the private sector more rewarding, they will concentrate there, so the cross-sector contrast we observe reflects not only the incentives a given physician faces but also who chose each sector. The aim of this subsection is to write that selection problem down explicitly and to show that, under the same primitives that drive P1–P3, this sorting amplifies the within-sector predictions rather than confounding them.

Formally, we work with a two-stage timing. In Stage 1, each physician anticipates the incentive environment she would face in each sector, namely the sector-specific net incentives S_t^s derived in Section 3.2 for $t \in \{w, 2\}$, and chooses the sector in which her expected utility is highest. In Stage 2, conditional on the sector she has chosen, cases arrive with characteristics (τ, d) and she applies the optimal decision rule of Section 3.3. A physician's Stage 1 choice depends on three components. Her own type ϕ enters because higher-ϕ physicians extract less rent from non-indicated cesareans in either sector. The expected discretionary rents she would earn in each sector enter because they depend on the $\{S_t^s\}$ environment that sector offers. Finally, any other sector-specific differences in utility that do not depend on the discretionary decision, such as salaries, amenities, academic standing, or the time cost of a routine caseload, shift the calculus on top of the rent considerations. The next three paragraphs build these three components into a sector-choice rule.

[8] A small residual of discretionary cases may remain in the nonworking-day intrapartum pool when spontaneous labor begins earlier than the physician anticipated, before a boundary-day scheduled cesarean can be arranged.

For each case the physician sees, the realized utility gain from performing a non-indicated cesarean depends on the calendar state of onset t and on the mother's preference τ. Combining the time-saving incentive of Section 3.2 with the override cost of Section 3.3, the net gain takes the form

$$X^s_{t,\tau}(\phi) \;=\; S^s_t \;-\; \phi \;-\; \mathbb{1}\{\tau = N\}\,\gamma.$$

Three pieces enter $X^s_{t,\tau}(\phi)$. The first, S^s_t, is the gross sector-specific gain from saving time on a day of calendar state t, derived in Section 3.2. The second, ϕ, is the physician's own professionalism cost, the same cost that governs her behavior within any sector. The third, the override γ, applies only when the mother explicitly prefers a vaginal delivery ($\tau = N$). A risk-neutral physician performs the cesarean exactly when this net gain is positive, so her realized rent from a case of type (t, τ) is $[X^s_{t,\tau}(\phi)]^+ \equiv \max\{X^s_{t,\tau}(\phi), 0\}$, strictly positive only when she intervenes. At the Stage 1 decision point she does not yet know τ for the patients she will see, so she averages over the population share $\theta = \Pr(\tau = C)$ of mothers who prefer a cesarean. The expected rent on a day with onset of type t is then

$$R^s_t(\phi) \;\equiv\; \theta\,[X^s_{t,C}(\phi)]^+ \;+\; (1-\theta)\,[X^s_{t,N}(\phi)]^+,$$

the average per-case rent a physician of type ϕ would earn from a stream of patients arriving on day-type t in sector s.

The physician also has to average over the calendar state of onset itself, which she does not control, and over whether the case is one of the $1 - \xi$ discretionary or ξ medically indicated cases introduced in Section 3.1. Let $p_w \equiv \Pr(t(d) \in \{0, 1\})$ and $p_2 \equiv 1 - p_w$ denote the calendar probabilities. Indicated cases contribute zero discretionary rent in the baseline model because the physician has no choice over mode of delivery, so the ex ante per-delivery discretionary rent in sector s is the calendar-weighted per-day rent on discretionary cases, scaled by the discretionary share,

$$V^s(\phi) \;=\; (1-\xi)\,\big[\,p_w\,R^s_w(\phi) \;+\; p_2\,R^s_2(\phi)\,\big]. \tag{7}$$

$V^s(\phi)$ is the object that drives sector choice in Stage 1. It captures, before knowing the mother's preference, the calendar state of onset, or whether the case will turn out to be discretionary, how much discretionary rent a physician of type ϕ expects to extract on average from each patient she sees in sector s. Because $[X^s_{t,\tau}(\phi)]^+$ is weakly decreasing in ϕ for every (s, t, τ), so is $V^s(\phi)$. The more professional a physician is, the less rent she earns from discretion, and for a physician with ϕ above all four sector-by-onset thresholds the discretionary rent is zero in either sector.

Discretionary rents are not the whole story. Each sector also carries components of utility that do not vary with the discretionary decision, and hence not with ϕ, such as salaries, hospital amenities, academic standing, and the expected time cost of the physician's routine caseload. We collect these into a sector-specific constant B_s.[9] A physician chooses the private sector exactly when private-sector total utility is higher than public-sector total utility,

$$B_{\text{priv}} + V^{\text{priv}}(\phi) \;>\; B_{\text{pub}} + V^{\text{pub}}(\phi).$$

Summarizing the ϕ-independent side of this trade-off in a single compensating differential in favor of the public sector,

$$\kappa \;\equiv\; B_{\text{pub}} - B_{\text{priv}},$$

the choice rule becomes

$$\Delta V(\phi) \;\equiv\; V^{\text{priv}}(\phi) - V^{\text{pub}}(\phi) \;>\; \kappa, \tag{8}$$

that is, a physician picks the private sector when the rent gap she would earn there over the public sector more than covers the ϕ-independent advantage that public practice carries.

This selection problem has a clean solution. Intuitively, the higher-ρ_s private sector is attractive precisely to physicians who would actually use the time savings it makes more valuable. A physician with a very low ϕ does many non-indicated cesareans and is much better off doing them in the sector where each one saves her more time, so her rent differential $\Delta V(\phi)$ is large. A physician with a very high ϕ does no non-indicated cesareans in either sector, so her rent differential is zero and the two sectors look the same to her on the rent dimension. Between these extremes, $\Delta V(\phi)$ slopes downward in ϕ in a single-crossing way that yields a sharp sorting cutoff.

Proposition 1. *The net rent differential $\Delta V(\phi)$ is weakly decreasing in ϕ, and therefore satisfies single crossing. Consequently, for any $\kappa \geq 0$, the set of physician types who choose the private sector is an interval $[0, \phi^*)$, sorting is positively assortative (lower-ϕ types match to the higher-ρ sector), and the marginal physician ϕ^* solves $\Delta V(\phi^*) = \kappa$ whenever a solution exists.*

Proof. Fix a calendar state t and maternal type τ, and define $a \equiv S_t^{\text{priv}} - \mathbb{1}\{\tau = N\}\gamma$ and $b \equiv S_t^{\text{pub}} - \mathbb{1}\{\tau = N\}\gamma$, with $a > b$ because $S_t^{\text{priv}} > S_t^{\text{pub}}$ (Lemma 2, part (ii)). The

[9] Public hospitals often carry greater academic prestige and serve as teaching institutions, with residency programs and research opportunities less available in the private sector. B_s also nets out the expected time cost of the physician's routine caseload, which is sector-specific because ρ_s multiplies all delivery time and is therefore larger in the higher-ρ private sector. This lowers B_{priv} relative to B_{pub} and so contributes positively to κ. Being independent of ϕ, every such component shifts the threshold κ without affecting the shape of $\Delta V(\phi)$ or Proposition 1.

difference in node-level rents is

$$f(\phi) = \max\{a - \phi, 0\} - \max\{b - \phi, 0\}.$$

If $\phi \leq b$, then $f(\phi) = a - b$, constant. If $b < \phi < a$, then $f(\phi) = a - \phi$, strictly decreasing. If $\phi \geq a$, then $f(\phi) = 0$. Thus $f(\phi)$ is weakly decreasing. By Equation (7), $V^s(\phi)$ is a nonnegative weighted sum over (t, τ) of terms of the form $[X^s_{t,\tau}(\phi)]^+$. Because the weights, namely the calendar probabilities p_t, the preference share θ, and the discretionary share $1 - \xi$, are common to both sectors, $\Delta V(\phi) = V^{\text{priv}}(\phi) - V^{\text{pub}}(\phi)$ is the same weighted sum applied to the node-level differences $f(\phi)$, and is therefore weakly decreasing. Monotonicity implies a cutoff characterization of the choice set. $\square$

The cutoff ϕ^* has a clean reading. It is the type for whom the rent advantage of private practice exactly equals the ϕ-independent advantage of public practice. Physicians below it have a private rent advantage that outweighs the public premium and sort into private, while physicians above it remain public, so the matching is positively assortative, with the lowest-ϕ physicians in the highest-ρ sector.

Figure 3 illustrates the sorting problem. The blue curve plots $\Delta V(\phi)$ against ϕ. It starts at a positive plateau for small ϕ, where every physician intervenes in every state in both sectors. There the rent gap is a constant set by the threshold differences, because the physician's own costs enter her private and public rents identically and cancel in the difference. The curve then declines piecewise linearly through the sector-by-onset thresholds $\hat{\phi}^s_t(C)$ inherited from Figure 2 as the calendar-by-sector incentives switch off one by one.[10] Within each region the slope of $\Delta V(\phi)$ is proportional to minus the calendar weight of the onset node(s) switching off there, with relative slopes $-p_w$ (only the working-onset node declining), $-p_2$ (only the nonworking-onset node declining), and $-(p_w + p_2)$ (both declining). The two outer declining segments therefore have different slopes whenever $p_w \neq p_2$, the shallower being the one driven by the less likely onset category. The sector cutoff ϕ^* is the value of ϕ at which the curve crosses the horizontal line at κ.

The two panels of Figure 3 correspond to the two possible orderings of the interior thresholds $\hat{\phi}^{\text{pub}}_2(C)$ and $\hat{\phi}^{\text{priv}}_w(C)$, whose comparison is parameter-dependent. The ordering $\hat{\phi}^{\text{pub}}_2(C) \lesseqgtr \hat{\phi}^{\text{priv}}_w(C)$ holds according as the calendar amplification $(\alpha L_N - L_C)/(L_N - L_C)$ falls short of or exceeds the sector gap $\rho_{\text{priv}}/\rho_{\text{pub}}$. In Panel A the sector gap dominates, and the two interior thresholds bracket a steeper segment of slope $-(p_w + p_2)$ in which both onset nodes decline together. In Panel B the calendar effect dominates, and the same middle region is instead flat, because the working-onset node has already switched off in both sectors while the nonworking-onset node

[10] The $\tau = N$ thresholds $\hat{\phi}^s_t(N) = \hat{\phi}^s_t(C) - \gamma$ add four further kinks, each shifted left by the override cost γ relative to its $\tau = C$ counterpart. They are suppressed in Figure 3 for clarity.

still has both intervening, leaving the gap independent of ϕ. Either way $\Delta V(\phi)$ is weakly decreasing, so the cutoff characterization of Proposition 1 applies in both panels. Appendix Figure A2 complements the schematic with calibrated plots that vary γ, α, $L_N - L_C$, and the sector gap $\rho_{\text{priv}} - \rho_{\text{pub}}$, confirming that $\Delta V(\phi)$ remains weakly decreasing across parameter choices and that any horizontal line at κ is crossed at most once.

Figure 3. Sector choice and the cutoff $\phi^*(\kappa)$

Panel A. Sector gap dominates: $\hat{\phi}_2^{\text{pub}}(C) < \hat{\phi}_w^{\text{priv}}(C)$

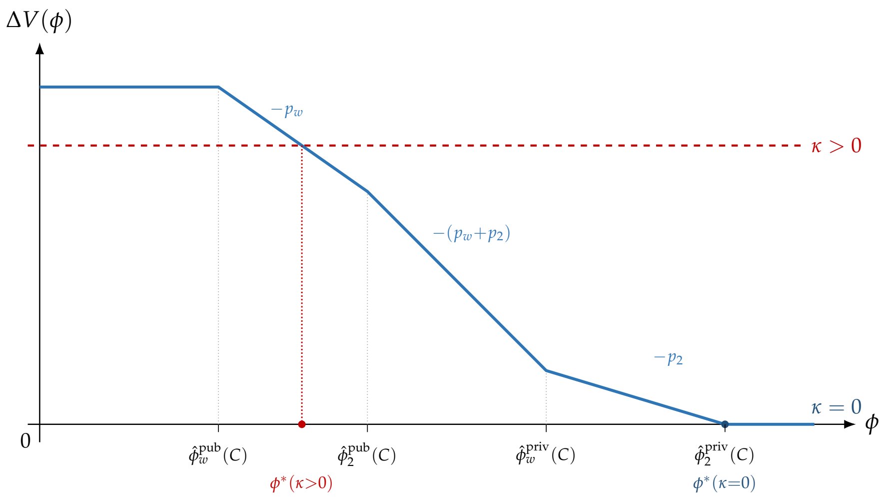

Panel B. Calendar effect dominates: $\hat{\phi}_w^{\text{priv}}(C) < \hat{\phi}_2^{\text{pub}}(C)$

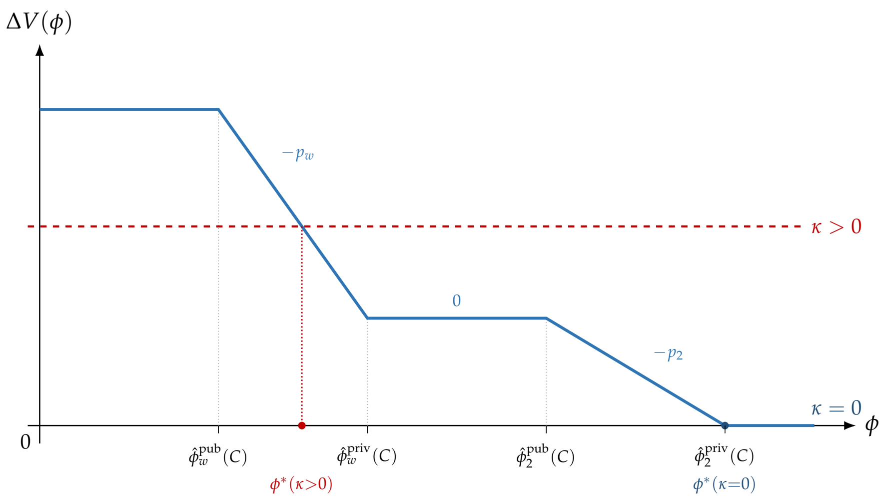

Notes: Schematic of the sector-choice problem in Section 3.5, drawn for $\tau = C$ so the thresholds coincide with S_t^s. The four $\tau = N$ kinks are suppressed, and the vertical scale is stylized and specific to each panel. The curve plots $\Delta V(\phi) = V^{\text{priv}}(\phi) - V^{\text{pub}}(\phi)$, with the slope of each segment given in the main text. Panels A and B are the two possible orderings of the interior thresholds $\hat{\phi}_2^{\text{pub}}(C)$ and $\hat{\phi}_w^{\text{priv}}(C)$. A physician chooses private when $\phi < \phi^*$, where $\Delta V(\phi^*) = \kappa$. The cutoff is $\hat{\phi}_2^{\text{priv}}(C)$ under $\kappa = 0$ (complete sorting) and lower under $\kappa > 0$ (partial sorting). Incorporating medically indicated cesareans (Section 3.6) shifts the curve upward by a constant, weakly raising ϕ^*.

The two panels vary the ordering of the thresholds, and hence the shape of the curve, whereas the value of κ fixes where the cutoff falls along whichever curve applies. Consider first the benchmark $\kappa = 0$. The compensating differential vanishes and the private-sector condition reduces to $\Delta V(\phi) > 0$. Because $\Delta V(\phi)$ is positive for every ϕ strictly below $\hat{\phi}_2^{\text{priv}}(C)$ and zero above it, the cutoff coincides with that largest threshold. In this benchmark, sorting is *complete*, in the sense that every physician who would ever perform a non-indicated cesarean in either sector strictly prefers private practice. The corollary is that, under $\kappa = 0$, all public-sector cesareans are medically indicated. This is a useful theoretical anchor but not what we expect to observe in the data.

We take $\kappa > 0$ as the empirically relevant case. Two forces push it up. Public practice carries a nonpecuniary premium through academic prestige and teaching, and private physicians, who value their time more, bear a higher time cost on their routine caseload. We take these to outweigh the higher pecuniary returns of private practice, which work in the opposite direction, so that $\kappa > 0$. A public sector that retains more than just the physicians who never intervene is consistent with this. Under $\kappa > 0$ the cutoff $\phi^*(\kappa{>}0)$ falls strictly below $\hat{\phi}_2^{\text{priv}}(C)$, so sorting is only partial, with the lower-ϕ types concentrating in private and the higher-ϕ types remaining in public. How much discretion survives in the public sector depends on the size of κ. Only when κ is large enough to push the cutoff below $\hat{\phi}_2^{\text{pub}}(C)$, the case drawn in Figure 3, does the public pool retain physicians who still perform some non-indicated cesareans, and the near-flat public calendar gradient documented in Section 6 suggests that any such residual public discretion is small.

Sorting in practice can fall further short of the partial separation just described. Full separation of any kind requires the absence of frictions in sector allocation, and our setting has several. Private-sector positions are not unlimited and may be capacity constrained, and hospitals hire and retain physicians under contracts that are not perfectly flexible. These frictions push the equilibrium inside the bounds traced out by the Proposition, so we treat the sorting result as an amplifying mechanism rather than a tight prediction about which physicians work where.

What the empirical analysis can recover from cross-sector data is also limited. The cross-sector difference in non-indicated cesareans we eventually document reflects two channels working in the same direction. The within-physician channel says that, for a given physician, moving to a higher-ρ sector raises her optimal cesarean use, which is what the within-sector predictions describe. The composition channel says that the physicians who actually practice in the private sector have systematically lower ϕ by the sorting result, so even at the same incentive environment they would do more non-indicated cesareans. Were we able to observe the same physician working in both sectors, we could decompose the cross-sector gap into these two pieces. Because our data do not link physicians across sectors, we cannot disentangle them empirically. We therefore treat sorting as a model-implied amplifier of the within-sector predictions whose magnitude is not separately identified in our data, rather than as a stand-alone prediction.

3.6 Extension: medically indicated cases

The model of Sections 3.2–3.5 focuses on cesareans without medical indication, the cases over which the physician genuinely has a choice and whose nonclinical trade-off is governed by ϕ and γ. Observed cesarean rates, however, include both kinds. Some women have a documented clinical reason for a cesarean and the physician has no choice over whether to operate, only over when. This extension asks whether incorporating those medically indicated cesareans changes the qualitative predictions of the baseline model, and in particular whether it changes the sorting result of Section 3.5. Indicated cases add a background component to observed rates, but they do not act through the professionalism cost ϕ or the override cost γ that govern the discretionary decision, so the within-sector predictions P1–P3 are unaffected. In the sector-choice problem, they strengthen the private-sector advantage, so the sorting result of Proposition 1 is reinforced rather than overturned.

Indicated cesareans differ from non-indicated ones in two ways. First, the decision to operate is taken on clinical grounds and does not respond to ϕ or γ. There is no professionalism cost to pay because there is no judgment call to override, and there is no maternal preference channel because clinical necessity is not negotiable. Indicated cases therefore add a fixed background rate of cesareans to both sectors without altering the within-sector decision rule derived in Section 3.3. Second, indicated cesareans can still be scheduled when the diagnosis is made in advance, and scheduling responds to calendar incentives in exactly the way Section 3.2 described for non-indicated cases. Moving a scheduled indicated cesarean from a nonworking day to a working one yields a strictly positive utility gain of $\rho_s(\alpha-1)L_C$, the same shift gain as in Lemma 1. This gain is common across all physicians within a sector, because it does not depend on ϕ, and it is larger in the private sector than in the public one, because $\rho_{\text{priv}} > \rho_{\text{pub}}$.

To see what indicated cases do to the sector-choice framework, recall from Section 3.1 that a fraction $\xi \in (0,1)$ of cases are medically indicated, and suppose that, absent scheduling, indicated cases would fall on nonworking days with the same probability p_2 as non-indicated ones. After scheduling, indicated cases contribute an expected rent of $\xi\, p_2\, \rho_s(\alpha-1)L_C$ per delivery, the product of the probability that the case is indicated, the probability that biological onset would have fallen on a nonworking day, and the per-case utility gain from rescheduling. Adding this to the per-delivery discretionary rent of Equation (7) gives the total expected rent in sector s,

$$\widetilde{V}^s(\phi) \;\equiv\; V^s(\phi) \;+\; \xi\, p_2\, \rho_s(\alpha-1)L_C. \tag{9}$$

The new term is a sector-specific constant. It depends on sector primitives through ρ_s but does not depend on the physician's own type ϕ, because every physician in a given sector schedules indicated cases the same way regardless of how willing she is to perform non-indicated ones.

The sector-choice rule now becomes $\widetilde{\Delta V}(\phi) > \kappa$, where the new rent differential is

$$\widetilde{\Delta V}(\phi) \;\equiv\; \widetilde{V}^{\text{priv}}(\phi) \;-\; \widetilde{V}^{\text{pub}}(\phi) \;=\; \Delta V(\phi) \;+\; \xi\, p_2\, (\rho_{\text{priv}} - \rho_{\text{pub}})(\alpha - 1)L_C. \qquad (10)$$

The extra term is strictly positive (because $\rho_{\text{priv}} > \rho_{\text{pub}}$ and $\alpha > 1$), uniform across ϕ, and depends only on sector primitives. It is the additional gain the private sector offers from being able to reschedule indicated cesareans away from nonworking days, a gain that scales with ρ_s and is therefore larger in private practice. Geometrically, this shifts the entire $\Delta V(\phi)$ curve of Figure 3 uniformly upward without changing its shape. The single-crossing property is preserved, the cutoff characterization of Proposition 1 continues to apply, and, because lifting a downward-sloping curve moves its crossing with κ to the right, the new cutoff satisfies $\widetilde{\phi}^* \geq \phi^*$. The implication is that incorporating indicated cases brings (weakly) more physicians into the private sector. Because the added rent is the same for every ϕ, it shifts the sorting cutoff to the right but does not sharpen the selection on low-ϕ types, so it does not necessarily reinforce the composition channel of Section 3.5. Predictions P1, P2, and P3 are themselves unaffected, because they are driven by the discretionary margin, which operates through ϕ and γ and which indicated cases do not involve.

A natural concern about taking this extension to the data is that women who anticipate a more medically complex delivery may sort into private care to access a more specialized system, which would raise the share of indicated cesareans in the private sector relative to the public sector and confound the within-sector tests of P1–P3. Two features of the empirical design address this concern. First, the analysis focuses on a cohort of low-risk first births, whose eligibility criteria exclude women with documented clinical risk factors that would predict such sorting on observable medical complexity. Second, results are reported separately for elective and intrapartum cesareans, which limits the extent to which any residual difference in the indicated background can contaminate the within-margin comparisons used to test P1–P3.

4 Data

We use individual-level survey and medical-record data from a prospective cohort study conducted by the Institute for Clinical Effectiveness and Health Policy (IECS) in the Buenos Aires metropolitan area. The cohort enrolled 382 pregnant women between 24 August 2010 and 20 December 2011 and followed them through delivery. Women were interviewed during the third trimester to elicit socioeconomic characteristics and their stated preference over delivery mode. Delivery outcomes were then extracted from clinical records, including the date of delivery, the realized mode of delivery, and whether a CS was scheduled prior to the onset of labor.

The cohort was designed to focus on relatively low-risk first births. Eligibility criteria included nulliparous women aged 18 to 35 years at enrollment with singleton pregnancies and a live fetus beyond 32 weeks of gestation. Women with assisted reproduction, major pre-existing conditions, pregnancy complications, or a documented medical indication for a pre-labor CS were excluded. Because some women turned 36 between enrollment and delivery, we retain women aged 18 to 36 years at delivery. Two women were removed during follow-up (one lost to follow-up and one with a multiple pregnancy), and one woman delivered outside the hospital setting. The main analysis therefore focuses on an analytic sample of 379 hospital deliveries, with 197 in public hospitals and 182 in private hospitals. The cohort was enrolled at five hospitals (two public and three private) selected for high delivery volume and long-standing collaboration with IECS (Mazzoni et al., 2016). Most deliveries occurred across six hospitals in and around the City of Buenos Aires, with the additional public hospital (Hospital Municipal Materno Infantil de Tigre Dr. Florencio Escardo) receiving 26 deliveries from women enrolled at other public sites. Figure 4 locates the six main delivery hospitals within the universe of inpatient facilities in the Buenos Aires metropolitan area.[11]

A key advantage is that the data separate delivery timing from intrapartum decisions. Medical records distinguish elective (scheduled) from intrapartum (unscheduled) CS. We define an elective CS as a procedure scheduled before labor begins, while intrapartum CS are those performed after labor onset. A small number of urgent pre-labor cesareans fall outside this dichotomy and are excluded from the intrapartum analysis. This mapping is central to the empirical design: in the model, shifting deliveries away from weekends and public holidays operates through scheduled, predictable interventions. The records also indicate whether labor was induced, which we use in sensitivity analyses. The data include the day, but not the hour, of delivery, which limits within-day analysis and implies that calendar-day measures are necessarily coarse, particularly for intrapartum cases that may begin late at night.

To operationalize calendar-based convenience, we classify dates using the Argentine holiday calendar. Our study period spans 459 consecutive days (25 September 2010 through 27 December 2011). We define nonworking days as weekends and national public holidays, and working days as all remaining dates. This classification underlies both the descriptive evidence on birth timing and the convenience measures that enter

[11]The hospital fixed effects comprise six hospital indicators. The single delivery recorded at another public site is grouped with a named public hospital, while the three deliveries at other private sites lack a hospital identifier and drop out of specifications that include hospital fixed effects.

Figure 4. Study hospitals within the Buenos Aires metropolitan area

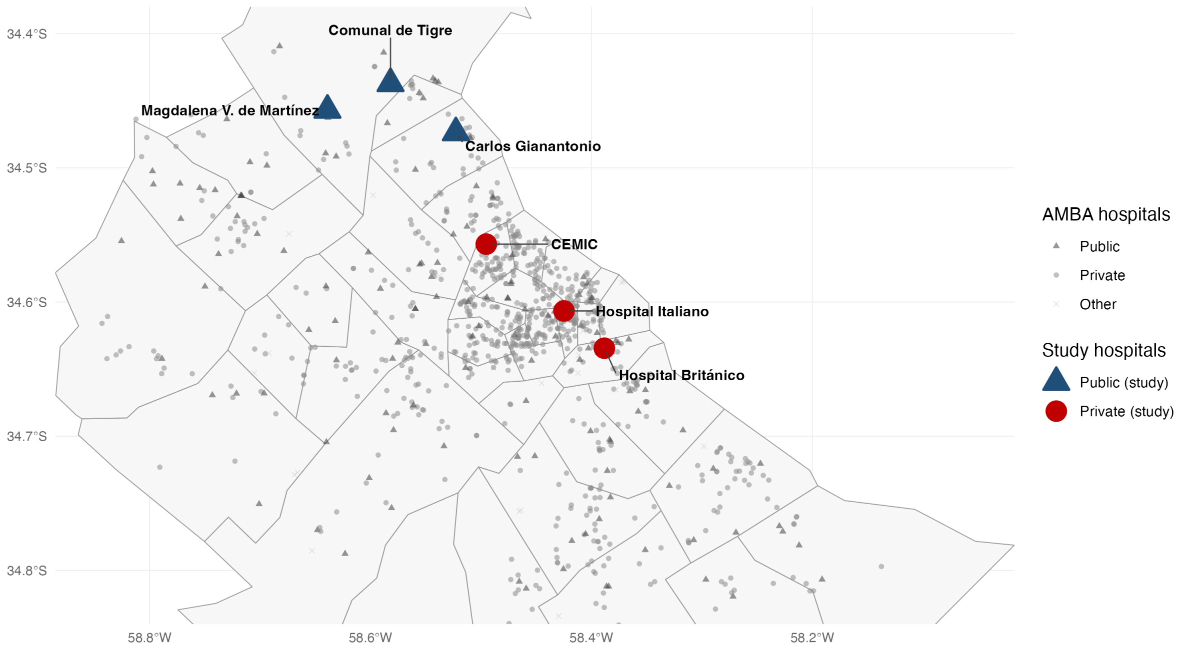

Notes: The map covers the Metropolitan Area of Buenos Aires. Grey markers denote all hospitals with inpatient services in the area, classified by financing origin (triangle = public, circle = private, cross = other/social security). The six delivery hospitals in the analytic sample are highlighted. The public ones (blue) are Hospital Zonal General de Agudos Magdalena V. de Martínez (Tigre), Hospital Materno Infantil Dr. Carlos Gianantonio (San Isidro), and Hospital Municipal Materno Infantil de Tigre Dr. Florencio Escardo (Comunal de Tigre). The private ones (red) are Hospital Italiano de Buenos Aires, CEMIC (sede Saavedra), and Hospital Británico de Buenos Aires. *Source:* Registro Federal de Establecimientos de Salud (REFES, December 2025), Ministry of Health. Administrative boundaries from INDEC.

the empirical specifications.

4.1 Summary statistics

Table 1 summarizes delivery modes by hospital sector, with counts disaggregated by individual hospital in Appendix Table A9. The cesarean rate in private hospitals (44.0%) exceeds the public-sector rate (34.5%) by nearly ten percentage points, but the composition of that gap is more informative than its size. Elective cesareans account for 14.8% of private-sector births against 4.6% in the public sector, a three-to-one difference concentrated on the timing margin of the model, where physicians schedule discretionary procedures in advance. Intrapartum cesareans, by contrast, are close in magnitude across sectors (34.2% versus 30.6% of women in labor) and mix medically indicated escalations with same-day conversions. Whatever drives the sectoral cesarean gap in Buenos Aires, it operates through the scheduling margin, which the model identifies as most exposed to physician discretion.

Table 1. Delivery modes by hospital sector

	Public	Private	Total
Total births	197	182	379
Vaginal births	129	102	231
Share of births	65.48%	56.04%	60.95%
Cesarean births	68	80	148
Share of births	34.52%	43.96%	39.05%
Elective (scheduled) cesareans	9	27	36
Share of births	4.57%	14.84%	9.50%
Intrapartum cesareans	57	53	110
Share of women in labor	30.65%	34.19%	32.26%

Notes: Intrapartum shares are computed among women who entered labor (excluding elective and urgent pre-labor cesareans). Elective and intrapartum categories do not fully exhaust cesareans because a small number of procedures are recorded as urgent pre-labor cesareans in the medical files.

Figure 5 shows the distribution of births by day of the week. Several patterns stand out. In private hospitals, cesareans cluster on Mondays and Fridays. Weekends are comparatively quiet, with a marked decline in CS. In public hospitals, births peak on Tuesdays, but this increase does not appear to be driven by CS. The cesarean share is elevated on public-sector Saturdays, but this reflects the small number of weekend deliveries there, which are disproportionately unscheduled, rather than any scheduling pattern. Figures 7 and 8 in Section 6.2 replicate the same plots separately for elective and intrapartum CS, respectively.

Nonworking days (weekends and public holidays) make up 32.9% of calendar days in the study window but host only 23.8% of deliveries. The compression onto working days is sharper in private hospitals, where 18.7% of births occur on nonworking days, against 28.4% in public hospitals (Appendix Table A2). The gap does not appear to be driven by a sector-specific distribution of biological onset: restricting the comparison to vaginal deliveries, which are unscheduled by construction, shrinks the sectoral gap in nonworking-day shares to 1.6 percentage points (27.1% in public versus 25.5% in private), statistically indistinguishable from zero. Because a minority of vaginal deliveries follow induction and some labors cross calendar days, this comparison is suggestive rather than decisive, but it shows no sign of the sectoral onset differences that could otherwise account for the compression. Appendix Figure A1 reports the day-of-week distribution and fails to reject the null of no sectoral difference. The

Figure 5. Births by day of the week

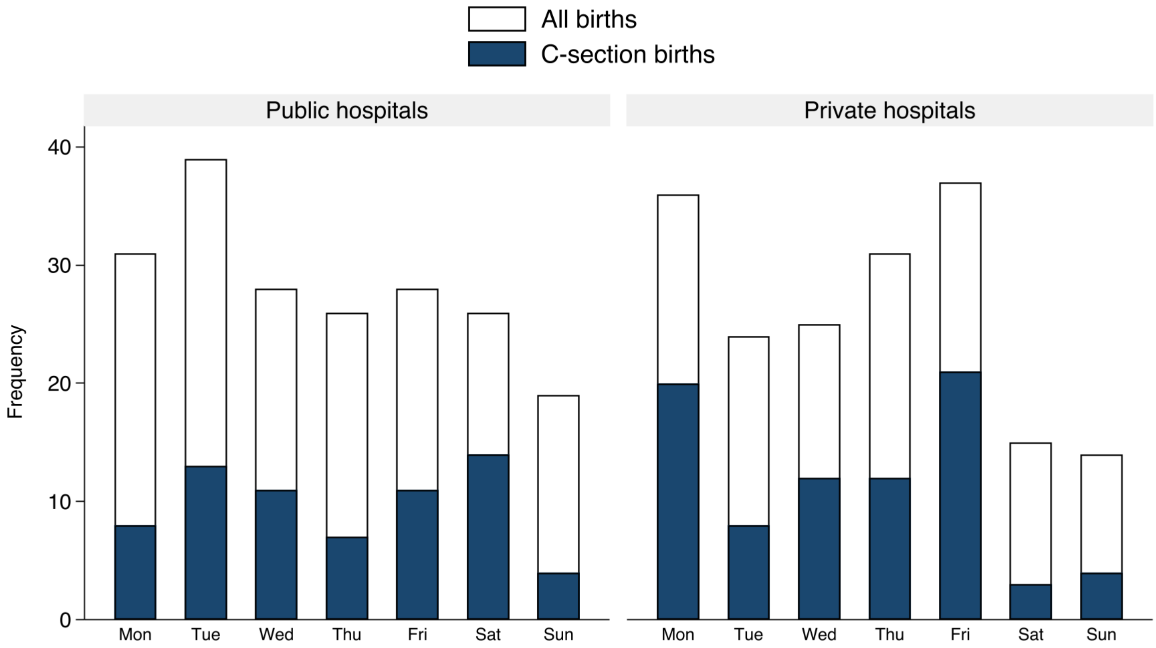

Notes: Counts of deliveries by day of the week and hospital sector, split by mode of delivery (cesarean vs vaginal). Sample: $N = 379$ hospital deliveries (197 public, 182 private) in the analytic sample described in Section 4.

sectoral differential in nonworking-day deliveries is therefore concentrated in scheduled cesareans rather than in the underlying distribution of spontaneous labor. Working-day bunching is thus a private-sector phenomenon operating through scheduling. It emerges precisely where, under the model's framing, physicians act as residual claimants of their own time and therefore internalize the higher opportunity cost α of weekend work.

Appendix Table A3 reports observed characteristics by hospital type, consistent with the stratification of insurance coverage described in Section 2. Women delivering in private hospitals are, on average, 29.9 versus 22.8 years old, have completed tertiary education at rates of 81.9% versus 10.7%, are employed at rates of 83.1% versus 15.7%, are married or cohabiting at rates of 93.4% versus 69.9%, and have private or social-security insurance coverage at rates of 96.7% versus 0%. Appendix Table A11 reports the census-block Sanitary Vulnerability Index of Rosati et al. (2020) at the nearest centroid and averaged over 2 km and 5 km buffers around each study hospital; the index differs across facilities, including within the public sector, and for several hospitals also across spatial scales (some sit in a higher-vulnerability immediate block than in the wider zone, others show the opposite gradient). The empirical specifications in Section 5 absorb these cross-hospital differences in maternal composition and catchment vulnerability through hospital fixed effects.

Table 2 reports stated delivery-mode preferences and realized cesarean rates. A small share of women state a preference for cesarean delivery, with similar frequency across sectors (8.1% in public hospitals, 6.5% in private hospitals), yet the mapping from stated preference to realized delivery mode diverges sharply both across and within sectors. Across sectors, among women who prefer a cesarean (the subset for whom the override cost γ does not bind), 72.7% receive one in private hospitals against 43.8% in public hospitals, a 29-percentage-point gap consistent with the higher time value ρ_{priv} pushing more physicians across the intervention threshold. Within each sector, the CS-rate gap between women who prefer a cesarean and women who prefer vaginal delivery is 32.5 percentage points in private hospitals and 10 percentage points in public hospitals, a three-fold difference consistent with the preference channel being empirically active in the private sector, where the higher intervention threshold S_t^{priv} places more physicians in the preference-responsive range $[S_t^{\text{priv}} - \gamma, S_t^{\text{priv}}]$ defined by the override cost. Both asymmetries are consistent with P1. Because Table 2 pools elective, intrapartum, and urgent pre-labor cesareans, these raw rates are descriptive rather than a direct test of the preference-to-elective channel, which we examine on elective cesareans in Section 6.1.

Table 2. Stated delivery-mode preferences and cesarean rates

	Preference share	**CS rate**
Preference for vaginal delivery		
Public hospitals	91.88%	33.70%
Private hospitals	93.53%	40.25%
Total	*92.64%*	*36.76%*
Preference for cesarean delivery		
Public hospitals	8.12%	43.75%
Private hospitals	6.47%	72.73%
Total	*7.36%*	*55.56%*
No preference		
Public hospitals	0.00%	–
Private hospitals	6.59%	66.67%
Total	*3.17%*	*66.67%*

Notes: Preferences are elicited during the third trimester. For "Preference for vaginal delivery" and "Preference for cesarean delivery," the reported preference shares are computed conditional on expressing a preference (i.e., excluding "No preference"). For "No preference," the reported share is computed over all women in the sector (including those who recorded no preference). CS rates are computed within each row's group. "No preference" responses occur only in the private-hospital subsample.

5 Empirical Strategy

The model in Section 3 highlights which nonmedical factors can affect observed cesarean use. There is a *timing* margin: when a delivery is anticipated to occur during nonworking time, a predictable technology such as a scheduled cesarean can shift the delivery date to a nearby working day. The model predicts that both the timing margin and the responsiveness of outcomes to maternal preferences should be stronger in private hospitals, where physicians personally capture the value of saved time, than in public hospitals, where team production within fixed shifts attenuates the return to time saved on any single case. The empirical strategy mirrors these mechanisms by separating elective from intrapartum cesareans and comparing outcomes across working and nonworking days.

5.1 Outcomes and analysis samples

Let $h(i)$ denote the hospital in which woman i delivers and let $Priv_{h(i)}$ be an indicator for whether that hospital is private. Using the medical records, we construct three outcomes. The indicator CS_i equals one if the delivery is by cesarean section. The indicator ECS_i equals one if the cesarean is scheduled prior to labor onset (elective cesarean). The indicator ICS_i equals one if a cesarean is performed after labor begins (intrapartum cesarean). These outcomes are complementary because they correspond to different decision points in the model. Scheduled cesareans are the natural empirical counterpart of timing shifts away from nonworking time, while intrapartum cesareans capture escalation decisions once labor is underway.

The elective analysis is conducted on the full delivery sample. The intrapartum analysis is conducted on the subsample of women who enter labor. Operationally, the labor sample used for the intrapartum analysis excludes both women with an elective cesarean and a small number of cesareans recorded in the files as urgent pre-labor procedures that are neither scheduled in advance nor performed after labor begins. Among the remaining women, who enter labor, we compare intrapartum cesareans to vaginal deliveries. Because the urgent pre-labor cases do not map cleanly to either decision margin, they do not enter the intrapartum specifications. Appendix Table A10 reports delivery modes in the labor sample by hospital sector.

Maternal preferences are measured in the third-trimester survey, before delivery occurs. Let $PrefCS_i$ be an indicator equal to one if the woman reports a preference for cesarean delivery and equal to zero if she reports a preference for vaginal delivery. Women who report no preference or have missing responses are excluded from regressions that use $PrefCS_i$. Let Age_i denote woman i's age at delivery.

5.2 Calendar-based measures of nonworking time

Our proxy for calendar-based convenience is whether delivery occurs during working time. Let D_i denote woman i's delivery date, and let $t(D_i)$ denote its calendar category in the model's partition of Section 3.1. In the baseline definition, a nonworking day ($t(D_i) = 2$) is a Saturday, Sunday, or national public holiday. Define

$$Work_i^{(1)} \equiv \mathbb{1}\{t(D_i) \in \{0, 1\}\} = \mathbb{1}\{D_i \text{ is a working day under definition (1)}\}.$$

This measure corresponds to treating weekends and public holidays as the set of days carrying the higher time weight $\omega_t = \alpha > 1$ in the physician's utility (1), against $\omega_t = 1$ on working days. Because leisure considerations may extend to days adjacent to nonworking time, we also consider two alternative definitions. Under definition (2), the set of nonworking days includes weekends and public holidays and also includes the working day immediately preceding any public holiday. When a public holiday falls on a Monday, the preceding day is defined as the prior Friday. When a holiday falls on a weekend, the preceding day is not reclassified. Under definition (3), the set of nonworking days is expanded only for long weekends, by adding the working day that bridges a public holiday to the adjacent weekend—the Friday preceding a Monday holiday, or the Thursday preceding a Friday holiday. For each definition $k \in \{1, 2, 3\}$ we define the corresponding indicator $Work_i^{(k)}$. Appendix Table A8 summarizes the three definitions.

A further implication of the model is that timing shifts should concentrate scheduled cesareans on boundary working days ($t = 1$) adjacent to nonworking time. In descriptive analysis and in selected specifications, we therefore also use an indicator for boundary days, defined as Monday and Friday working days (the working days that flank a weekend).

A limitation of the data is that we observe the day of delivery but not the hour and do not observe the precise onset of labor. For elective cesareans, this is not consequential for identification because timing is part of the decision. For intrapartum cesareans, the model's primitive is the onset day, whereas the data provide the delivery day. The baseline analysis therefore interprets working versus nonworking classifications as applying to the realized delivery date. Because labor can span calendar days and intrapartum management can affect the interval from onset to delivery, we read the resulting intrapartum gradient as a delivery-date pattern consistent with P3 rather than as a clean test of onset-day selection.

5.3 Econometric specifications

The institutional mechanism in Section 3 implies that convenience incentives and preference responsiveness should differ by sector. We estimate specifications that allow both working-day effects and preference effects to vary between private and public hospitals, while absorbing time-invariant differences across facilities through hospital fixed effects. These absorb not only the sectoral split that drives the core mechanism of the model, but also the within-sector heterogeneity in catchment-area composition documented in Section 4.1 (Appendix Table A11); the variability in catchment vulnerability across facilities and spatial scales motivates absorbing this heterogeneity through fixed effects rather than through a small set of area-level controls. Since sector is constant within a hospital, the main effect of $Priv_{h(i)}$ is absorbed by the fixed effects, and the parameters of interest are the interaction terms that capture differential responses in private hospitals.

For elective cesareans, the delivery date is an outcome of scheduling rather than a predetermined regressor. We therefore focus on how the propensity to schedule a cesarean varies with maternal preferences across sectors. The estimating equation is

$$\Pr(ECS_i = 1 \mid PrefCS_i, Age_i, h(i)) = \Phi\Big(\mu_{h(i)} + \delta_1\, PrefCS_i + \delta_2\,(PrefCS_i \times Priv_{h(i)}) + g(Age_i)\Big), \tag{11}$$

where $\Phi(\cdot)$ is the standard normal cdf, $\mu_{h(i)}$ denotes hospital fixed effects, and $g(Age_i)$ is a quadratic in maternal age. The coefficient δ_2 is the interaction term in the probit index, so its sign indicates whether stated preferences translate more strongly into scheduled procedures in the private sector, while the corresponding magnitude on the probability scale is given by the predicted probabilities reported in Section 6.1. A positive δ_2 is consistent with the within-sector mechanism of P1 operating more strongly in the private sector, which the sector-choice extension of Section 3.5 rationalizes through the sorting of low-ϕ physicians into private practice.

For intrapartum cesareans, we restrict the sample to women who enter labor and compare outcomes across working and nonworking days within the same hospital. For each convenience definition $k \in \{1, 2, 3\}$ we estimate

$$\begin{aligned}\Pr(ICS_i = 1 \mid Work_i^{(k)}, PrefCS_i, Age_i, h(i)) = \Phi\Big(&\mu_{h(i)} + \beta_1\, Work_i^{(k)} + \beta_2\,(Work_i^{(k)} \times Priv_{h(i)}) \\ &+ \beta_3\, PrefCS_i + \beta_4\,(PrefCS_i \times Priv_{h(i)}) + g(Age_i)\Big).\end{aligned} \tag{12}$$

Here β_1 is the working-day versus nonworking-day differential in the probit index for public hospitals, while β_2 captures the additional working-day differential in private hospitals. The model predicts that the selection induced by timing shifts makes intrapartum cesareans relatively more likely on working days in private hospitals, so that

β_2 should be positive. Conditional on labor, the model also predicts limited sensitivity to maternal preferences. The elective margin is the physician's primary channel for non-indicated intervention: a physician willing to act on a stated preference for cesarean will typically do so by scheduling an ECS in advance rather than by converting during labor. What remains in the labor sample is therefore a pool of physicians who, regardless of the mother's stated preference, are less inclined to perform a non-indicated cesarean. The coefficients β_3 and β_4 should therefore be small relative to the preference effects for elective cesareans. The interaction β_4 may even turn negative, because in the private sector women who prefer cesarean disproportionately receive an ECS in advance, leaving the intrapartum pool depleted of the preference-responsive types whose preferences would otherwise map to cesarean during labor.

We estimate (11) and (12) by probit and report average marginal effects and predicted probabilities for ease of interpretation. Standard errors are heteroskedasticity-robust. Given the small number of hospitals, the empirical emphasis is on within-hospital contrasts and on the stability of the estimated private-sector differentials across alternative convenience definitions and sample restrictions.

6 Results

This section presents evidence on the three testable implications laid out in Section 3.4. We analyze them in two blocks. Section 6.1 examines the preference margin (P1), combining the interaction of stated preference and sector in the elective cesarean equation with predicted probabilities that summarize the magnitudes. Section 6.2 examines the timing margin (P2) and its selection effect on intrapartum outcomes (P3), combining the calendar distribution of elective cesareans with the working-day versus nonworking-day gradient in intrapartum cesareans among women who enter labor. A specification that combines both margins into a single regression for overall cesarean use is reported for completeness in Appendix B. Because scheduled cesareans directly determine the delivery date, that specification mechanically combines the two margins and is not the primary object of interest.

6.1 The preference margin

We test P1 by focusing on elective cesareans. Table 3 reports probit estimates for the probability of scheduling a cesarean, omitting delivery-day convenience because the delivery date is itself part of the scheduling decision. The interaction between private hospitals and a stated preference for cesarean delivery is positive in both specifications and statistically significant at the 5 percent level when hospital fixed effects are included, indicating that stated preferences translate into scheduled procedures more strongly in private hospitals than in public ones and matching the sign predicted by P1. The observed asymmetry is consistent with the within-sector mechanism of P1 operating more strongly in private. Sector selection (Section 3.5) reinforces this pattern. Because lower-ϕ physicians sort into private practice, the public pool is left relatively more concentrated in physicians reluctant to perform non-indicated cesareans, so the preference response in public hospitals is weaker and statistically undetectable.

Table 3. Elective cesarean: probit estimates

Variable	(1)	(2)
Private hospital (PrH)	0.339	–
	(0.264)	–
Age (centered at 18)	-0.109	-0.109
	(0.075)	(0.079)
Age squared	0.008*	0.007*
	(0.004)	(0.004)
Prefers CS	0.106	0.041
	(0.498)	(0.515)
PrH × Prefers CS	1.087*	1.388**
	(0.650)	(0.643)
Constant	-1.491***	-1.813***
	(0.255)	(0.343)
Observations	367	338
Hospital fixed effects	NO	YES

Notes: Dependent variable is an indicator for elective (scheduled) cesarean section. Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. The fixed-effects column drops 26 deliveries at one public hospital with no elective cesareans (perfect separation) and 3 deliveries without a hospital identifier.

To quantify the magnitude, Table A4 reports predicted probabilities under both specifications (we report probabilities from the specification without fixed effects in the prose because hospital fixed effects absorb level differences in the public-sector baseline). Averaging over maternal preferences and age, the predicted probability of elective cesarean is 14.0% in private hospitals against 4.6% in public hospitals. Cesarean use is far more responsive to stated preference in the private sector: among women who prefer a cesarean, the predicted probability is 48.0% in private against 5.6% in public, whereas among women who prefer vaginal delivery it is 11.5% versus 4.5%. This differential responsiveness is the signature of the preference channel in P1, rather than the overall level gap, to which the large majority who prefer vaginal delivery contribute most. Figures 6a and 6b plot these predicted probabilities by age, preferences, and sector. In public hospitals, the curves for preference-for-cesarean and preference-for-vaginal women nearly coincide at every age, whereas in private hospitals they diverge sharply, with the preference-for-cesarean curve substantially higher.

A natural concern is that the cross-sector difference in preference responsiveness reflects unobservable patient characteristics, such as assertiveness, prior medical knowl-

Figure 6. Predicted probability of elective cesarean by age and hospital sector

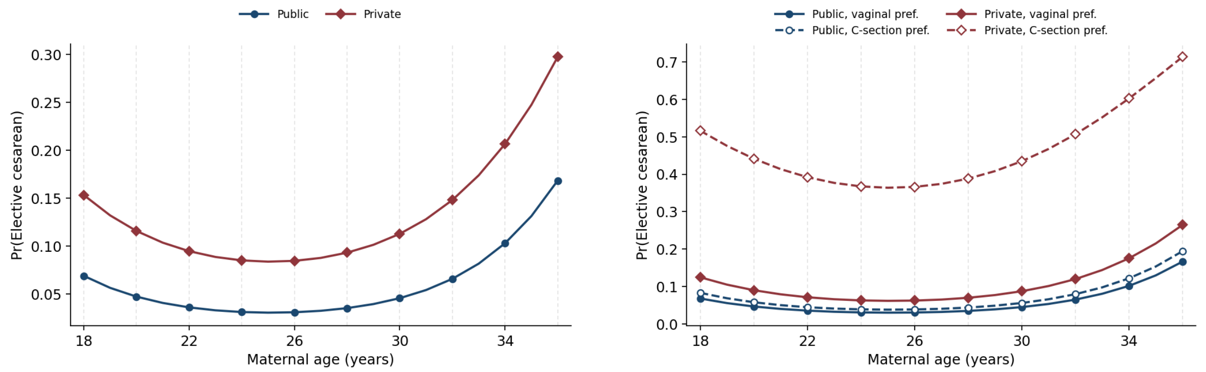

(a) By age and hospital sector

(b) By age, preferences, and hospital sector

Notes: Predicted probabilities of elective (scheduled) cesarean by maternal age, computed from the probit model in Column (1) of Table 3 (without hospital fixed effects). Panel (a) averages over preferences, Panel (b) conditions on stated preference. Sample: $N = 367$ women with a recorded specific preference.

edge, or the strength of the patient-physician relationship, that are correlated with both sector choice and the probability of obtaining a cesarean and that are not captured by age, stated preferences, or hospital fixed effects.[12] Appendix C assesses robustness to selection on unobservables using the coefficient-bounds approach of Oster (2019). Applied to the linear probability analogue of the specification in Table 3, the bias-adjusted coefficient on PrefCS $\times$ Private remains positive, and unobservable selection would have to be more important than the full set of observed covariates to overturn the private-sector differential. Additionally including an interaction between education and stated preference, intended as a direct proxy for patient agency within the preference channel, *increases* the controlled coefficient rather than attenuating it, so no positive proportional selection on unobservables can reconcile the estimate with a zero effect.

A related concern is that stated preferences, elicited in the last trimester of gestation, may partially reflect physician advice received during the pregnancy rather than preferences formed independently. We cannot rule this out. However, even under this interpretation the differential pattern across sectors remains informative: if private physicians are more likely to both shape patient preferences toward cesarean delivery and follow through on those preferences, the interaction coefficient captures the combined effect of these two channels, both of which are expressions of physician agency under private organizational design. The result is thus conservative with respect to the paper's broader argument, even if the preference channel alone cannot be cleanly separated from the advice channel.

[12]The canonical example is a more educated woman who is better able to assert her delivery-mode preference in discussions with her physician. Education is strongly correlated with private-hospital choice (Table A3), so if patient agency varies with education and operates independently of our observed controls, it would confound the cross-sector comparison of preference responsiveness.

6.2 The timing margin

We now test P2 and P3, which together describe the timing margin and its selection effect on intrapartum outcomes. P2 predicts that elective cesareans bunch on boundary working days (the Monday or Friday adjacent to a weekend or holiday) in the private sector. P3 predicts that, among women who enter labor, the intrapartum cesarean rate is higher on working days than on nonworking days in the private sector and flat by calendar type in the public sector, as a consequence of selection induced by P2 rather than of independent intrapartum discretion. Figure 5 already shows that private hospitals compress deliveries into working days, with visible bunching around Mondays and Fridays. Figures 7 and 8 decompose this pattern. Elective cesareans in private hospitals cluster on Mondays and Fridays and are almost absent from weekends and public holidays. The public sector shows no comparable pattern (Figure 7). Among women who enter labor, the residual intrapartum cesareans in private hospitals are concentrated on working days, while the distribution is roughly flat in the public sector (Figure 8). The two panels match the descriptive content of P2 and P3.

Figure 7. Elective cesarean births by day of the week

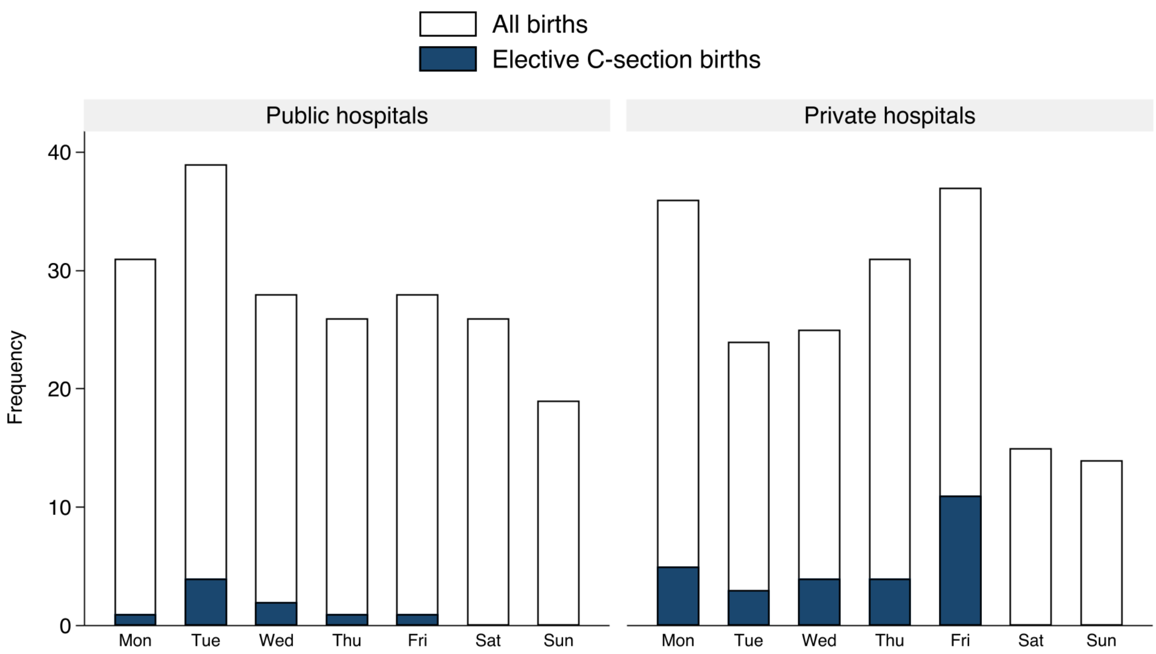

Notes: Counts of elective (scheduled) cesareans by day of the week and hospital sector. Sample: $N = 36$ elective cesareans (9 public, 27 private) in the analytic sample.

Table 4 reports two complementary probit tests. Column (1) restricts to working days and regresses the ECS indicator on a boundary-day dummy interacted with private hospital, with hospital fixed effects and an age quadratic. The Boundary $\times$ Private coefficient is 0.87 ($p = 0.053$), matching the sign predicted by Lemma 1 and P2. On the probability scale, within private hospitals the predicted ECS rate is sharply higher on boundary working days than on interior working days; within public hospitals the two rates are similar (Appendix Table A6). Column (2) of Table 4 asks a complementary question: conditional on a scheduled cesarean, how much more likely is it to fall on a boundary day in the private sector than in the public sector? The private-hospital coefficient is positive and statistically significant. Of the 27 ECS performed in private hospitals, 15 (55.6%) are scheduled on a boundary working day (Monday or Friday), against 2 of 9 (22.2%) in public (Fisher exact test, $p = 0.128$).

Table 4. Elective cesareans concentrate on boundary working days in the private sector

	(1) Pr(*ECS*)	(2) Pr(*Boundary*)
Boundary × Private	0.867*	
	(0.448)	
Boundary day	−0.511	
	(0.371)	
Private hospital	0.793*	0.904*
	(0.451)	(0.532)
Hospital fixed effects	Yes	No
Sample	Working days	ECS only
Observations	271	36

Notes: Probit estimates. Column (1): dependent variable is an indicator for an elective (scheduled) cesarean, the sample is restricted to working days under convenience definition (1), the specification controls for a quadratic in maternal age (coefficients not shown) and hospital fixed effects, and one small public hospital with zero ECS is dropped by perfect separation. Because a second delivery hospital is collinear with the private-sector indicator in this subsample, that indicator is retained as part of the hospital fixed-effects parameterization in Column (1) and does not represent a sector effect net of those fixed effects. Column (2): dependent variable is an indicator for delivery on a Monday or Friday, and the sample is restricted to elective cesareans. "Boundary day" equals one for deliveries on a Monday or Friday that is a working day. Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

We next quantify the intrapartum calendar gradient. The onset of spontaneous labor is not under the control of either the patient or the physician, so calendar-day differences in intrapartum cesarean rates do not reflect patients selecting their delivery day. They arise instead from the composition of the labor sample, which the timing margin shapes differently across day types. In the model, a positive working-day gradient in intrapartum cesareans does not require a separate intrapartum decision margin: the same discretionary choice of the physician (Section 3.3), interacted with the timing technology of Lemma 1, produces the observed gradient through a composition effect on the labor sample.

After excluding elective cesareans, weekend and holiday deliveries in the labor sample are a selected group. Whenever a nonworking-day onset is anticipated ($t(d) = 2$) and a discretionary cesarean is chosen, it is optimal to schedule it on an adjacent

Figure 8. Intrapartum cesarean births by day of the week

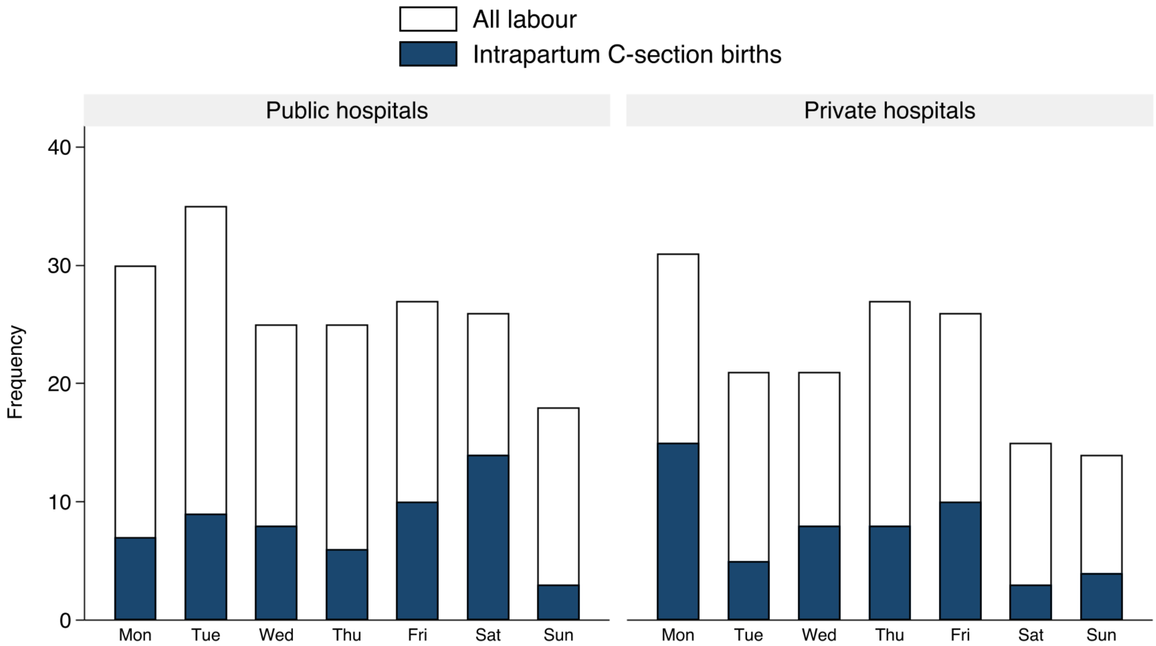

Notes: The difference between "all labor" and "intrapartum CS births" corresponds to vaginal deliveries.

boundary working day (Shift-C) rather than perform it on the nonworking day. As a result, discretionary cases that would otherwise generate weekend cesareans are disproportionately shifted out of nonworking days and therefore do not appear as intrapartum weekend cesareans.

The intrapartum cesareans that remain on weekends and holidays are thus composed of two groups: medically indicated escalations, which are present every day, and a small residual of discretionary cases for which scheduling was not feasible. The residual consists mostly of cases in which spontaneous labor began earlier than the physician anticipated, before a boundary-day elective cesarean could be arranged. This residual should be small, so weekend intrapartum cesareans are disproportionately concentrated in indicated escalations.

Working-day intrapartum cesareans, in contrast, are more heterogeneous: they combine indicated escalations with discretionary cesareans that were not scheduled in advance and are therefore executed during labor. Because $\rho_{\text{priv}} > \rho_{\text{pub}}$ makes the discretionary decision more frequent in the private sector, the working-day intrapartum pool is larger in private hospitals, which together with the weekend selection effect produces the P3 gradient.

Table 5 reports probit estimates of the probability of intrapartum cesarean as a function of delivery-day convenience and its interaction with private hospitals, under the three convenience definitions. Across specifications, the interaction term is positive and statistically significant, indicating that intrapartum cesareans are more likely to occur on working days in private hospitals than on weekends and public holidays. The negative level coefficient on *Private hospital* in the specifications without fixed effects reflects the compositional effect documented above: because discretionary cases in private hospitals are disproportionately scheduled as elective cesareans and therefore excluded from the labor sample, the private labor pool contains a lower share of discretionary cesareans and a correspondingly higher share of vaginal deliveries, lowering the baseline intrapartum rate relative to public hospitals. The interaction term captures the differential calendar gradient that arises despite this lower baseline.

Table 5. Intrapartum cesarean: probit estimates by convenience definition

Variable	(1) *ICS*	(1) *ICS*	(2) *ICS*	(2) *ICS*	(3) *ICS*	(3) *ICS*
Private hospital (PrH)	-0.565*	–	-0.553*	–	-0.597*	–
	(0.331)	–	(0.301)	–	(0.314)	–
Working day	-0.217	-0.263	-0.208	-0.261	-0.176	-0.205
	(0.210)	(0.214)	(0.203)	(0.207)	(0.206)	(0.209)
PrH × Working day	0.610*	0.677*	0.635*	0.698**	0.663**	0.707**
	(0.349)	(0.361)	(0.325)	(0.333)	(0.333)	(0.342)
Age	0.0684	0.0634	0.0667	0.0610	0.0658	0.0604
	(0.055)	(0.056)	(0.055)	(0.056)	(0.055)	(0.056)
Age squared	-0.00251	-0.00256	-0.00240	-0.00240	-0.00226	-0.00228
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Prefers CS	0.294	0.317	0.288	0.309	0.297	0.320
	(0.350)	(0.346)	(0.348)	(0.342)	(0.351)	(0.346)
PrH × Prefers CS	-0.123	-0.148	-0.146	-0.186	-0.157	-0.185
	(0.665)	(0.680)	(0.663)	(0.676)	(0.663)	(0.676)
Constant	-0.615**	-1.090***	-0.628***	-1.084***	-0.646***	-1.116***
	(0.242)	(0.377)	(0.234)	(0.376)	(0.237)	(0.377)
Observations	332	329	332	329	332	329
Hospital fixed effects	NO	YES	NO	YES	NO	YES
Convenience definition	(1)	(1)	(2)	(2)	(3)	(3)

Notes: Sample excludes elective cesareans and is restricted to women who enter labor. "Working day" indicates a delivery classified as working under the corresponding convenience definition. Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Appendix Table A7 translates the probit estimates into predicted probabilities by sector and day type. In private hospitals, the intrapartum cesarean rate is 35.7 percent on working days and 22.4 percent on nonworking days, a 13.3 percentage-point gap. In public hospitals, the corresponding rates are 28.3 and 36.0 percent, a difference of -7.7 percentage points. The private–public difference in this working-day gradient is statistically significant at the 10 percent level (the interaction in Table 5). Figures 9a and 9b replicate these patterns across maternal age and by stated preference, and confirm that preferences are much less predictive for intrapartum cesareans than for elective procedures, in line with the model's emphasis on scheduling as the main channel through which nonmedical factors operate.

Figure 9. Predicted probability of intrapartum cesarean by age and hospital sector

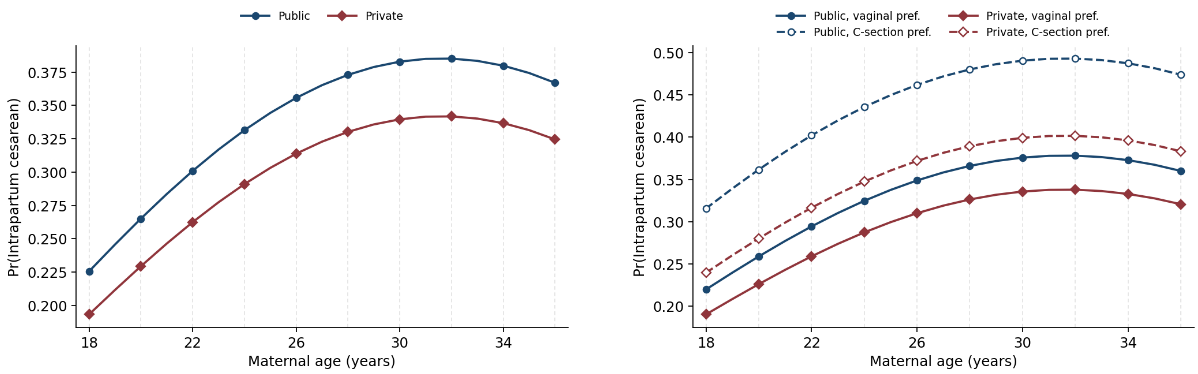

(a) By age and hospital sector

(b) By age, preferences, and hospital sector

Notes: Predicted probabilities of intrapartum cesarean by maternal age, computed from the probit specification in Column (1) of Table 5 (without hospital fixed effects) under convenience definition (1). Panel (a) averages over preferences, Panel (b) conditions on stated preference. Sample: women who enter labor ($N = 332$).

A potential concern for the P3 test is that the working-day gradient in intrapartum cesareans could reflect physician-scheduled inductions of labor rather than the selection mechanism of Lemma 1, since induction is itself a timing technology available to the physician. The medical records identify 55 induced labors in the labor sample (16%), with similar sectoral shares (15.6% in public versus 16.9% in private). Restricting the sample to labors with a spontaneous onset sharpens rather than attenuates the private-sector calendar effect. Under convenience definition (1) with hospital fixed effects, the PrH $\times$ Working day coefficient rises from 0.68 ($p = 0.060$, $N = 329$) in the full labor sample to 0.77 ($p = 0.050$, $N = 272$) in the spontaneous subsample. The analogous estimates under definitions (2) and (3) are 0.81 ($p = 0.028$) and 0.82 ($p = 0.029$), both significant at the 5% level. Appendix Table A5 reports the full probit specifications. Induction scheduling is therefore not driving the intrapartum calendar gradient.

The calendar patterns in Figures 7 and 8 and the regression results in Tables 4 and 5 match P2 and P3, and the predicted-probability curves in Figure 9 confirm that the intrapartum patterns are not driven by maternal age or stated preference.[13]

[13]The evidence is consistent with the selection mechanism formalized in P3, but does not rule out a complementary direct channel in which private physicians are more likely to perform a cesarean during labor on working days independently of any prior scheduling decision. Distinguishing the two would require direct measures of clinical indication severity by day type, which are not available in our data. We therefore interpret the working-day gradient as consistent with the model's prediction rather than as a test of the selection mechanism per se. The descriptive pattern in Figure 8 offers some suggestive evidence on this: intrapartum cesareans in private hospitals are most concentrated on Mondays, which

Our estimates identify the total reduced-form effect of organizational design, which operates through two channels: within-sector incentives (a higher ρ_{priv} shifting the cutoffs $\hat{\phi}_t^s$ to the right) and between-sector sorting of physicians with lower professionalism costs into private practice. Both channels are endogenous to organizational form (as formalized in Section 3.5), and both are therefore part of the organizational design effect we set out to measure.

The distinction nonetheless matters for the effectiveness of specific policy tools. If the sectoral gap is driven primarily by within-physician incentives, reforms that dilute individual schedule control can compress discretionary cesarean use in the short run without altering the composition of the physician pool. If, instead, the gap is driven primarily by sorting, comparable reforms will have attenuated effects unless they are complemented by changes in how physicians are recruited, trained, and monitored in each sector.

Data with a consistent physician identifier that tracks the same doctor across public and private hospitals would allow a direct decomposition of the two channels: within-physician variation (the same doctor attending deliveries in both sectors) identifies the incentive channel under physician fixed effects, while the distribution of behavior across physicians who practice in only one sector versus both sectors disciplines the sorting channel. Building such a linked dataset is a natural avenue for future work.

7 Conclusion

This paper has investigated how organizational design shapes physician discretion over cesarean delivery. Our setting, the Buenos Aires metropolitan area, offers a sharp contrast between individualized private maternity care and team-based public maternity care within the same labor market, with monetary incentives that do not vary with delivery mode. A simple model of physician decision-making identifies two discretionary margins, a *preference margin* and a *timing margin*, and yields three testable predictions that should be more visible in the sector with higher schedule control: stated maternal preferences translating into elective cesareans (P1), elective cesareans concentrating on working days (P2), and a working-day gradient in intrapartum cesareans (P3). The third prediction arises as a selection effect of the timing margin rather than as an independent intrapartum channel.

Our empirical findings align with all three predictions. On the preference margin, stated preferences for cesarean translate into scheduled procedures strongly in private

is consistent with the direct channel operating alongside selection. One clinical instance of this channel is the conservative management of spontaneous weekend labor until Monday, when the cesarean is then performed as an intrapartum conversion rather than as a scheduled elective procedure. However, the small number of cases per day precludes a formal test.

hospitals and have no detectable effect in public hospitals. This differential survives a sensitivity analysis for selection on unobservables. On the timing margin, elective cesareans in private hospitals cluster on Mondays and Fridays and are almost absent from weekends and public holidays, while the public sector shows no comparable pattern. Among women who enter labor, intrapartum cesareans in private hospitals are about 13 percentage points more likely on working days than on nonworking days, while the public-sector difference is small and statistically indistinguishable from zero. We read the private-sector gradient as consistent with a selection effect of the timing margin: discretionary cases that would otherwise generate weekend labors have already been moved to scheduled, working-day procedures. The two margins therefore point to the same underlying mechanism: in the private sector, discretionary cesareans operate primarily through the decision to schedule an elective procedure in advance, rather than through intrapartum conversions.

Two broader lessons emerge. First, analyses of day-of-week effects that do not distinguish scheduled from intrapartum procedures can be misleading, because the delivery date is itself an outcome of discretionary scheduling. In this setting, the timing margin is not a side detail. It is a central channel through which convenience incentives operate and selection arises. Second, the institutional environment governs whether convenience incentives meaningfully distort care. The same procedure that provides clinical benefits in appropriate cases can also become a tool for schedule management when providers internalize the value of time saved. Conversely, when production is organized around shifts and teams, convenience incentives translate less strongly into differential procedure use across the calendar.

These results have direct policy implications for reducing medically unnecessary cesareans without compromising access to clinically indicated surgery. Discretionary cesareans are not welfare-neutral relative to vaginal delivery. When a cesarean is performed for scheduling convenience rather than clinical indication, it replaces a delivery pathway that would likely have resulted in vaginal birth, exposing the mother to the additional surgical risks and longer recovery associated with cesarean section (Betrán et al., 2018) and generating higher costs for the insurer or public system (Gibbons et al., 2010), even when physician fees do not vary by delivery mode. Interventions that focus narrowly on procedure prices may be insufficient if the main distortion operates through the value of physician time and the organization of coverage. Policies that reduce the extent to which any single physician is the residual claimant of her schedule, such as shared-call arrangements, group practices, or hospitalist models, may attenuate the incentive to concentrate scheduled deliveries on boundary working days. Complementary approaches include systematic monitoring of elective cesarean use, audit and feedback at the facility level, and decision support aimed at aligning scheduled procedures with clinical indications and informed preferences. Because a portion of cesarean use is medically necessary, the relevant objective is not to minimize cesareans outright, but to reduce the discretionary component that arises from nonclinical incentives.

Several limitations point to priorities for future work. The empirical evidence relies on a small sample from six hospitals, and our ability to separate medical indication from discretion is imperfect. Because we do not observe the same physicians practicing in both sectors, we also cannot decompose whether the sectoral gap reflects within-physician incentives or between-sector sorting of physicians, a distinction that matters for the effectiveness of specific policy tools. In addition, the day of delivery is an imperfect proxy for the onset of labor, and financial arrangements between hospitals and physicians are difficult to observe directly. Richer administrative data linking clinical risk, labor onset, provider identifiers across sectors, and maternal and neonatal outcomes would allow sharper tests of the model and clearer welfare conclusions. More broadly, evaluating reforms that change weekend staffing, coverage arrangements, or reporting requirements would help quantify the extent to which institutional design can curb unnecessary cesarean use while preserving high-quality maternity care.

References

Altonji, J. G., Elder, T. E., and Taber, C. R. (2005). Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools. *Journal of Political Economy*, 113(1):151–184.

Ariovich, A., Hasicic, C., Marin, G., and Librandi, J. M. (2025). Evolución de las tasas de cesáreas en efectores públicos y privados del Instituto de Obra Médico Asistencial (2020–2023). *SALUD pública: Revista del Ministerio de Salud de la Provincia de Buenos Aires*, 1. COSAPRO 2025 special issue. ISSN 3008-7074.

Arrow, K. J. (1963). Uncertainty and the welfare economics of medical care. *American Economic Review*, 53(5):941–973.

Béhague, D. P., Victora, C. G., and Barros, F. C. (2002). Consumer demand for caesarean sections in Brazil: Informed decision making, patient choice, or social inequality? *BMJ*, 324(7343):942–945.

Belizán, J. M., Althabe, F., Barros, F. C., and Alexander, S. (1999). Rates and implications of caesarean sections in Latin America: Ecological study. *BMJ*, 319(7222):1397–1400.

Betrán, A. P., Temmerman, M., Kingdon, C., Mohiddin, A., Opiyo, N., Torloni, M. R., Zhang, J., Musana, O., Wanyonyi, S. Z., Gülmezoglu, A. M., and Downe, S. (2018). Interventions to reduce unnecessary caesarean sections in healthy women and babies. *Lancet*, 392(10155):1358–1368.

Boerma, T., Ronsmans, C., Melesse, D. Y., Barros, A. J. D., Barros, F. C., Juan, L., Moller, A.-B., Say, L., Hosseinpoor, A. R., Yi, M., de Lyra Rabello Neto, D., and Temmerman, M. (2018). Global epidemiology of use of and disparities in caesarean sections. *Lancet*, 392(10155):1341–1348.

Borra, C., González, L., and Sevilla, A. (2019). The impact of scheduling birth early on infant health. *Journal of the European Economic Association*, 17(1):30–78.

Bound, J. (1991). Self-reported versus objective measures of health in retirement models. *Journal of Human Resources*, 26(1):106–138.

Brown, H. S. I. (1996). Physician demand for leisure: implications for cesarean section rates. *Journal of Health Economics*, 15(2):233–242.

Card, D., Fenizia, A., and Silver, D. (2023). The health impacts of hospital delivery practices. *American Economic Journal: Economic Policy*, 15(2):42–81.

Clemens, J. and Gottlieb, J. D. (2014). Do physicians' financial incentives affect medical treatment and patient health? *American Economic Review*, 104(4):1320–1349.

Costa-Ramón, A., Kortelainen, M., Rodríguez-González, A., and Sääksvuori, L. (2022). The long-run effects of cesarean sections. *Journal of Human Resources*, 57(6):2048–2085.

Costa-Ramón, A. M., Rodríguez-González, A., Serra-Burriel, M., and Campillo-Artero, C. (2018). It's about time: Cesarean sections and neonatal health. *Journal of Health Economics*, 59:46–59.

Currie, J. and MacLeod, W. B. (2008). First do no harm? tort reform and birth outcomes. *Quarterly Journal of Economics*, 123(2):795–830.

de Elejalde, R. and Giolito, E. (2021). A demand-smoothing incentive for cesarean deliveries. *Journal of Health Economics*, 75:102411.

Dickert-Conlin, S. and Chandra, A. (1999). Taxes and the timing of births. *Journal of Political Economy*, 107(1):161–177.

Foo, P. K., Lee, R. S., and Fong, K. (2017). Physician prices, hospital prices, and treatment choice in labor and delivery. *American Journal of Health Economics*, 3(3):422–453.

Frakes, M. and Gruber, J. (2019). Defensive medicine: Evidence from military immunity. *American Economic Journal: Economic Policy*, 11(3):197–231.

Gans, J. S. and Leigh, A. (2009). Born on the first of july: An (un)natural experiment in birth timing. *Journal of Public Economics*, 93(1–2):246–263.

Gibbons, L., Belizán, J. M., Lauer, J. A., Betrán, A. P., Merialdi, M., Althabe, F., et al. (2010). The global numbers and costs of additionally needed and unnecessary caesarean sections performed per year: Overuse as a barrier to universal coverage. World Health Report (2010) Background Paper 30, World Health Organization.

Green, J. (1978). Physician-induced demand for medical care. *Journal of Human Resources*, 13(Supplement):21–34.

Gruber, J., Kim, J., and Mayzlin, D. (1999). Physician fees and procedure intensity: the case of cesarean delivery. *Journal of Health Economics*, 18(4):473–490.

Grytten, J., Skau, I., and Sørensen, R. (2013). Do mothers decide? The impact of preferences in healthcare. *Journal of Human Resources*, 48(1):142–168.

Halla, M., Mayr, H., Pruckner, G. J., and García-Gómez, P. (2020). Cutting fertility? effects of cesarean deliveries on subsequent fertility and maternal labor supply. *Journal of Health Economics*, 72:102325.

Hopkins, K. (2000). Are Brazilian women really choosing to deliver by cesarean? *Social Science & Medicine*, 51(5):725–740.

Jena, A. B., Seabury, S., Lakdawalla, D., and Chandra, A. (2011). Malpractice risk according to physician specialty. *New England Journal of Medicine*, 365(7):629–636.

Johnson, E. M. and Rehavi, M. M. (2016). Physicians treating physicians: Information and incentives in childbirth. *American Economic Journal: Economic Policy*, 8(1):115–141.

Kessler, D. and McClellan, M. (1996). Do doctors practice defensive medicine? *The Quarterly Journal of Economics*, 111(2):353–390.

Librandi, J. M., Fasano, M. V., Malzone, M. P., Giles, H. F., Belizán, J. M., Marin, G., and Cormick, G. (2023). Tendencia de cesáreas en entidades privadas de la provincia de Buenos Aires, Argentina (2017–2021). *Ginecología y Obstetricia de México*, 91(4):249–255.

Lo, J. C. (2003). Patients' attitudes vs. physicians' determination: implications for cesarean sections. *Social Science & Medicine*, 57(1):91–96.

Mazzoni, A., Althabe, F., Gutierrez, L., Gibbons, L., Liu, N. H., Bonotti, A. M., Izbizky, G. H., Ferrary, M., Viergue, N., Vigil, S. I., Zalazar Denett, G., and Belizán, J. M. (2016). Women's preferences and mode of delivery in public and private hospitals: a prospective cohort study. *BMC Pregnancy and Childbirth*, 16(1):34.

Mazzoni, A., Althabe, F., Liu, N. H., Bonotti, A. M., Gibbons, L., Sánchez, A. J., and Belizán, J. M. (2011). Women's preference for caesarean section: A systematic review and meta-analysis of observational studies. *BJOG*, 118(4):391–399.

McGuire, T. G. (2000). Physician agency. In Culyer, A. J. and Newhouse, J. P., editors, *Handbook of Health Economics*, volume 1, chapter 9, pages 461–536. Elsevier, Amsterdam.

Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, 37(2):187–204.

Potter, J. E., Berquó, E., Perpétuo, I. H. O., Leal, O. F., Hopkins, K., Souza, M. R., and Formiga, M. C. d. C. (2001). Unwanted caesarean sections among public and private patients in Brazil: Prospective study. *BMJ*, 323(7322):1155–1158.

Rosati, G. F., Olego, T. A., and Vazquez Brust, H. A. (2020). Building a sanitary vulnerability map from open source data in Argentina (2010–2018). *International Journal for Equity in Health*, 19:206.

Sandall, J., Tribe, R. M., Avery, L., Mola, G., Visser, G. H. A., Homer, C. S. E., Gibbons, D., Kelly, N. M., Kennedy, H. P., Kidanto, H., Taylor, P., and Temmerman, M. (2018). Short-term and long-term effects of caesarean section on the health of women and children. *The Lancet*, 392(10155):1349–1357.

Schulkind, L. and Shapiro, T. M. (2014). What a difference a day makes: quantifying the effects of birth timing manipulation on infant health. *Journal of Health Economics*, 33:139–158.

Spetz, J., Smith, M. W., and Ennis, S. F. (2001). Physician incentives and the timing of cesarean sections: Evidence from california. *Medical Care*, 39(6):536–550.

Triunfo, P. and Rossi, M. (2009). The effect of physicians' remuneration system on the caesarean section rate: The Uruguayan case. *International Journal of Health Care Finance and Economics*, 9(4):333–345.

Tussing, A. D. and Wojtowycz, M. A. (1997). Malpractice, defensive medicine, and obstetric behavior. *Medical Care*, 35(2):172–191.

Villar, J., Valladares, E., Wojdyla, D., Zavaleta, N., Carroli, G., Velazco, A., Shah, A., Campodónico, L., Bataglia, V., Faundes, A., et al. (2006). Caesarean delivery rates and pregnancy outcomes: The 2005 WHO global survey on maternal and perinatal health in Latin America. *Lancet*, 367(9525):1819–1829.

Wachter, R. M. and Goldman, L. (1996). The emerging role of "hospitalists" in the American health care system. *New England Journal of Medicine*, 335(7):514–517.

World Health Organization (2015). WHO statement on caesarean section rates. https://www.who.int/reproductivehealth/publications/maternal_perinatal_health/cs-statement/en/.

A Appendix: Tables

Table A1. Model notation

Symbol	Description
d	Biological onset day.
D	Delivery day.
$t \in \{0, 1, 2\}$	Calendar category: regular, boundary, or nonworking day.
$t(\cdot)$	Mapping from a calendar date to its category.
ω_t	Opportunity-cost multiplier on a day of category t.
$\alpha > 1$	Premium on physician time on nonworking days.
$L_N > L_C$	Time required for a vaginal and a cesarean delivery.
$\tau \in \{C, N\}$	Maternal preference type: cesarean or vaginal.
$\theta \in [0, 1]$	Share of women who prefer cesarean.
$\xi \in (0, 1)$	Share of deliveries that are medically indicated.
$s \in \{\text{pub}, \text{priv}\}$	Sector: public or private.
$\rho_{\text{priv}} > \rho_{\text{pub}}$	Marginal value the physician places on time saved in sector s.
$\phi \in [0, \bar{\phi}]$	Professionalism cost of performing a non-indicated cesarean.
$\gamma > 0$	Additional professionalism cost when the woman prefers vaginal.
U^s	Physician utility in sector s.
S^s_w	Time-saving gain from a cesarean on a working onset day.
S^s_2	Time-saving gain on a nonworking onset day.
$T(\tau, \phi)$	Total professionalism cost on a preference-τ woman.
$\hat{\phi}^s_t(\tau)$	Physician-type cutoff in sector s at onset category t for type τ.
p_2, p_w	Probability that onset falls on a nonworking or working day.
$X^s_{t,\tau}$	Net gain from performing a non-indicated cesarean.
$R^s_t(\phi)$	Expected per-case rent on a type-t onset day, averaged over maternal preference τ.
$V^s(\phi)$	Expected per-delivery discretionary payoff in sector s.
B_s	Sector-specific non-discretionary amenity.
κ	Relative amenity of the public sector.
$\Delta V(\phi)$	Sorting index, weakly decreasing in ϕ.
ϕ^*	Sorting threshold between public and private sectors.
$\widetilde{V}^s(\phi)$	Expected per-delivery rent including indicated-case rescheduling (Section 3.6).
$\widetilde{\Delta V}(\phi)$	Sorting index with indicated cases, equal to $\Delta V(\phi)$ plus a positive sector constant.
$\widetilde{\phi}^*$	Sorting threshold with indicated cases, with $\widetilde{\phi}^* \geq \phi^*$.

Notes: Summary of variables and parameters used in Section 3.

Table A2. Births by working versus nonworking days

	Working days	Weekends and public holidays	Total
Total calendar days	308	151	459
	67.10%	32.90%	100.00%
All hospitals			
Births	289	90	379
	76.25%	23.75%	100.00%
Average births per day	0.938	0.596	0.826
Public hospitals			
Births	141	56	197
	71.57%	28.43%	100.00%
Average births per day	0.458	0.371	0.429
Private hospitals			
Births	148	34	182
	81.32%	18.68%	100.00%
Average births per day	0.481	0.225	0.397

Notes: The study window spans 459 consecutive days (25 September 2010 to 27 December 2011). Nonworking days are Saturdays, Sundays, and national public holidays.

Figure A1. Vaginal deliveries by day of the week and hospital sector

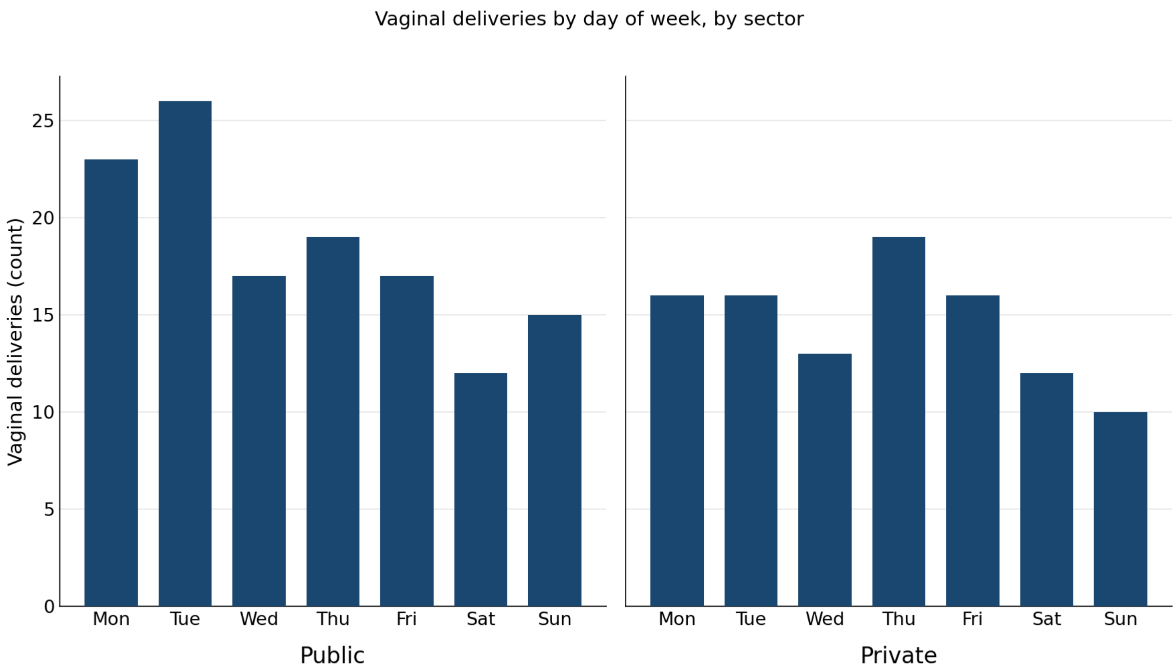

Notes: Counts of vaginal deliveries ($N = 231$) by day of the week and hospital sector. Pearson $\chi^2(6) = 2.07$, $p = 0.91$ for the day-of-week by sector contingency test. Within-sector uniformity across days of the week is also not rejected (public: $\chi^2(6) = 7.36$, private: $\chi^2(6) = 3.82$, both well below the 5% critical value of 12.59).

Table A3. Maternal characteristics by hospital sector

All hospitals	N	mean	sd	min	max
Age	379	26.214	5.375	18	36
Education					
Elementary school	378	0.093	0.290	0	1
High school	378	0.458	0.499	0	1
Tertiary or university	378	0.450	0.498	0	1
Married or in couple	378	0.812	0.391	0	1
Working	375	0.477	0.500	0	1
Insured	379	0.464	0.499	0	1
Prefers CS	367	0.074	0.261	0	1
Had CS	379	0.391	0.489	0	1
Public hospitals					
Age	197	22.792	4.104	18	35
Education					
Elementary school	196	0.179	0.384	0	1
High school	196	0.714	0.453	0	1
Tertiary or university	196	0.107	0.310	0	1
Married or in couple	196	0.699	0.460	0	1
Working	197	0.157	0.365	0	1
Insured	197	0.000	0.000	0	0
Prefers CS	197	0.081	0.274	0	1
Had CS	197	0.345	0.477	0	1
Private hospitals					
Age	182	29.918	3.946	18	36
Education					
Elementary school	182	0.000	0.000	0	0
High school	182	0.181	0.386	0	1
Tertiary or university	182	0.819	0.386	0	1
Married or in couple	182	0.934	0.249	0	1
Working	178	0.831	0.375	0	1
Insured	182	0.967	0.179	0	1
Prefers CS	170	0.065	0.247	0	1
Had CS	182	0.440	0.498	0	1

Notes: "Prefers CS" is measured in the third-trimester survey. Insurance coverage is nearly perfectly correlated with hospital type, reflecting the institutional organization described in Section 2.

Table A4. Elective cesarean: predicted probabilities by preferences and hospital sector

	(1)	(2)
Prefers vaginal delivery, public hospital	0.0452***	0.0527***
	(0.0156)	(0.0178)
Prefers vaginal delivery, private hospital	0.115***	0.118***
	(0.0251)	(0.0245)
Prefers cesarean delivery, public hospital	0.0561	0.0570
	(0.0527)	(0.0526)
Prefers cesarean delivery, private hospital	0.480***	0.512***
	(0.152)	(0.113)
Observations	367	338
Hospital fixed effects	NO	YES

Notes: Predicted probabilities are average predictive margins computed from the probit model in Table 3. Standard errors (delta method) in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A5. Intrapartum cesarean: probit estimates restricted to spontaneous labors

Variable	(1) *ICS*	(1) *ICS*	(2) *ICS*	(2) *ICS*	(3) *ICS*	(3) *ICS*
Private hospital (PrH)	-0.348	–	-0.327	–	-0.378	–
	(0.346)	–	(0.318)	–	(0.329)	–
Working day	-0.508**	-0.558**	-0.492**	-0.559**	-0.479**	-0.510**
	(0.236)	(0.243)	(0.229)	(0.237)	(0.232)	(0.238)
PrH × Working day	0.702*	0.772*	0.719**	0.809**	0.766**	0.818**
	(0.373)	(0.395)	(0.351)	(0.367)	(0.359)	(0.376)
Age	0.0494	0.0619	0.0463	0.0583	0.0480	0.0592
	(0.063)	(0.064)	(0.064)	(0.064)	(0.064)	(0.064)
Age squared	-0.00253	-0.00346	-0.00229	-0.00319	-0.00235	-0.00321
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Prefers CS	0.118	0.182	0.0797	0.132	0.133	0.196
	(0.404)	(0.417)	(0.387)	(0.389)	(0.406)	(0.417)
PrH × Prefers CS	–	–	–	–	–	–
Constant	-0.577**	-0.902**	-0.604**	-0.890**	-0.612**	-0.922**
	(0.264)	(0.381)	(0.256)	(0.380)	(0.258)	(0.381)
Observations	275	272	275	272	275	272
Hospital fixed effects	NO	YES	NO	YES	NO	YES
Convenience definition	(1)	(1)	(2)	(2)	(3)	(3)

Notes: Sample is restricted to women who experienced spontaneous onset of labor ($N = 275$; induced labors dropped). "Working day" indicates a delivery classified as working under the corresponding convenience definition. The PrH × Prefers CS interaction is omitted across all columns because the two cases in the cell (Private × Prefers CS) predict failure perfectly in the spontaneous subsample. Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A6. Elective cesarean predicted probabilities by calendar category and sector

	Boundary (Mon/Fri)	Interior (Tue–Thu)
Private hospitals	0.218	0.140
	(0.044)	(0.036)
Public hospitals	0.038	0.095
	(0.024)	(0.033)

Notes: Predicted probabilities of elective cesarean from the probit specification in Column (1) of Table 4, with delta-method standard errors in parentheses. Sample is restricted to working days under convenience definition (1) ($N = 271$). The cross-sector difference-in-differences between boundary and interior working days is $(0.218 - 0.140) - (0.038 - 0.095) = 0.135$, or 13.5 percentage points.

Table A7. Predicted probability of intrapartum cesarean on working versus nonworking days

	Predicted probability
Public hospitals	
Working day	0.283***
Nonworking day	0.360***
Difference (working minus nonworking)	-0.077
Private hospitals	
Working day	0.357***
Nonworking day	0.224***
Difference (working minus nonworking)	0.133

Notes: Predicted probabilities are average predictive margins computed from the probit specification in Column (1) of Table 5 (without hospital fixed effects). A working day is defined using convenience definition (1). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A8. Convenience definitions

	Day	(1)	(2)	(3)
Normal week	Tuesday	x	x	x
	Wednesday	x	x	x
	Thursday	x	x	x
	Friday	x	x	x
	Saturday			
	Sunday			
	Monday	x	x	x
Week with long weekend (holiday on Monday)	Tuesday	x	x	x
	Wednesday	x	x	x
	Thursday	x	x	x
	Friday	x		
	Saturday			
	Sunday			
	Monday			
Week with midweek holiday (e.g., Wednesday)	Tuesday	x		x
	Wednesday			
	Thursday	x	x	x
	Friday	x	x	x
	Saturday			
	Sunday			
	Monday	x	x	x

Notes: "x" indicates that the day is classified as working under the corresponding definition, and shaded cells indicate nonworking days. Definition (1) treats weekends and public holidays as nonworking. Definition (2) additionally treats the day immediately preceding a public holiday as nonworking (when a holiday falls on Monday, the preceding day is the prior Friday). Definition (3) treats only the Friday preceding a Monday public holiday as nonworking, without reclassifying the day preceding midweek holidays.

Table A9. Number and type of deliveries per hospital

Type	Name	Mode of Delivery CS	VD	Total
Public	Hospital Magdalena V. de Martínez	24	53	77
	Hospital Materno Infantil Dr. Carlos Gianantonio	39	54	93
	Hospital Comunal de Tigre	4	22	26
	Other public	1	0	1
	Subtotal	*68*	*129*	*197*
Private	Hospital Italiano de Buenos Aires	43	46	89
	Centro de Educación Médica e Investigaciones Clínicas "Norberto Quirno" (CEMIC)	17	20	37
	Hospital Británico de Buenos Aires	20	33	53
	Others	0	3	3
	Subtotal	*80*	*102*	*182*
	At home	0	1	1
	Total	**148**	**232**	**380**

Notes: The table reports all 380 deliveries in the cohort after follow-up exclusions. The single home delivery is excluded from the analytic sample, which comprises the 379 hospital deliveries used throughout the paper.

Table A10. Mode of delivery among women who enter labor

Mode of delivery	Public hospitals	Private hospitals	Total
Vaginal delivery	69.4%	65.8%	67.7%
Intrapartum cesarean	30.6%	34.2%	32.3%
N (labor sample)	186	155	341

Notes: Percentages are computed within each column over women who delivered either vaginally or by intrapartum cesarean, and the last row reports the corresponding sample sizes. Two urgent pre-labor cesareans (both in public hospitals) fall into neither category and are excluded from the labor sample.

Table A11. Health vulnerability in the catchment area of each study hospital

Hospital	Location	Sanitary Vulnerability Index		
		Nearest block	2 km buffer	5 km buffer
Public Hospital 1	Tigre (GBA)	0.036	0.274	0.449
Public Hospital 2	San Isidro (GBA)	0.002	0.142	0.214
Public Hospital 3	Tigre (GBA)	0.327	0.265	0.360
Private Hospital 1	CABA, Comuna 5	0.065	0.111	0.170
Private Hospital 2	CABA, Comuna 12	0.305	0.243	0.196
Private Hospital 3	CABA, Comuna 4	0.120	0.214	0.185

Notes: The Sanitary Vulnerability Index (IVS) is a census-block level indicator developed by Rosati et al. (2020) and published by Fundación Bunge y Born. It combines socioeconomic conditions from the 2010 Census (materials, overcrowding, schooling, employment) with routed travel distances to the nearest health facility (2018 hospital and health-center registries plus the OpenStreetMap road network). Values are expressed as percentiles from 0 (lowest vulnerability, best access) to 1 (highest); the national median is 0.5 by construction. "Nearest block" is the IVS of the single census-block centroid closest to the hospital. "2 km buffer" and "5 km buffer" are unweighted means of the IVS across all census-block centroids within the corresponding buffer around each hospital (total of 52,406 census blocks countrywide). Hospitals are listed in the same order as in Table A9.

Figure A2. Single-crossing of $\Delta V(\phi)$ under parameter variation

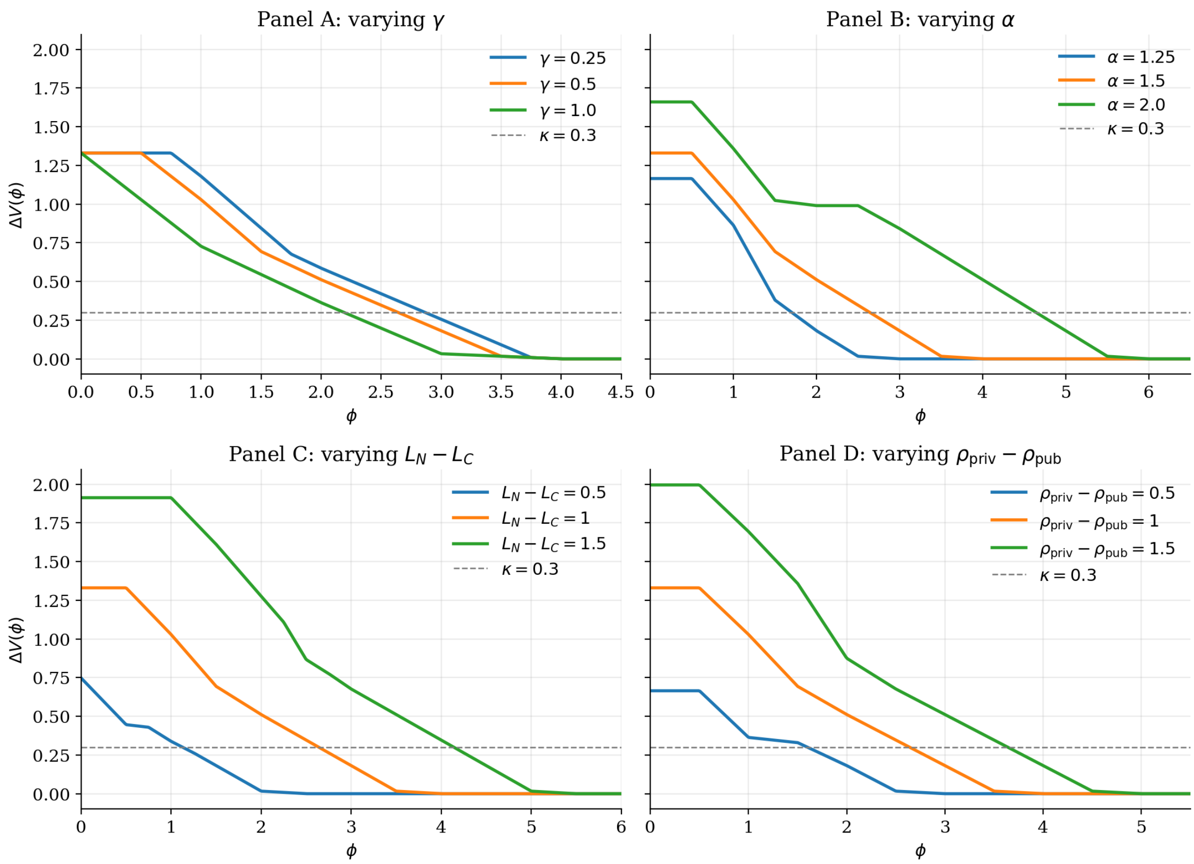

Notes: Each panel plots $\Delta V(\phi) = V^{\text{priv}}(\phi) - V^{\text{pub}}(\phi)$ against the physician's professionalism cost ϕ, for three values of one parameter while holding the others at their base values ($\rho_{\text{priv}} = 2$, $\rho_{\text{pub}} = 1$, $\alpha = 1.5$, $L_N = 2$, $L_C = 1$, $\gamma = 0.5$, $\theta = 0.1$, $p_2 = 0.33$). Panel D varies ρ_{priv} while holding $\rho_{\text{pub}} = 1$ fixed, so the label $\rho_{\text{priv}} - \rho_{\text{pub}}$ tracks the sector gap. The dashed horizontal line marks an illustrative compensating differential $\kappa = 0.3$; its single-crossing with $\Delta V(\phi)$ pins down the marginal type ϕ^* whenever a solution exists. $\Delta V(\phi)$ is weakly decreasing across all parameterizations (Proposition 1).

B Appendix: Overall cesarean use

This appendix reports a specification that combines both margins into a single regression for the probability of any cesarean delivery, treating delivery-day convenience and stated preference symmetrically as right-hand-side variables. Because the timing of elective cesareans is itself part of the scheduling decision, the working-day interaction in this specification mechanically combines P2 (scheduled interventions shifted to working days) with P3 (selection into working-day intrapartum labor), and we therefore report it only as a summary rather than as the primary test of either implication.

Table B1 reports probit estimates for the probability of cesarean delivery as a function of hospital sector, delivery-day convenience, and their interaction, under the three convenience definitions introduced in Section 5.2. Across specifications, the interaction between delivery on a working day and private hospitals is positive and statistically significant.

Table B1. Cesarean delivery: probit estimates by convenience definition

Variable	(1) *CS*	(2) *CS*	(3) *CS*	(4) *CS*	(5) *CS*	(6) *CS*
Private hospital (PrH)	-0.580*	-0.606**	-0.652**	–	–	–
	(0.324)	(0.295)	(0.306)	–	–	–
Working day	-0.0882	-0.108	-0.0866	-0.139	-0.151	-0.123
	(0.204)	(0.195)	(0.199)	(0.206)	(0.198)	(0.201)
PrH × Working day	0.651*	0.729**	0.762**	0.750**	0.786**	0.831**
	(0.340)	(0.315)	(0.323)	(0.348)	(0.322)	(0.330)
Prefers CS	0.279	0.278	0.281	0.291	0.289	0.294
	(0.339)	(0.338)	(0.339)	(0.332)	(0.330)	(0.332)
PrH × Prefers CS	0.464	0.428	0.423	0.491	0.449	0.446
	(0.526)	(0.526)	(0.526)	(0.513)	(0.513)	(0.512)
Age (centered at 18)	0.0327	0.0321	0.0306	0.0277	0.0263	0.0254
	(0.0510)	(0.0511)	(0.0510)	(0.0516)	(0.0517)	(0.0516)
Age squared	5.28e-05	0.000108	0.000265	-4.64e-05	6.20e-05	0.000188
	(0.00283)	(0.00284)	(0.00284)	(0.00284)	(0.00285)	(0.00285)
Constant	-0.522**	-0.513**	-0.524**	-1.073***	-1.063***	-1.078***
	(0.233)	(0.224)	(0.226)	(0.373)	(0.371)	(0.371)
Observations	367	367	367	364	364	364
Hospital fixed effects	NO	NO	NO	YES	YES	YES
Convenience definition	(1)	(2)	(3)	(1)	(2)	(3)

Notes: "Working day" indicates a delivery classified as working under the corresponding convenience definition. Columns (1) to (3) report specifications without hospital fixed effects, and columns (4) to (6) include hospital fixed effects. Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Figures B1a and B1b plot predicted probabilities of any cesarean delivery by age, hospital sector, and either preferences or convenience. The separation between preference types and between calendar types is wider in the private sector in both panels, broadly consistent with the elective-cesarean results in the main text. Because these figures pool elective and intrapartum cesareans, the correspondence is not exact. In particular, among women who prefer vaginal delivery the private any-cesarean curve lies slightly below the public one, because for this group the higher private elective rate is outweighed by the lower private intrapartum baseline (Section 6.2).

Figure B1. Predicted probability of cesarean delivery by age, preferences, convenience, and hospital sector

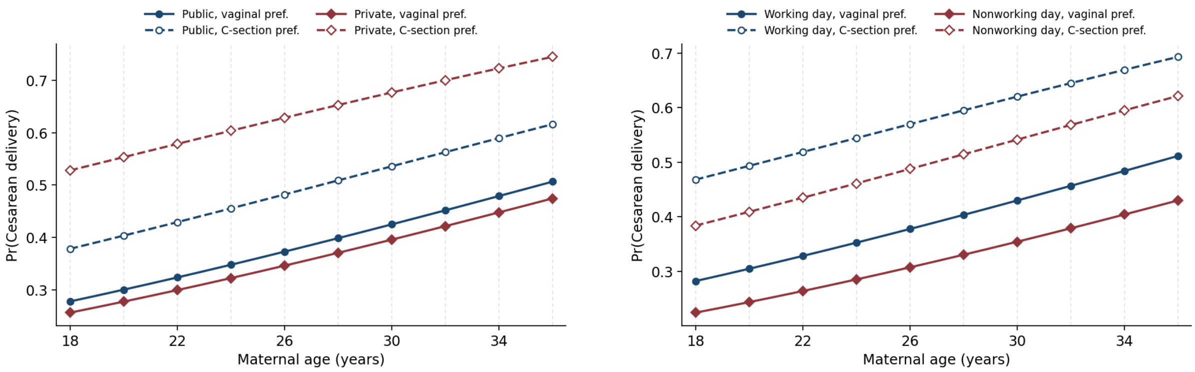

(a) By age, preferences, and hospital sector

(b) By age, preferences, and convenience

Notes: Predicted probabilities of any cesarean delivery by maternal age, computed from the probit specification in Column (1) of Table B1 (without hospital fixed effects) under convenience definition (1). Panel (a) plots predicted probabilities by hospital sector and stated preference, Panel (b) by delivery-day convenience and stated preference. Sample: full analytic sample with a recorded specific preference ($N = 367$).

C Appendix: Oster (2019) sensitivity analysis

The interpretation of the interaction PrefCS $\times$ Private in Table 3 as evidence that private-sector physicians accommodate maternal preferences rests on the assumption that the sector gradient is not driven by selection on unobservables. Women delivering in private hospitals may differ from those in public hospitals along dimensions, such as assertiveness, prior medical knowledge, or the strength of the patient-physician relationship, that are correlated with both sector choice and the probability of obtaining a cesarean, and that are only imperfectly captured by age, stated preferences, and hospital fixed effects. This appendix evaluates the extent to which selection on such unobservables could plausibly overturn the estimated interaction, following the coefficient-bounds approach of Oster (2019).

The method exploits the joint movement of the coefficient of interest and of the regression R^2 as observable controls are added. Under the assumption that selection on unobservables is proportional to selection on observables, with proportionality parameter δ, and under a user-specified maximum achievable R^2 denoted $R_{\max}$, the method yields a bias-adjusted coefficient β^* and a breakdown value of δ at which the coefficient would be driven to zero. A breakdown value exceeding one indicates that unobservables would have to be more important than the full set of observed covariates to nullify the result. Following the recommendation of Oster (2019), we report results for $\delta = 1$ and $R_{\max} = 1.3\tilde{R}^2$, where $\tilde{R}^2$ is the R^2 of the controlled regression.

Because the standard implementation is designed for linear models, we estimate the sensitivity on a linear probability analogue of the elective cesarean equation rather than on the probit specification reported in Table 3. The dependent variable is an indicator for elective cesarean and the coefficient of interest is the interaction PrefCS $\times$ Private. We report three specifications of increasing richness. Specification (A) is a baseline including hospital fixed effects, a quadratic in maternal age, and a main effect for the stated preference. Specification (B) additionally includes an indicator for having completed a tertiary or university degree. Specification (C) further interacts this education indicator with the stated preference for cesarean, which we view as a direct proxy for the component of patient agency that operates within the preference channel. In each specification, the uncontrolled benchmark against which the controlled coefficient is compared is a regression of elective cesarean on the interaction alone (PrefCS $\times$ Private, with no other regressors), so its coefficient and R^2 are common across the three columns.

Table C1 reports the results. In specifications (A) and (B), the bias-adjusted coefficient β^* is positive and of a substantive magnitude, and the breakdown δ is approximately 1.53, so unobservable selection would have to be more than fifty percent larger than the full set of observed covariates to overturn the private-sector differential. The inclusion of the completed-degree indicator in specification (B) has essentially no additional effect on the coefficient or on R^2, which reflects the fact that education is strongly correlated with sector (Table A3), so the hospital fixed effects absorb its between-hospital component and little within-hospital variation remains to predict elective cesarean. In specification (C), adding the interaction between education and stated preference *increases* the controlled coefficient rather than decreasing it. Under Oster's proportionality assumption, selection on unobservables resembling this education proxy would push the estimate away from zero rather than toward it, though this does not by itself rule out other forms of selection.

Table C1. Oster (2019) Sensitivity Analysis: PrefCS × Private

	(A) Baseline	(B) + Educ	(C) + Educ × Pref
LPM estimates			
$\hat{\beta}$ (uncontrolled)	0.469	0.469	0.469
R^2 (uncontrolled)	0.078	0.078	0.078
$\hat{\beta}$ (controlled)	0.400	0.400	0.486
R^2 (controlled)	0.171	0.171	0.173
Oster (2019) results ($\delta = 1$, $R_{\max} = 1.3\tilde{R}^2$)			
β^*	0.289	0.291	0.496
Breakdown δ	1.531	1.533	∞[a]
Observations	363	363	363
Hospital FE	Yes	Yes	Yes

Notes: Dependent variable: elective cesarean (ECS). All models estimated by LPM with robust standard errors and hospital fixed effects. All columns, including the uncontrolled benchmark, are estimated on the same sample. The coefficient of interest is PrefCS × Private in all columns. Specification (C) additionally controls for Educ × PrefCS. β^* is the bias-adjusted estimate assuming $\delta = 1$ and $R_{\max} = 1.3\tilde{R}^2$ (Oster, 2019). The breakdown δ is the degree of selection on unobservables (relative to observables) required to drive the coefficient to zero. Values of $\delta > 1$ indicate that unobservables would have to be more important than all observed covariates to nullify the result. Oster bounds for (A) and (B) use the exact quadratic formula. For (C) the linear approximation (Oster, 2019, eq. 4) is used. [a]The coefficient *increases* when adding controls, so observable selection works against the coefficient, implying the true effect is even larger.