Paying for Progress: Innovation Diffusion Under Regulation

By Zakaria El Amrani and Anastasiia Evdokimova

Northwestern University, Yale University

[preliminary draft, please do not circulate]

November 3, 2026

Latest version

ABSTRACT. A regulated price breaks the link between the value a technology creates and the return to adopting it. We study what motivates adoption when the price does not reward it and how a planner can use that motive, taking transcatheter aortic valve replacement (TAVR) under Medicare as the example. In 100 percent Medicare claims, TAVR cost more than surgical replacement but was paid the same rate until 2015, yet half of eventual adopters entered earlier. Demand estimates show that adoption raises the value of the hospital itself, even to patients ineligible for TAVR: an adopter is worth 22 minutes of travel in 2012 and over an hour by 2019. Competition for patients, not the price, finances early entry. The 2015 increase raised volume at existing programs by 8 percent, but advancing it to 2012 adds 0.4 percent more hospitals, at $55 million each. The delayed rate sorts hospitals by private return, so delaying the increase and spending the saving on targeted lump sums and travel vouchers raises adopters from 619 to 634, or TAVR patients by 13 percent, at equal outlay.

KEYWORDS. Healthcare innovation, technology diffusion.

JEL CODES. I110, I180, D83, D7.

I. Introduction

A regulated price breaks the link between the value a technology creates and the return to adopting it. Where prices are set in the market, the adopter can charge for part of the value it creates, so its return moves with that value and adoption follows profitability (Griliches, 1957); competition then shapes the timing, as when cannibalization delays incumbents' innovation (Igami, 2017). Where a regulator sets one payment for every provider and revises it on a schedule that lags the technology, a new technology can be paid the incumbent's rate

CONTACT. zakaria.elamrani@kellogg.northwestern.edu, anastasiia.evdokimova@yale.edu.

for years although it costs more. The level and structure of regulated payment shape which technologies providers adopt (Finkelstein, 2007; Acemoglu and Finkelstein, 2008; Chandra and Skinner, 2012; League, 2023); while the new technology earns less per case than the one it replaces, adoption must be financed by a return outside the payment. We show that competition for patients provides that return: a regulated price fixes what a hospital earns per patient but not the number of patients it treats, and offering the new technology increases that number.

The same return determines how a planner should intervene. A planner who seeks faster diffusion faces a limited budget and limited information. Raising the regulated price is the natural instrument, but whether a higher rate induces adoption or rewards adoption that would have occurred regardless depends on what each provider earns beyond the payment. The question is therefore not only how much to pay and when, but also whom to pay, hospitals or the patients who travel to them, and in what form; the adoption-subsidy literature finds large differences in cost-effectiveness across these choices (Xing et al., 2021; De Groote and Verboven, 2019; Springel, 2021; Langer and Lemoine, 2022).

This paper asks two questions. First, what motivates early adoption, net of the direct financial stimulus? Second, can a planner with the same budget reach more hospitals offering the technology, or more patients treated, than the observed policy? The two objectives are treated separately. To answer both we use transcatheter aortic valve replacement (TAVR) under Medicare as the example: on 100 percent Medicare claims we estimate patient demand for hospitals, which gives the return to adoption, and an adoption model, which gives its cost. The setting has three features that make it well suited: three years in which TAVR was paid the rate of the incumbent procedure although it cost more, a payment revision on a date common to all hospitals, and claims that record every patient's choice of hospital.

We document four key findings. First, the timing of adoption is not the one financial incentives predict: more than two fifths of the hospitals that chose to adopt entered while TAVR was paid the surgical rate although it cost more,[1] and entry did not jump when Medicare raised the payment. Second, neither the price nor learning by doing, which shortens stays but does not lower the cost of the device, explains this pattern; the evidence points to patient steering. Third, adoption raises the value patients place on the hospital, not only on the procedure: patients who are not eligible for TAVR also prefer an adopter, and adoption is worth about 22 minutes of travel to the median patient in 2012 and 75 minutes

[1]Counting the 97 pivotal-trial sites, which ran TAVR programs before Medicare coverage, half of eventual adopters entered before the revision.

by 2019. Fourth, an earlier reimbursement increase acts on volume far more than on entry: TAVR volume at programs already operating rose by about 8 percent when the increase came in 2015, while in the estimated model moving the increase from 2015 to 2012 adds 2.2 hospitals, 0.4 percent of adopters, at about $55 million each. Because the delayed rate sorts hospitals by the steering return, a planner who reallocates the payment to targeted lump sums and travel vouchers reaches more hospitals, 634 against 619, or treats more patients, 22.1 against 19.6 thousand a year, than the observed policy at the same budget.

TAVR replaces the aortic valve through a catheter and substitutes for open-heart surgical replacement (SAVR). It costs a hospital more than SAVR, primarily because of the device (Reddy et al., 2023; McCarthy et al., 2017). Medicare covered TAVR from fiscal 2012 but paid roughly the surgical rate until fiscal 2015, when it created separate payment groups for transcatheter valve replacement (Dranove et al., 2022). Hospitals could bill the procedure throughout, so for three years TAVR was billable at the rate of the procedure it replaced while costing more. We observe this in 100 percent Medicare inpatient claims for 2010 to 2020. TAVR is performed almost entirely on patients 65 and older: the median patient in the national registry was 84 in 2011–13 and 80 in 2019 (Mack et al., 2013; Carroll et al., 2020). Fee-for-service Medicare claims therefore cover nearly every procedure outside Medicare Advantage, and because Medicare sets both what hospitals receive and what patients pay, neither price varies across hospitals.

The three years in which TAVR was paid the SAVR rate identify the first question. A hospital that entered then earned less on a TAVR case than on a SAVR case, so early entry required a return the payment did not provide. Hospitals do respond to the payment: TAVR volume at existing adopters rose by about 8 percent at the revision. Learning by doing is present, but it operates on length of stay rather than on the device, and hospitals with no local rival enter no earlier than others. The evidence points instead to steering. Patients travel farther to reach an adopter, including patients who subsequently receive SAVR, and a hospital's adoption hazard declines as its rivals adopt.

To measure what steering is worth and whether it comes from the procedure or from the hospital that offers it, we estimate patient demand for hospitals. We model demand for the hospital rather than the procedure because the heart team, not the patient, assigns the procedure after admission. Because Medicare fixes the price and standardizes cost sharing, a patient expresses her preference for an adopter through how far she travels, so travel time is the operative price (Gaynor et al., 2016). Variation in eligibility separates the two sources. Coverage expanded by surgical-risk tier on dates common to all hospitals, and the procedure

is unaffordable outside coverage, so at any date the patient groups defined by their risk score below the threshold had close to no chance of receiving TAVR; their preference for an adopter is a preference for the hospital, which we interpret as a signal of quality, the role that models of hospital competition assign to new technology (Luft et al., 1986; Devers et al., 2003; Horn et al., 2022). A mixed logit with hospital fixed effects and a random preference for treatment (Berry et al., 1995) gives a travel disutility flatter than in studies of emergency and maternity admissions (Chandra et al., 2016; Ho and Pakes, 2014; Raval and Rosenbaum, 2021), as expected for a planned, referral-driven procedure: each ten minutes of driving lowers the odds of choosing a hospital by about 16 percent. It also shows that patients value an adopter well below the eligibility threshold, from a risk score of 0.6 percent in fiscal 2013 against a threshold of 8 percent, and at every risk level from fiscal 2015. Adoption therefore brings a hospital two groups of patients: eligible patients who travel for the procedure, and patients who are not eligible for TAVR but choose the hospital that offers it.

The demand estimates give the return to adoption; an adoption model gives its cost. The model has two stages: hospitals choose when to enter, or never, given the paths of rivals, and, once they operate a program, how to assign admitted patients between TAVR and SAVR. We model entry as a choice among adoption years under perfect foresight of the payment path, as in models of the timing of hospital technology adoption (Schmidt-Dengler, 2025). The value of entering in a given year combines the estimated demand, the observed reimbursement path and cost-to-charge margins, and the model recovers the fixed cost of adoption in dollars: \$5.0M for the median adopter, in the range of industry estimates for a program that requires a hybrid operating room, about \$1M for the program and \$3–5M for the room (Tangorra, 2015; ECRI Institute, 2013; Fornell, 2011), and \$9.4M for the median hospital that never adopts.

Demand and adoption costs together let us ask whether a planner can improve on the observed policy at the same budget. We focus on three instruments. The first is the timing of the rate increase. An earlier update raises volume at hospitals that already operate a program but has little effect on entry, since a hospital whose fixed cost exceeds its return at the revised rate stays out whenever that rate arrives. The second is a lump sum paid at the update to hospitals that have not entered. Its rationale is fiscal: the delayed rate has already separated hospitals with a return to early entry from the rest, so the transfer can exclude hospitals that would enter without it. The third is a travel voucher for patients who receive TAVR. If the objective is patients treated, it treats patients at existing programs rather than waiting for a local hospital to adopt, and it avoids subsidizing hospitals that remain outside

for reasons unrelated to cost. By enlarging the pool an adopter can capture, the voucher also intensifies competition for patients and so raises the return to adoption indirectly. At the observed outlay of $502M, the observed schedule leaves 619 hospitals offering TAVR by 2019 and 19.6 thousand TAVR patients a year. The same budget buys 634 hospitals with the rate rise delayed to 2018 and the saving spent on a targeted lump sum and the voucher, or 22.1 thousand patients a year with the rise delayed to 2016. The rise is delayed rather than advanced because the rate is paid on every case at every adopter, so moving it is almost entirely inframarginal for entry: advancing it by three years costs about $55M per additional hospital, against about $10.5M for targeted lump sums and the voucher. Most of the gain in patients treated comes from the voucher, which moves patients to existing programs: with the rise delayed to 2019 and no lump sum, it reaches 21.7 thousand patients a year for a quarter of the observed outlay.

Related literature. The paper contributes to three literatures. The first studies how regulated reimbursement shapes hospitals' technology choices (Finkelstein, 2007; Acemoglu and Finkelstein, 2008; Chandra and Skinner, 2012; Dranove et al., 2022) and how sharply providers respond to it (Dafny, 2005; Clemens and Gottlieb, 2014; Eliason et al., 2018; Einav et al., 2018; Alexander and Schnell, 2024; Cabral et al., 2025); within TAVR, prior work covers learning and volume requirements, physician incentives, geographic access and hospital margins (Yang, 2024; Vemulapalli et al., 2019; Hoagland, 2024; Groeneveld et al., 2020; Strobel et al., 2019; Nathan et al., 2021; Damluji et al., 2020; Sakowitz et al., 2025; McCarthy et al., 2017; Reddy et al., 2025). We study the margin this work leaves open, diffusion while the regulated price pays the new technology no more than the one it replaces, and show that adoption pays through the patients it brings to the hospital's whole program.

The second is non-price competition and structural models of adoption. Hospitals under regulated prices compete on quality and technology (Robinson and Luft, 1985; Dranove et al., 1992; Kessler and McClellan, 2000; Cooper et al., 2011; Gaynor et al., 2013; Bloom et al., 2015; Gaynor et al., 2015), and new technologies reallocate patients toward adopters (Aggarwal et al., 2017; Chandra et al., 2016; Horn et al., 2022). We link the hospital-choice literature (Capps et al., 2003; Ho and Pakes, 2014; Gaynor et al., 2016) to models of adoption and entry timing (Reinganum, 1981; Fudenberg and Tirole, 1985; Ackerberg and Gowrisankaran, 2006; Ryan, 2012; Dunne et al., 2013; Igami, 2017; Schmidt-Dengler, 2025) by estimating the value of adoption from patient choice and embedding it in the timing model. The demand response extends to patients who are not eligible for the new procedure, so the hospital-level return to innovation exceeds the value generated by its use.

The third is adoption-subsidy design, where the central problem is inframarginal spending and the form, targeting and timing of a subsidy determine its cost-effectiveness (Xing et al., 2021; De Groote and Verboven, 2019; Sheldon and Dua, 2019; Hughes and Podolefsky, 2015; Springel, 2021; Langer and Lemoine, 2022). We show that a delayed rate screens hospitals by private return, in the spirit of screening by ordeal and regulation under asymmetric information (Nichols and Zeckhauser, 1982; Besley and Coate, 1992; Baron and Myerson, 1982; Laffont and Tirole, 1993), and quantify the adoption per dollar that targeted transfers buy relative to an earlier increase in the regulated price.

The remainder of the paper proceeds as follows. Section II describes the institutional setting and Section III the data. Sections IV.A–IV.D document the descriptive evidence and develop the two-stage model of entry under a delayed rate. Section V estimates patient demand for hospitals and Section VI reports the estimates. Section VII presents the adoption-timing model, Section VIII the planner's problem and the policy comparison, and Section IX concludes.

II. Institutional setting

Aortic stenosis is a narrowing of the aortic valve that, once severe and symptomatic, is fatal within a few years unless the valve is replaced. For decades the only treatment was surgical aortic valve replacement (SAVR), an open-heart operation under cardiopulmonary bypass, which many elderly or frail patients cannot withstand. Transcatheter aortic valve replacement (TAVR) delivers a replacement valve through a catheter, usually via the femoral artery, without opening the chest. Randomized trials established its benefit in the order of surgical risk: TAVR reduced mortality relative to medical management for patients who could not undergo surgery (Leon et al., 2010), matched surgery for high-risk patients (Smith et al., 2011), and later matched it for intermediate-risk patients (Reardon et al., 2017). It also shortens recovery: in our data the mean length of stay is about five days for TAVR against ten for SAVR (Table 2). The FDA approved the first device in late 2011 and Medicare covered the procedure from fiscal 2012; fiscal 2011 activity is confined to clinical trial sites.

TAVR costs a hospital more than SAVR, largely because of the valve, and Medicare's payment did not reflect this for three years. Hospitals could code and bill the procedure from the moment Medicare covered it, because TAVR was assigned to the existing payment groups for cardiac valve procedures, MS-DRGs 216–221. The payment attached to those groups reflected the surgical procedures already in them and not the cost of a transcatheter

valve, so hospitals were reimbursed at approximately the surgical rate until fiscal 2015, when Medicare created separate payment groups for endovascular valve replacement, MS-DRGs 266 and 267, with a higher payment (Centers for Medicare & Medicaid Services, 2014; Dranove et al., 2022). Fiscal years 2012 through 2014 are therefore three years in which TAVR was billable and paid at about the surgical rate although it cost more, so a TAVR case earned a lower margin than a SAVR case; physicians, who billed under a dedicated CPT code, were paid for the procedure (Reddy et al., 2023; McCarthy et al., 2017). That the constraint was price and not coding matters for what follows: hospitals that did not adopt in this window were able to bill for the procedure and chose not to perform it, so the window reveals which hospitals were willing to adopt before per-case margins cleared their fixed cost. Figure 1 summarizes the timing.

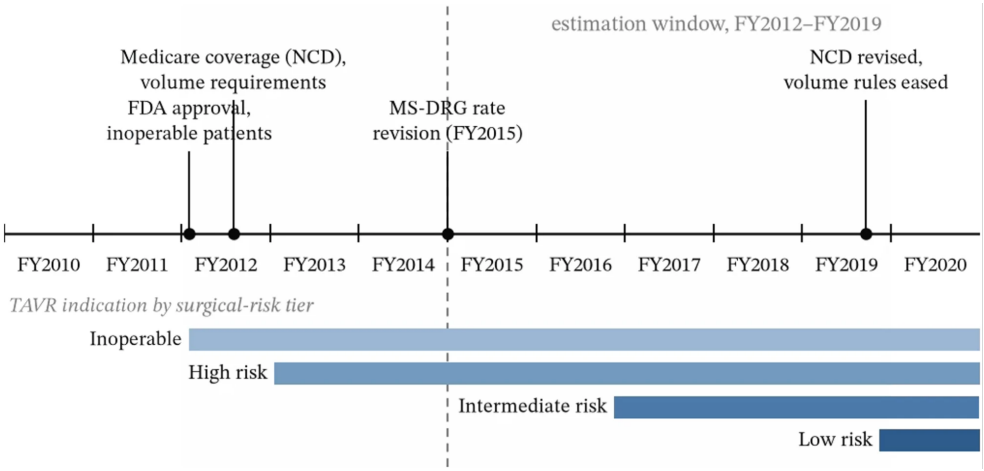


Figure 1: Coverage, coding and payment for TAVR under Medicare, fiscal years 2011–2019. The procedure was billable from the coverage decision under the payment groups for cardiac valve procedures; separate groups for endovascular valve replacement (MS-DRG 266/267) with a higher payment took effect in fiscal 2015.

Who could receive TAVR was set by Medicare coverage, which follows the FDA indication and expanded by surgical-risk tier: inoperable or extreme-risk patients from fiscal 2012 (Edwards SAPIEN), high-risk patients from fiscal 2013, intermediate-risk patients from the end of fiscal 2016 (SAPIEN 3) and in fiscal 2017 (Evolut), and low-risk patients from the end of fiscal 2019. Risk is measured by the Society of Thoracic Surgeons (STS) predicted operative mortality. We assign each fiscal year the threshold in force for most of it, 8 percent through fiscal 2016 and 4 percent from fiscal 2017; fiscal 2016 and 2019 end with an expansion, so in their last months the assigned threshold overstates the risk required for coverage. Coverage is the binding margin for Medicare patients: outside the indication

Medicare does not pay, and the procedure is unaffordable out of pocket. Two consequences follow. The pool of patients who could receive TAVR at a given hospital changed over the sample, but on dates common to all hospitals, so the change is not confounded with any hospital's adoption date. And at every date a large share of AVR patients had close to no chance of receiving TAVR; Section V uses this to separate patients' preference for the procedure from their preference for the hospital that offers it.

Finally, Medicare's national coverage determination conditions payment on program requirements. Until 2019 a hospital opening a program needed at least 50 aortic valve replacements in the prior year, including 10 high-risk patients, at least 1,000 catheterizations including 400 PCIs, and two surgeons with cardiac privileges, and an operating program had to sustain 20 AVRs a year; the 2019 revision lowered the requirements for new programs. The start-up requirement can bind for small hospitals, and Section VIII treats the volume requirement as a constraint on the set of hospitals a payer can reach.

III. Data

The paper needs four things from the data: every hospital's adoption date and volume, to study entry; every patient's diagnosis and choice of hospital, including the choice not to be treated, to estimate demand; hospital costs and Medicare payments, to value adoption in dollars; and the characteristics of hospitals and patients that enter the three models. Most of these come from 100 percent Medicare fee-for-service claims for fiscal years 2010–2020: the Inpatient file, which records for every stay the beneficiary's residential ZIP code, demographics, diagnoses and procedures, length of stay, total charges, DRG and Medicare payment, and a hospital identifier; the Outpatient file, which we use only to identify diagnosed patients who are not treated; and the Master Beneficiary Summary File. All aortic valve replacements for Medicare beneficiaries in this period occur in inpatient settings, and TAVR is performed almost entirely on patients aged 65 and older (Mack et al., 2013; Carroll et al., 2020), so the claims cover nearly every procedure outside Medicare Advantage. We supplement them with hospital characteristics from the American Hospital Association Annual Survey and the CMS Provider of Services file (beds, teaching status, ownership, system membership, location), with hospital-year cost-to-charge ratios from the Hospital Cost Reports, and with driving times from ZIP-code centroids to hospitals, computed for every patient–hospital pair used in the demand model.

An AVR episode is an inpatient stay with a TAVR or SAVR procedure code.[2] A hospital's adoption year is the first fiscal year in which it performs at least one TAVR procedure, and TAVR and SAVR volumes are the number of procedures with each code in a hospital-year, including patients who die in hospital. Because the procedure codes distinguish the transcatheter from the surgical approach throughout the period, adoption and volume are measured consistently on both sides of the fiscal 2015 payment revision. The sample of hospitals is short-stay general acute-care hospitals that perform AVR, after excluding those with sporadic or incomplete cost-reporting records: 946 hospitals, which we use throughout the descriptive evidence. The adoption-timing model of Section VII uses 945 of them: 97 pivotal-trial sites that performed TAVR before fiscal 2012, and 848 hospitals that choose when to enter, of which 521 adopt by 2019. The demand model is estimated on 933 hospitals.

Table 1 describes the hospitals. They are large, a mean of 434 beds; most are teaching (70 percent) and nonprofit (70 percent); and the typical hospital performed about 38 SAVR cases a year before TAVR arrived. Markets are not thin: the average hospital shares its referral region with five other sample hospitals and has eight others within an hour's drive, so for most hospitals adoption takes place in the presence of rivals that patients can reach. Section IV reports these characteristics by adoption cohort.

TABLE 1: Hospital characteristics

	All hospitals
Hospitals (N)	946
Beds	434
Teaching (share)	0.70
Government-owned (share)	0.10
Nonprofit (share)	0.70
For-profit (share)	0.16
Ownership missing (share)	0.05
SAVR cases per year, FY2010–11	38
Sample hospitals in the HRR	6
Sample hospitals within 60 minutes	8

Notes: Means unless noted. Pre-period volumes are annual averages over FY2010–2011. Hospitals within 60 minutes counts other sample hospitals reachable within 60 minutes of driving from the hospital's ZIP.

[2]TAVR: ICD-9-CM 35.05 and 35.06; ICD-10-PCS 02RF38H, 02RF38Z, 02RF3KH, 02RF3KZ, 02RF37H, 02RF37Z, 02RF3JH and 02RF3JZ. SAVR: ICD-9-CM 35.21 and 35.22; ICD-10-PCS 02RF07Z, 02RF08Z, 02RF0JZ and 02RF0KZ.

Patients enter the analysis in two ways. For the descriptive evidence the unit is the AVR episode. For the demand model, and through it for the supply model, the unit is the patient at risk, because a diagnosed patient can choose not to be treated, and the outside option is what lets the model count patients that adoption brings into treatment rather than only patients it moves between hospitals. The population at risk in a fiscal year is therefore every fee-for-service beneficiary aged 65 or older with a diagnosis of aortic valve disease that year, defined as at least one inpatient claim or at least two outpatient claims carrying the diagnosis,[3] who has not previously had a valve replaced, together with every beneficiary treated that year. Over fiscal years 2012 to 2019 this population contains 2.71 million patient-years, of which 14 percent receive a valve replacement in the year.

To place each patient against the eligibility thresholds we recover the STS predicted operative mortality as it is computed in clinical practice. The claims do not carry the score, but they carry its inputs: predicted mortality is read from the STS risk grid, a table of predicted operative mortality by age, sex, dialysis, hypertension and heart failure with $1{,}600$ cells ranging from 0.3 to 30 percent, and each patient is matched to a cell on age, sex and the three comorbidity flags, read from the diagnosis codes on the beneficiary's inpatient and outpatient claims in the same fiscal year. The same measure is used wherever risk enters, in the demand model of Section V and in the risk buckets of Section VII. For each patient we also compute driving time from the residential ZIP-code centroid to each hospital in the consideration set; travel time, not administrative market boundaries, defines choice sets, since about a third of AVR patients are treated outside their home hospital referral region.

Table 2 describes the treated patients by procedure. TAVR patients are older (81.5 against 73.7 years), more often women, and sicker, with a mean STS risk score of 3.2 percent against 1.7 for SAVR, as the risk-tiered indication implies; they stay half as long (4.9 against 10.0 days) and die in hospital half as often. Despite the difference in case mix, the DRG weight and the Medicare payment are almost identical across the two procedures, about $50,000 per case, which is the payment delay described in Section II. Travel is long for both groups: the mean trip to the treating hospital is about 96 minutes, two thirds of TAVR patients and three fifths of SAVR patients bypass their nearest hospital, and more than a quarter of each group travels over an hour. Patients already travel for aortic valve replacement; the question for the demand model is how much of that travel adoption redirects.

[3] ICD-9 424.1 and ICD-10 I35, which cover all nonrheumatic aortic valve disease, not only severe symptomatic stenosis. The outside option absorbs the breadth of the pool in levels, but the pool includes patients with no indication for replacement, who would not respond to access, so the model may overstate how many untreated patients a new program brings into treatment.

TABLE 2: Patient characteristics by procedure

	TAVR	SAVR
Episodes (N)	182,327	238,387
Age	81.5	73.7
Female (share)	0.46	0.37
STS predicted mortality (percent)	3.2	1.7
DRG weight	6.90	6.79
Length of stay (days)	4.9	10.0
In-hospital mortality (share)	0.02	0.04
Travel time to treating hospital, mean (minutes)	95.7	96.4
Bypassed the nearest hospital (share)	0.65	0.57
Travelled more than 60 minutes (share)	0.29	0.27
Total charges ($)	217,741	232,081
Medicare payment ($)	50,187	49,271

Notes: Medicare fee-for-service AVR episodes, FY2012–2019. Means unless noted. STS predicted mortality from the age–sex–comorbidity grid. Travel time from the patient's residential ZIP to the treating hospital.

The claims report total charges and the Medicare payment for each stay but not the hospital's cost. We approximate cost following Reddy et al. (2025), applying hospital-year cost-to-charge ratios from the cost reports; the device is the bulk of the variable cost of a transcatheter valve.[4] Per-case reimbursement is the payment amount on each claim. Figure 2 plots the distribution of charges and payments by procedure and year. Charges per TAVR case are higher than per SAVR case in every pre-revision year, a median of $178 thousand against $155 thousand in 2012, and estimated costs are higher until 2016 (Appendix Figure 10), yet Medicare paid the two procedures within a few thousand dollars of each other: the median TAVR payment is $50 thousand in 2012 and $45 thousand in 2014, against $44 and $42 thousand for SAVR. The 2015 revision raises the median TAVR payment to $52 thousand, a gap of about $10 thousand over SAVR; the gap then narrows as TAVR moves to lower-risk patients coded without major complications, and by 2019 the medians coincide at about $44 thousand, while SAVR charges rise over the decade and overtake TAVR charges from 2016. In the pre-2015 years a TAVR case was therefore paid about what a SAVR case was paid while costing more, so its margin was below the surgical margin. The revision widened the median payment gap from about $3 thousand in 2014

[4]Cost-to-charge ratios are hospital-wide. They apply the hospital's average markup to the device, so where devices are marked up less than other inputs they understate the TAVR–SAVR cost difference; and they measure fully allocated average cost, which can exceed the cost of an additional admission.

to about $10 thousand in 2015, and the gap is largest in the years right after it. These are medians across episodes; Section VII builds the hospital-, year- and risk-specific margins the adoption model uses.

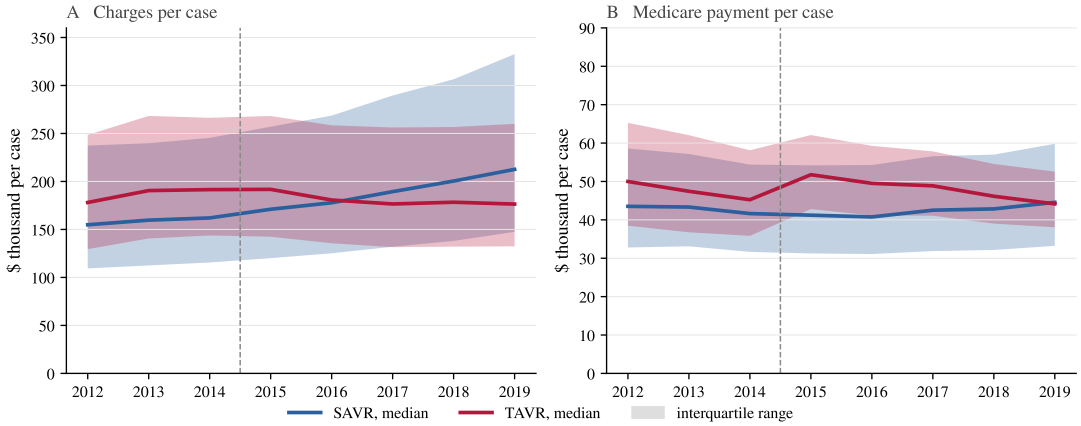


Figure 2: Charges and Medicare payment per case for TAVR and SAVR by fiscal year: median and interquartile range across episodes. The dashed line marks the fiscal 2015 payment revision. Source: Medicare fee-for-service AVR episodes, FY2012–2019.

IV. Descriptive Evidence

The claims record every hospital's adoption date and every patient's choice of hospital. This section uses both, and finds the following. More than two fifths of the hospitals that chose to adopt entered while TAVR was paid the SAVR rate, and entry did not jump when the rate rose. Hospitals respond to the rate: TAVR volume at existing adopters rises by about 8 percent at the revision. Learning exists but shortens stays rather than lowering the device cost, and hospitals with no local rival enter no earlier than others. Steering does appear: SAVR patients bypass their nearest hospital for one that offers TAVR, and a hospital's adoption hazard falls by 2.3 percent for every ten percentage points of rival adoption in its region. A two-stage model of entry under a delayed rate organizes these four tests. The reduced forms establish the sign of each mechanism; the sections that follow estimate its size.

IV.A. Two adoption puzzles

Of the 946 hospitals in our sample, roughly 65% adopted TAVR by the end of 2019; the remainder had not adopted by the end of the sample window. The fiscal-year distribution of

entrants (Figure 3) does not follow a smooth S-curve.[5] Two features make it puzzling.

Puzzle 1: a pre-reimbursement wave. The first wave of broad-based adoption comes in fiscal 2012, the first year the procedure could be billed outside clinical trial sites: excluding the 97 pivotal-trial sites, which ran programs before coverage, 117 of the 523 hospitals that chose to adopt enter that year, 22 percent, and 223, 43 percent, enter within the three years before the rate revision, when a TAVR case earned less than a SAVR case.[6] A model in which hospitals adopt when the procedure pays does not predict a discrete mass of entry at a relative loss.

Puzzle 2: a diffused post-reimbursement response. The 2015 revision raised the TAVR payment by about $10 thousand per case relative to SAVR (Section III) and raised the TAVR margin above the SAVR margin. Entry did not jump again; it rose gradually over the following five years. A model in which the revision releases a queue of hospitals waiting for the rate predicts a discrete jump in 2015, which the data do not show. By the end of the sample roughly 35% of hospitals had still not adopted. For some of them the setup cost exceeds the return even at the revised rate, which is why Section VIII considers transfers as well as the rate.

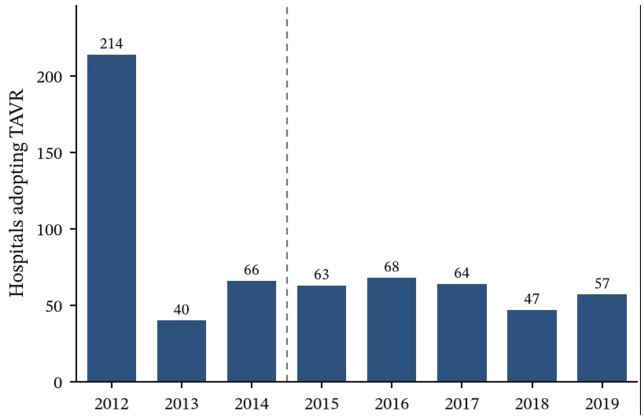


Figure 3: Distribution of hospitals by fiscal year of TAVR adoption.

IV.B. A two-stage model of entry under a delayed rate

The two puzzles are about timing: why hospitals enter before the rate pays, and why the rate, once it pays, draws in so few. To organize the evidence, we write the simplest model

[5]We use fiscal rather than calendar years because Medicare reimbursement claims are submitted on a fiscal-year basis.

[6]The cost of the valve kept early TAVR margins low (Reddy et al., 2023).

in which a delayed rate matters. Stage 1 is the window before the rate revision, $t < t_R$, in which TAVR is reimbursed at the surgical rate; stage 2 begins at t_R and continues indefinitely at the revised rate, with stage-2 quantities in present value at t_R and $\beta \in (0,1)$ discounting stage 2 to stage 1. Hospitals know t_R and the revised rate.[7]

Hospital i in market m pays a fixed cost F_i to adopt. In stage s it performs $v_{i,s}$ TAVR cases on its own AVR patients, each earning the TAVR margin over SAVR, μ_s, with $\mu_1 < 0 < \mu_2$: before the revision every TAVR case earns less than the surgical alternative, after it more. Two further primitives can raise the value of adopting. Learning: a hospital that operates in stage 1 enters stage 2 with lower per-case cost, raising its stage-2 margin by $\eta_i \geq 0$. Steering: under regulated rates the only margin on which an adopter competes is attracting patients. Market m contains a pool S_m of AVR patients who will travel to a hospital that offers TAVR, whether or not they receive it,[8] and an adopter captures a share $\chi_i \in [0,1]$ of the pool available when it enters. Captured patients stay, so a stage-2 entrant draws on the residual pool $Q_m = S_m\big(1 - \sum_{j \in E_m} \chi_j\big) \leq S_m$, where E_m is the set of stage-1 adopters in m, and each captured patient earns the SAVR margin $\mu_s^S > 0$, since most receive the incumbent procedure. The incremental profit from adoption therefore has a steering component and a procedure-margin component,

$$\Delta\pi_{i,1} = \chi_i S_m\, \mu_1^S + v_{i,1}\, \mu_1, \qquad \Delta\pi_{i,2}^E = \chi_i S_m\, \mu_2^S + v_{i,2}\, (\mu_2 + \eta_i), \qquad \Delta\pi_{i,2}^L = \chi_i Q_m\, \mu_2^S + v_{i,2}\, \mu_2,$$

where E and L index early (stage-1) and late (stage-2) adoption. The three options are worth

$$V_i^E = -F_i + \Delta\pi_{i,1} + \beta\, \Delta\pi_{i,2}^E, \qquad V_i^L = \beta\big(-F_i + \Delta\pi_{i,2}^L\big), \qquad V_i^N = 0.$$

Hospital i enters early rather than late when $V_i^E \geq V_i^L$,

$$\underbrace{\chi_i S_m\, \mu_1^S + \beta\, \chi_i\, (S_m - Q_m)\, \mu_2^S}_{\text{steering rent}} + \underbrace{\beta\, v_{i,2}\, \eta_i}_{\text{learning gain}} \;\geq\; \underbrace{-\, v_{i,1}\, \mu_1}_{\text{stage-1 loss}} + \underbrace{(1-\beta)\, F_i}_{\text{cost of paying early}}, \tag{1}$$

and enters late rather than never when $V_i^L \geq 0$,

$$\chi_i Q_m\, \mu_2^S + v_{i,2}\, \mu_2 \;\geq\; F_i. \tag{2}$$

[7] This is the perfect-foresight assumption maintained in the adoption-timing model of Section VII, which states it formally.

[8] Patients may prefer the minimally invasive procedure, or may read adoption as a signal of hospital quality. Section V estimates this preference.

Both conditions are cutoffs in F_i. Write R_i for the left side of (1), the rent to early entry, $\ell_i = -v_{i,1}\mu_1$ for the stage-1 loss, and L_i for the left side of (2), the stage-2 return of a late entrant. Hospitals with F_i below $F_i^E = (R_i - \ell_i)/(1-\beta)$ enter in stage 1, those between F_i^E and $F_i^L = L_i$ enter at t_R, and the rest never enter; if F_i^E exceeds L_i, no hospital enters at t_R and the early entrants are those with $V_i^E \geq 0$ (Appendix B, Lemma 1). Fixed-cost heterogeneity sorts hospitals into cohorts, and the rents in (1) and the residual pool in (2) shift the cutoffs.

The model says what early entry needs. With neither rent ($\eta_i = 0$, $\chi_i = 0$), a hospital enters before the revision only if $F_i \leq -\ell_i/(1-\beta) < 0$, that is, only if returns outside the Medicare margin exceed its setup cost. Within the Medicare margin, Puzzle 1 therefore requires learning or steering; returns outside that margin are a third source, which the adoption-timing model of Section VII absorbs in the net cost K. The two rents differ in where they come from. The learning gain $\beta v_{i,2}\eta_i$ does not depend on rivals and grows with a hospital's volume, as does the stage-1 loss $v_{i,1}|\mu_1|$; because the per-case loss before the revision is small (Section III), the learning gain is worth more to a hospital with no local competitor, which keeps all of its local volume. The steering rent requires patients who sort toward adopters and is larger the more of the pool a first mover can lock in before rivals enter. Puzzle 2 is the same primitive seen from the other side: every stage-1 capture lowers Q_m in (2), so where early entrants took much of the pool, fewer late hospitals clear their cutoff even though $\mu_2 > 0$. Steering rents for early entrants and residual-demand exhaustion for late entrants are the same mechanism.[9] Figure 4 draws the three options against F_i. The value of entering early falls faster in F_i than the value of waiting, because an early entrant pays the fixed cost before discounting, so the two cross at F^E; the value of waiting crosses zero at F^L. The steering rent raises the early line and moves F^E to the right, which is Puzzle 1; rivals' capture lowers the late line and moves F^L to the left, which is Puzzle 2. The learning gain enters V_i^E in the same way as the steering rent, so the figure does not separate the two; the market-structure test below does.

This gives four hypotheses, and each maps to an observable sign:

(H1) *Hospitals respond to margins.* The puzzles are puzzles only if the payment enters hospitals' decisions. If it did not, volume at existing programs would not respond to

[9]Heterogeneous beliefs about the revised rate add little. A hospital that is optimistic about μ_2 but faces neither rent has nothing to set against the stage-1 loss, so optimism alone cannot generate early entry; belief dispersion spreads entry dates among hospitals with positive rents, which in the reduced forms is indistinguishable from dispersion in F_i. The model is also silent on the pace of post-t_R entry; that is what the adoption-timing model of Section VII produces from the estimated distribution of F_i.

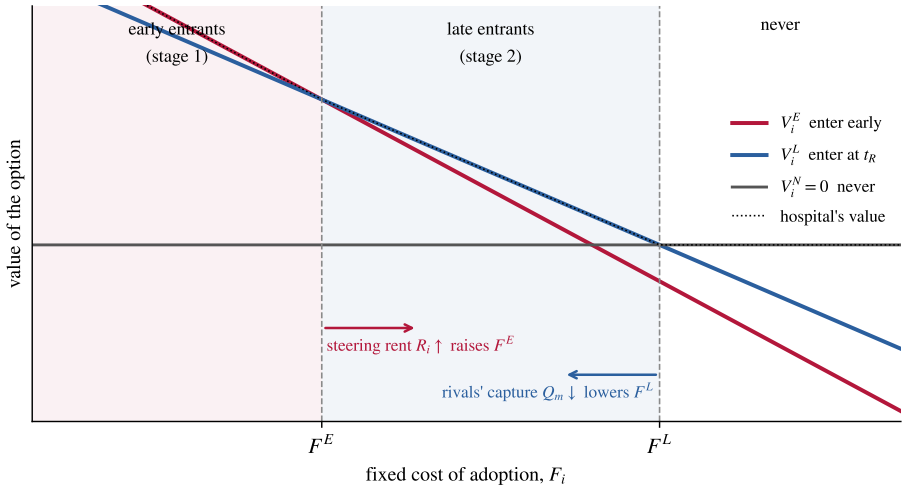


Figure 4: Entry cutoffs in the two-stage model. The lines are the values of entering early, V_i^E, entering at the revision, V_i^L, and never entering, against the fixed cost F_i; the dotted line is the hospital's value, the upper envelope. Hospitals below F^E enter in stage 1, those between F^E and F^L at t_R, and the rest never. The steering rent R_i shifts V_i^E up and F^E right; capture by stage-1 rivals lowers Q_m, shifts V_i^L down and F^L left. Illustrative parameters.

the rate revision.

(H2) *Learning.* Hospital experience reduces per-patient resource use, the necessary condition for $\eta_i > 0$; and if learning carries early entry, hospitals with no local rival enter earlier than others, since they keep all of the volume over which experience pays.

(H3) *Steering, rent side.* Patients sort toward adopters, including patients who receive SAVR.

(H4) *Steering, exhaustion side.* Post-t_R adoption hazards decrease in cumulative within-market rival entry.

H1–H3 concern early entry and are tested in Section IV.C; H4 concerns late entry and is tested in Section IV.D.

IV.C. Puzzle 1: margins, learning or steering?

We start with the first hypothesis. The test is whether hospitals that already offer TAVR change their volume when the rate rises, on hospitals that adopted in 2012, for which the 2015 cutoff isolates the reimbursement change: formal Medicare eligibility did not expand

materially until fiscal 2017 (Section II). We estimate

$$\log(\text{TAVR volume})_{it} = \alpha_i + \beta\,(\text{year}_t - 2012) + \tau\,\mathbf{1}_{t\geq 2015} + \delta\,\mathbf{1}_{t\geq 2015}(\text{year}_t - 2012) + \varepsilon_{it}, \tag{3}$$

over fiscal years 2013–2018,[10] and test $\Delta_{2015} \equiv \tau + 3\delta = 0$. From Table 3, $\widehat{\Delta}_{2015} = 0.471 + 3(-0.131) \approx 0.078$, an increase of roughly 8% in TAVR volume at the cutoff.[11] Hospitals already offering TAVR respond to the rate, so the payment enters their decisions. Two further tests point the same way and are reported in Appendix A: early adopters did not concentrate TAVR on higher-benefit patients after 2015, and adopters shifted their within-hospital AVR mix toward TAVR once its relative margin improved. These tests do not exclude other motives, since a hospital that valued patient benefit could also expand when a higher payment relaxes its budget; we maintain profit maximization in the models that follow.

	$\log$(TAVR volume)
Fiscal year $-$ 2012	0.289^{***}
	(0.049)
Post-2015	0.471^{***}
	(0.105)
Fiscal year $-$ 2012 $\times$ Post-2015	-0.131^{**}
	(0.051)
Observations	994
Hospital fixed effects	Yes (190)
R^2	0.805

Notes: Hospitals that adopted TAVR in fiscal 2012, fiscal years 2013–2018. Standard errors in parentheses, clustered at the hospital level. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

TABLE 3: TAVR volume at 2012 adopters around the 2015 rate revision

Learning is the first candidate. We estimate the within-hospital learning slope on length of stay and log hospital charges from

$$Y_{ihts} = \beta_0 + \beta_1 \cdot \text{Years since adoption}_{ht} + \mathbf{C}'_{iht}\boldsymbol{\theta} + \delta_h + \lambda_s + \varepsilon_{ihts}, \tag{4}$$

[10] We exclude 2012 because adoption may occur mid-year, yielding partial-year volumes.

[11] New valve generations, SAPIEN 3 and Evolut R, were approved in June 2015, the last quarter of fiscal 2015, so they can account for at most a small part of the fiscal-2015 jump.

where Y_{ihts} is length of stay in days or log total hospital charges for patient i in hospital h, year t, and severity category s. Hospital fixed effects δ_h absorb time-invariant hospital heterogeneity and DRG-based severity fixed effects λ_s absorb clinical case mix. The patient-level vector $\mathbf{C}_{iht}$ collects the Charlson Comorbidity Index and indicators for hypertension, congestive heart failure, chronic pulmonary disease, diabetes, renal disease, and prior cardiac surgery; this absorbs the secular shift toward lower-risk TAVR patients as eligibility expanded, the leading confound for $\hat{\beta}_1$. Standard errors are clustered at the hospital level.

TABLE 4: Effect of TAVR Adoption on Utilization and Charges

	Dependent variable:	
	Length of stay (days)	Log(Total charges)
	(1)	(2)
Years since adoption	-0.996^{***}	-0.030^{***}
	(0.074)	(0.005)
Constant	-2.202^{*}	11.697^{***}
	(1.180)	(0.073)
Observations	40,025	40,025
R^2	0.346	0.683
Adjusted R^2	0.333	0.676
Residual Std. Error (df = 39,207)	4.675	0.288
F Statistic (df = 817; 39,207)	25.438^{***}	103.396^{***}

Note: Includes hospital and severity-category fixed effects and patient-level Charlson Comorbidity Index, hypertension, CHF, COPD, diabetes, renal disease, and prior cardiac surgery indicators.

$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

Table 4 reports the estimates. Each additional year since adoption reduces length of stay by close to one day (Column 1), about 12% of the baseline mean of 8 days, and reduces total hospital charges per patient by about 3% (Column 2; $100 \cdot (e^{-0.03} - 1) \approx -2.95\%$). Learning exists, but it operates on the intensive margin: the prosthetic valve is a device cost that experience cannot erode, which caps η_i.

A second check uses market structure. If learning carried early entry, a hospital that is the only one in its HRR, which keeps all of its local volume, would enter earlier than others; steering predicts the opposite, since such a hospital has no local rivals to draw patients from.

Using hospitals observed in fiscal 2012, we estimate the linear probability model

$$\mathbf{1}\{\tau_h < 2015\} = \alpha + \beta\,\text{Monopoly}_h + \gamma\,\text{Beds}_h + \delta_{\text{type}(h)} + \varepsilon_h,$$

where Monopoly_h equals one if h is the only hospital in its HRR, restricted to HRRs with fewer than ten hospitals. Conditional on size and type, monopoly status is not associated with earlier adoption (Table 5): the estimate is -3.4 percentage points (s.e. 4.5), which excludes a monopoly advantage larger than about 5 points. HRR boundaries are crossed by patients (Section III), so the measure is coarse. The early-entry advantage that learning predicts does not appear.

TABLE 5: Monopoly status and pre-t_R adoption

	(1) Pre-t_R adoption
Monopoly	-0.034
	(0.045)
Beds	0.001^{***}
	(0.000)
Hospital type FE	Yes
Observations	857
R^2	0.21

Sample: hospitals observed in fiscal year 2012 in HRRs with $n < 10$ hospitals. Linear probability model; standard errors in parentheses. *** $p < 0.01$.

Steering requires $S_m > 0$: patients who travel to a hospital because it offers TAVR. SAVR patients provide the test. SAVR is available at every hospital, so a SAVR patient who bypasses her nearest hospital for one that offers TAVR is not traveling for the procedure she receives. She may be traveling for the hospital, read through its program, or for the chance of TAVR before the heart team assigned her procedure; Section V separates the two by risk. Figure 5 shows that in 2012–2014, 41 to 46 percent of SAVR patients who live within an hour of a program bypass their nearest hospital for a TAVR hospital, and in 2015–2019, 47 to 57 percent at every distance; where the nearest program is 60 to 120 minutes away, the median SAVR patient travels 13 to 16 minutes beyond her nearest hospital after 2015, against about 8 before. Under regulated prices this is steering: adoption is a non-price attribute on which hospitals compete, and the demand model in Section V measures how much patients will travel for it.

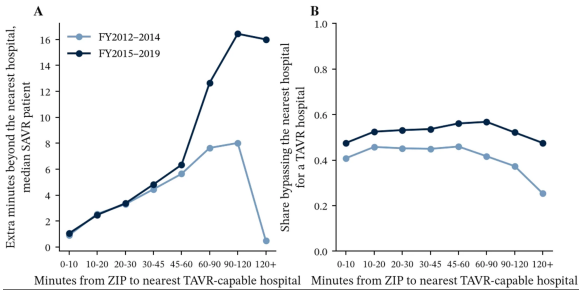


Figure 5: SAVR patients' travel toward TAVR hospitals, by driving time from the patient's ZIP code to the nearest TAVR-capable hospital, fiscal 2012–2014 and 2015–2019. Panel A: median extra minutes beyond the nearest hospital. Panel B: share bypassing the nearest hospital for a TAVR hospital.

IV.D. Puzzle 2: residual demand for late entrants

Condition (2) says that a hospital's value of entering at or after t_R falls with the share of the contestable pool that rivals captured earlier (H4). Fixed-cost heterogeneity alone spreads post-t_R entry over time but predicts no relation between a market's early entry and its later hazards; exhaustion does. We test it with a Cox proportional-hazards model stratified by HRR. The framework is discrete-time, so a complementary log-log or discrete-time logit would be equally appropriate; we use Cox for the flexibility of the HRR-specific baseline hazard $\lambda_{0,m(h)}(t)$, which absorbs arbitrary HRR-level time profiles. A discrete-time logit yields identical signs and similar magnitudes. For each hospital h that has not yet adopted by year t,

$$\lambda_h(t) \;=\; \lambda_{0,m(h)}(t)\,\exp\Big(\beta\,\text{Competitor Rate}_{ht} \;+\; \mathbf{X}_{ht}\,\boldsymbol{\gamma}\Big), \tag{5}$$

where Competitor Rate$_{ht}$ is the fraction of other hospitals in h's HRR that have adopted by t and $\mathbf{X}_{ht}$ collects hospital size, ownership, teaching status, and age. Standard errors are clustered at the HRR level. β identifies exhaustion under the assumption that cumulative within-HRR adoption is uncorrelated with unobserved time-varying hospital shocks given $\mathbf{X}_{ht}$ and the HRR baseline; the conventional confound, that cumulative entry proxies a common positive HRR shock, would produce a positive β.

The Competitor Rate coefficient is -0.232 and significant (Table 6): conditional on the HRR baseline and hospital controls, a 10-percentage-point increase in within-HRR cumulative adoption lowers the adoption hazard by about 2.3%. Earlier entry leaves less residual demand and lowers the value of entering for those who remain.

The hospital controls in Table 6 are individually significant and consistent with het-

TABLE 6: Cox proportional hazard for first TAVR adoption (HRR-stratified)

	(1)
Competitor Rate	-0.232**
	(0.0999)
Number of beds	-0.000125***
	(0.0000393)
Teaching hospital	0.117***
	(0.0304)
Ownership type	-0.0584***
	(0.0218)
Year hospital is active	0.00243
	(0.0309)
Observations	5114

Standard errors clustered at the HRR level in parentheses.
Baseline hazard stratified by HRR.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

erogeneity in F_i. The teaching indicator enters positively, consistent with a lower cost of assembling a cross-specialty heart team at teaching hospitals, and beds enter negatively with a small per-bed magnitude. At this stage such sorting is indistinguishable from variation in beliefs, and the covariates should be read as an unconditional sorting fact rather than as estimates of cost differences. The adoption-timing model in Section VII recovers the adoption cost from adoption dates directly; its estimate is the fixed cost net of returns to a program that accrue outside the Medicare margin, the counterpart of F_i here.

The descriptive evidence establishes that hospitals respond to margins, that learning operates on length of stay rather than on the device cost, and that patient steering, identified from patient sorting and from the negative effect of cumulative rival entry on later hazards, is the mechanism that generates both puzzles. These reduced forms establish signs, not sizes. Section V estimates how much demand adoption shifts toward a hospital; Section VII converts that shift into dollars and recovers the adoption cost that rationalizes observed entry dates.

V. Patient demand for hospitals and technology

The descriptive evidence shows that adoption attracts patients. To measure how many, and which, we estimate a model of hospital choice on the population of Medicare beneficiaries diagnosed with aortic valve disease. The model finds that a TAVR program raises the value of a hospital to patients well below the risk at which they could receive the procedure, by an amount worth about 22 minutes of travel to the median patient in 2012 and 75 minutes by 2019; Section VI reports the estimates. Each patient chooses among nearby hospitals or no treatment, and a TAVR program raises the value of a hospital by an amount that depends on the patient's surgical risk. We estimate demand for the hospital, not for the procedure, because that is the choice the patient makes: the heart team assigns the modality after admission. The variation in eligibility by risk tier, on dates common to all hospitals, is what lets us separate a preference for the procedure from a preference for the hospital that offers it, since at every date a large share of patients could not receive TAVR anywhere. The estimates give the change in a hospital's admissions when it adopts, holding rivals' adoption fixed,

$$\Delta Y_{j,t}^{\text{steer}} \;:=\; Q_{j,t}\big(a_{j,t}=1,\, a_{-j,t}\big) \;-\; Q_{j,t}\big(a_{j,t}=0,\, a_{-j,t}\big), \qquad Q_{j,t}=\sum_{i\in\mathcal{I}_t} P_{ijt}, \quad (6)$$

where $\mathcal{I}_t$ is the market in year t, defined below, and P_{ijt} is the probability that patient i is treated at hospital j. The capture share χ_j of Section IV.B is $\Delta Y_{j,t}^{\text{steer}}/Q_{m(j),t}$, the gain in j's flow as a share of contestable demand. Three sets of parameters govern it: the travel-cost gradient, the value of a TAVR program at each level of risk, and the dispersion of patients' preference for treatment, which determines whether the gain comes from rival hospitals or from patients who would not otherwise have been treated. The counterfactuals report the two parts separately.

Adoption is not random; large, high-volume hospitals adopt first. Two features of the specification prevent selection into adoption from biasing the TAVR coefficients. Hospital fixed effects absorb permanent differences in quality, scale, and referral relationships, and year effects absorb aggregate changes in treatment rates. The TAVR coefficients are then identified from the change in a hospital's patient flow when its program opens and from the difference in the risk composition of patients at TAVR and non-TAVR hospitals within a year. We maintain that, conditional on the hospital and year effects, the timing of adoption is unrelated to hospital-specific demand shocks. Section V.B discusses this assumption and

the two features of the specification that limit the bias if it fails; an event study of hospital admissions around adoption would test it directly.

In the counterfactuals the estimated technology terms are held fixed. This is appropriate for policies that change when and where TAVR is available and what patients are paid to reach it, and not for policies that act on hospital quality. The modality an admitted patient receives is chosen by the hospital and is modeled in Section VII.

V.A. Environment and consideration sets

Throughout, i denotes a patient-year, j a hospital, t a fiscal year, and $j = 0$ the outside option of no aortic valve replacement in year t. Let $a_{jt} \in \{0, 1\}$ indicate whether hospital j performs TAVR in year t, and let r_i denote patient i's predicted operative mortality in percent, the risk measure defined in Section III.

The population at risk in year t is the one defined in Section III: fee-for-service beneficiaries aged 65 or older diagnosed with aortic valve disease that year who have not previously had a valve replaced, together with every beneficiary treated that year.[12] This is the population a hospital can draw on when it adopts.

The utility of patient i from hospital j in her consideration set $\mathcal{J}_{it}$, defined below, in year t is

$$U_{ijt} = \delta_j + \gamma_t + \beta_{1,i}\,\mathrm{tt}_{ij} + \beta_2\,\mathrm{tt}_{ij}^2 + \beta_c\,\mathbf{1}\{j = c(i)\} + \kappa_t\,r_i + a_{jt}\left(\alpha_{0t} + \alpha_{1t}\,r_i\right) + \sigma_0\,\zeta_i + \varepsilon_{ijt}, \qquad U_{i0t} = \varepsilon_{i0t}, \tag{7}$$

where tt_{ij} is the driving time in minutes from i's ZIP code to hospital j, $c(i)$ is the hospital nearest to i, and ε_{ijt} is i.i.d. type-I extreme value. The outside option is normalized to zero. The fixed effect δ_j absorbs hospital attributes common to all patients and years; the year effect γ_t shifts the value of treatment at any hospital relative to no treatment, with γ_{2012} normalized to zero. Two terms vary across patients. Travel disutility is $\beta_{1,i} = \beta_1 + \sigma\eta_i$, and the value of treatment at any hospital relative to none carries a patient-specific component $\sigma_0\zeta_i$, with η_i and ζ_i independent standard normal, so the model is a mixed logit. The second term is the random-coefficients counterpart of a nest that contains every hospital and excludes the outside option: patients differ in their propensity to be treated at all, so hospitals are closer substitutes for one another than for no treatment, and a hospital that adopts draws relatively more patients from its rivals and fewer from the untreated than a

[12] Treated patients are counted in the fiscal year of their first procedure.

plain logit would imply. Medicare regulates prices and standardizes cost sharing, and no price term enters.

Risk enters twice. The term $\kappa_t r_i$ shifts the value of being treated anywhere. It captures that sicker patients are, in the data, less likely to undergo a valve replacement at all. The TAVR term $\alpha_{0t} + \alpha_{1t} r_i$ is the additional value of a hospital that offers TAVR, allowed to vary with the patient's risk and with the year. The year dependence matters because the indication for TAVR expanded by risk tier over the sample (Section II), so the set of patients for whom a TAVR program means access to the procedure, rather than a signal about the hospital, changes over time. Both coefficients are common to all patients who consider the hospital, including those who go on to receive SAVR.

With the extreme-value error, the choice probability is

$$P_{ijt} \;=\; \int \frac{\exp\big(U_{ijt}(\eta,\zeta)\big)}{1+\sum_{k\in\mathcal{J}_{it}} \exp\big(U_{ikt}(\eta,\zeta)\big)}\, dF(\eta)\, dF(\zeta), \tag{8}$$

and predicted flows are $Q_{jt} = \sum_{i\in\mathcal{I}_t} P_{ijt}$ over the market.

Choice sets are based on travel time rather than on administrative market boundaries, since SAVR patients routinely bypass their nearest hospital for one that offers TAVR (Figure 5). Each patient's set consists of the twenty hospitals nearest the patient's residential ZIP code by driving time, the hospital actually chosen when it lies outside that set, and the outside option. The median treated patient travels about 30 minutes, far inside the reach of the twentieth-nearest hospital, so the sets are wide relative to behavior. The modality a treated patient receives is determined by the hospital and is modeled on the supply side, so consideration sets do not condition on clinical eligibility.

V.B. Identification

Hospital fixed effects are the standard device in this literature and we use them for the same reason: δ_j absorbs everything about a hospital that patients value in every year, so the TAVR term is identified from changes within a hospital and from differences in composition within a market. The TAVR term itself has two meanings a patient may attach to a TAVR hospital, the hospital, with the program read as a signal of its quality, and access to the procedure. The specification does not separate them, because modality is not a patient choice in the model, but the regulatory thresholds do: for a patient below the risk threshold in force that year a TAVR hospital could not offer the procedure, so any positive value at that risk level is

a preference for the hospital. Section VI uses this reading; the counterfactuals do not require the decomposition, since they use the estimated change in flows rather than its behavioral source.

The parameters to identify are the travel disutility and its dispersion (β_1, β_2, σ), the preference for the nearest hospital (β_c), the year effects, the risk gradient of treatment (κ_t), the dispersion of the preference for treatment (σ_0), and the technology terms (α_{0t}, α_{1t}). The travel parameters are identified by variation in patients' locations. Patients in different ZIP codes face different travel-time vectors to the same hospitals, and the rate at which choice probabilities fall with travel time traces out the gradient, its curvature, and, through the share of long trips, its dispersion. The nearest-hospital coefficient is identified by the excess probability mass on the closest hospital relative to what the smooth travel terms predict.

The year effects and the risk gradient are identified from the treatment margin. Within a year, the share of the market that is treated at all pins down γ_t given the hospital effects, and the way that share falls with the risk score pins down κ_t. These terms matter for what follows because they absorb two forces that would otherwise be attributed to the technology terms, the secular rise in treatment over the sample and the fact that high-risk patients are treated less often everywhere. The dispersion σ_0 is identified from how the treated share responds when a patient's choice set changes: when a hospital in the set adopts, a large σ_0 means that its gain comes mostly from the other hospitals in the set, a small one that it comes in larger part from patients who would not otherwise have been treated.

The technology terms are identified from two sources. The first is staggered adoption. When hospital j opens a program, its patient flow from every ZIP code in its market area changes, while δ_j is fixed; 405 hospitals adopt between fiscal 2013 and 2019 in the estimation sample, each supplying a before-and-after comparison within the hospital. The second is the cross-section of risk. Within a year, a TAVR hospital draws patients of different risk scores in different proportions than a non-adopter in the same market; α_{1t} is identified from that difference, net of κ_t. Because both sources are within-hospital or within-market, permanent quality differences between adopters and non-adopters do not contribute.

The threat to this design is that adoption may coincide with other changes that affect patient choice. A capacity expansion, a marketing campaign, or a strengthening of referral relationships timed to the launch of a TAVR program would raise the hospital's patient flow and be attributed to α_{0t}. The maintained assumption is that, conditional on the hospital and year effects, such changes are not systematically timed to adoption. Two features of the setting limit the scope for such bias. The year effects absorb market-wide trends, and the

risk slope α_{1t} is identified from the composition of a hospital's patients rather than from their number, so a shock that raises a hospital's flow uniformly across risk levels does not move it.

V.C. Estimation

The parameters are estimated by simulated maximum likelihood on the market defined above, over fiscal years 2012 to 2019, 2.71 million patient-years and 933 hospitals. Choice probabilities are simulated with two-dimensional scrambled Halton draws, 15 per patient during optimization and 100 at the converged estimates. For a candidate parameter vector, the hospital effects are concentrated out: δ solves the condition that predicted and observed patient shares coincide for every hospital, with the outside option held at zero, and we maximize the log likelihood over the remaining parameters by a quasi-Newton method. With random coefficients, share matching is not the likelihood's first-order condition in δ, so the estimator is not exact maximum likelihood. It is consistent, because at the true parameters expected and population shares coincide, and it is the standard approach when the number of fixed effects is large (Berry et al., 1995; Goolsbee and Petrin, 2004).

VI. Demand results

Three results come out of the estimates (Table 7). Travel costs are lower than in studies of emergency and maternity admissions, as expected for a planned procedure, and patients place a large premium on the nearest hospital. A TAVR program raises the value of a hospital to patients whether or not they are eligible for the procedure, from 2015 at every level of risk, and the value grows with risk. And the value a program adds rose over the sample as the indication widened.

Travel is costly and patients care about the closest hospital. Every additional ten minutes of driving lowers the odds of choosing a hospital by about 16 percent.[13] The quadratic term is positive, so the disutility of travel flattens at long distances: it offsets about 3 percent of the linear term at 100 minutes and 6 percent at 200. The option of the closest hospital is worth about 90 minutes of travel on its own, so proximity is a major determinant of hospital choice. The travel gradient is several times flatter than in hospital choice for childbirth, where a

[13]The travel coefficient is -0.0174 per minute, so ten minutes change the log-odds by -0.174, $e^{-0.174} \approx 0.84$. With price fixed by Medicare, travel time is the price patients pay, and all other hospital attributes are measured against it in minutes.

TABLE 7: Patient demand for hospitals: estimated parameters

A. Travel and choice-set parameters				
Travel time (minutes)				-0.0174
Travel time squared ($\times 10^6$)				5.73
Nearest hospital				1.574
Standard deviation of the travel coefficient, σ				0.0095
Standard deviation of the inside-option constant, σ_0				0.4420
B. Year-specific parameters				
Fiscal year	TAVR level, α_{0t}	TAVR $\times$ risk, α_{1t}	Risk, κ_t	Year, γ_t
2012	0.364	0.011	-0.309	0 (norm.)
2013	-0.188	0.294	-0.558	0.635
2014	-0.070	0.269	-0.520	0.517
2015	0.085	0.265	-0.475	0.336
2016	0.037	0.353	-0.541	0.468
2017	0.242	0.396	-0.574	0.247
2018	0.245	0.449	-0.597	0.170
2019	0.439	0.458	-0.609	-0.036
Log likelihood				-1,807,657.0
Patient-years				2,714,812
Hospitals				933

Notes: Mixed logit with hospital fixed effects, concentrated out, and an outside option normalized to zero. Risk is predicted operative mortality in percentage points. The year effect for 2012 is normalized to zero. Standard errors to be added.

standard logit gives -0.121 per minute (Raval and Rosenbaum, 2021) and the distance coefficient is -0.215 per mile (Ho and Pakes, 2014), which imply odds reductions of 50 to 80 percent per ten minutes. Aortic valve replacement is a planned, referral-driven procedure, and patients referred for tertiary care travel much farther than emergency admissions: among Medicare heart-attack patients, transfers travel a median of 34 miles against 5 for emergency admissions (Chandra et al., 2016). The nearest-hospital premium, 1.57, is larger than the closest-hospital effect of 1.07 for English cardiac bypass patients (Gaynor et al., 2012), consistent with the finding that a large share of the distance gradient in hospital choice reflects a preference for the home hospital rather than transport cost (Raval and Rosenbaum, 2021). Travel disutility also varies across patients: the standard deviation of the travel coefficient is 0.0095, about half the mean. Under the normal distribution about 3 percent of patients have a nonnegative coefficient. Patients with flat gradients are the ones a distant adopter can attract, and the voucher values subsidized minutes at the absolute value of each

patient's coefficient, so it never lowers utility. The preference for treatment varies across patients, with a standard deviation of 0.44. Because this term is common to every hospital in a patient's set, it governs how much of an adopter's gain comes from rival hospitals rather than from patients who would otherwise go untreated (Section V.B). The year effects are positive in most years relative to 2012, and the risk gradient κ_t is negative throughout, -0.31 per point of the risk score in 2012 and between -0.48 and -0.61 after: holding hospitals fixed, a patient one point riskier is less likely to decide to undergo treatment at all. We know of no hospital-choice model that reports this gradient, but the clinical evidence has the same sign: among patients with severe aortic stenosis, those who do not undergo valve replacement have higher surgical risk, a median STS score of 2.7 against 1.8 percent for those operated (Bach et al., 2009), and older patients and those with poor ventricular function are more often denied surgery (Iung et al., 2005). These coefficients absorb who gets treated; without them the TAVR coefficients would pick up the fact that TAVR hospitals treat sicker patients rather than the value patients place on a TAVR hospital.

The TAVR coefficients are evidence that both preferences exist, for the procedure and for the hospital that offers it, and Figure 6 is the way to read them. For each year it draws the utility of a TAVR hospital relative to a non-adopter, $\alpha_{0t} + \alpha_{1t} r$, against the STS risk score r, over the risk distribution of treated and untreated patients, with the eligibility threshold in force marked. The TAVR term combines a preference for the procedure and a preference for the hospital. The eligibility threshold separates them: a patient below the threshold could not receive TAVR at any hospital,[14] so where the utility of a TAVR hospital is positive below the threshold, choice is driven solely by a preference for the hospital. The range of risk between the point at which the line crosses zero and the threshold is therefore the hospital preference, which we read as a quality signal.

The first-order pattern is that in every year the utility of a TAVR hospital is positive over almost the whole range of risk below the eligibility threshold. In 2013 and 2014 the line is steep and crosses zero at a risk score of 0.6 and 0.3 percent, against a threshold of 8 percent; from 2015 it is positive at every risk level, close to zero only at the lowest scores until 2017. The median treated patient, at a risk score of 1.9 percent, is above the break-even point in every year, so patients with no access to the procedure were drawn to the hospitals that offered it throughout. The first year differs: in 2012 the line is flat and positive, $\alpha_{0,2012} = 0.36$ and $\alpha_{1,2012} = 0.01$, so a TAVR hospital is preferred at every risk

[14] Eligibility is not only a near-zero chance of receiving the procedure. Below the threshold Medicare does not cover TAVR, so the procedure is not available to a fee-for-service patient at any price she could pay.

level by 22 to 26 minutes of travel. In the year the procedure arrived, with little information about what it meant, patients treated adoption as a positive signal about the hospital at every level of risk rather than as access to a procedure for a defined group.

Beyond the break-even point the gain rises with risk, and the slope steepens over time. The risk gradient α_{1t} rises from 0.01 per point in 2012 to 0.29 in 2013 and 0.46 in 2019, and the level α_{0t} rises from -0.19 in 2013 to 0.44 in 2019, so each year more of the line lies in the positive domain and the line is steeper: patients learn more about what a program means for the hospital and value it more, and those at higher risk, who also gain access to the procedure, value it most. Converted into travel time, the gain at a risk score of 4 percent is about 23 minutes in 2012, and by 2019 it is about two hours, with about 78 minutes at 2 percent. The fall in the break-even risk is the demand-side counterpart of the diffusion documented in Section IV.A; late adopters entered a market in which the technology mattered to a larger share of patients than it had for the 2012 cohort. Two readings of the levels are therefore consistent with the estimates, and both matter for what follows. Below the threshold the program is a signal about the hospital, the steering mechanism of Section IV.B; above it, and increasingly after the indication expanded, the value includes access to the procedure itself. The counterfactuals use the full estimated pull, and the supply model's modality rule determines which of these patients receive TAVR once admitted. Because the market includes the untreated, about 340,000 patient-years in 2019 of which 57,000 are treated, the pull accrues on every patient in a program's market area in every year of its operation; part of the resulting gain in admissions comes from rival hospitals and part from patients who would otherwise have gone untreated; the counterfactual TAVR counts are net of patients a new program takes from existing programs.

VI.A. Use in counterfactuals

The demand system enters the counterfactual analysis (Section VIII) at two points. First, for each hospital and year it supplies the expected number of market patients the hospital treats with and without a TAVR program, holding rivals at their observed status, and the market size $|\mathcal{I}_t|$; these feed the per-period profit of adoption in the supply model. Second, under any candidate adoption path $\{T_j^*(\mathcal{P})\}$ we recompute $\{a_{jt}(\mathcal{P})\}$, the choice probabilities $P_{ijt}(\mathcal{P})$ from (8), and the flows $Q_{jt}(\mathcal{P})$. The steering primitive and the capture share are as in (6),

$$\chi_j(\mathcal{P}) \;=\; \frac{\Delta Y_{j,t}^{\text{steer}}(\mathcal{P})}{Q_{m(j),t}(\mathcal{P})},$$

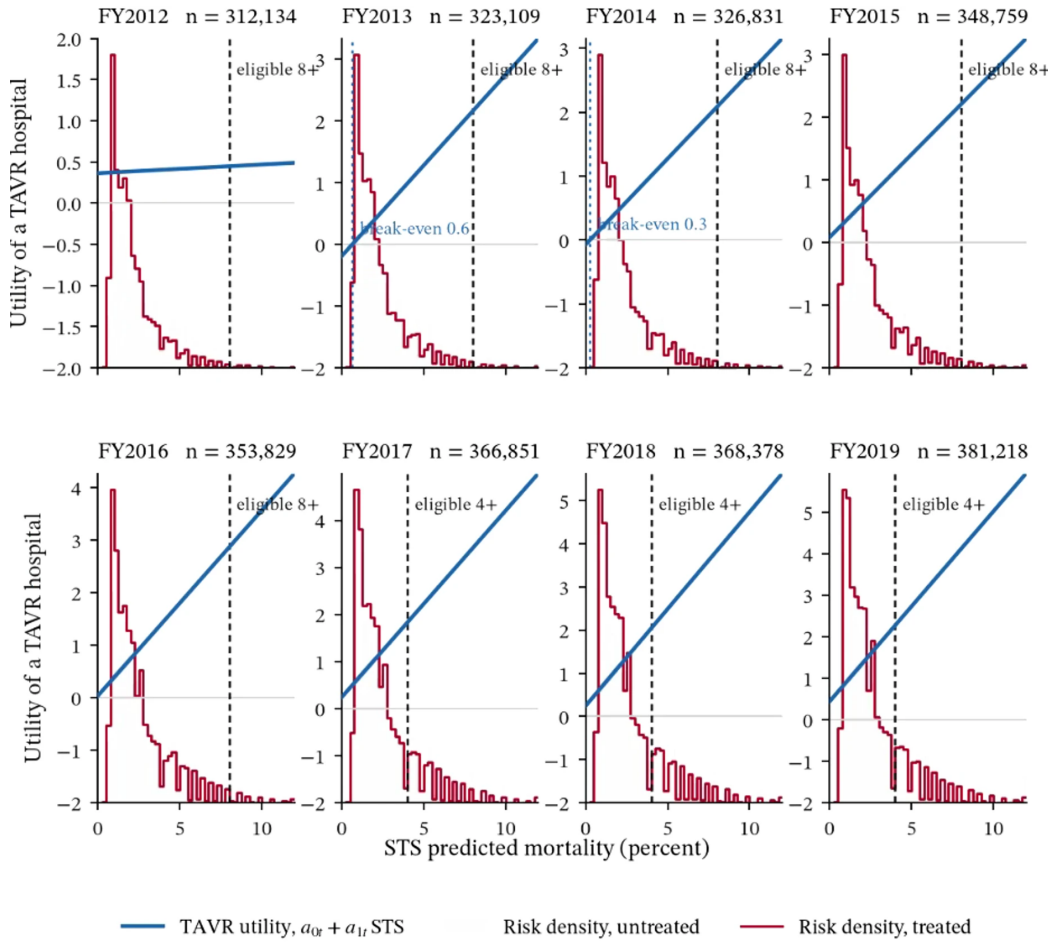


Figure 6: The utility of a TAVR hospital by patient risk and fiscal year. In each panel the line is $a_{0t} + a_{1t}\,\mathrm{STS}$ from Table 7, the utility a patient at the STS risk score on the horizontal axis attaches to a hospital offering TAVR relative to a non-adopter. The shaded area is the risk distribution of untreated patients in the market and the red step line that of treated patients. The dotted line marks the break-even risk at which the utility is zero, and the dashed line the eligibility threshold in force, 8 percent through fiscal 2016 and 4 percent from fiscal 2017. Patients between the two lines were drawn to TAVR hospitals although they could not receive TAVR there.

which feeds the residual-demand law and the adoption value in the counterfactual section. The travel-voucher policy enters through the same system. Paying a patient's extra travel to a TAVR hospital lowers tt_{ij} in (7) by the subsidized minutes, and the consumer-surplus change is computed from the same utilities in travel-time units.

Three restrictions apply. The hospital fixed effects are held fixed across counterfactuals, in line with the maintained assumption that the policies considered change the timing and location of adoption rather than hospital quality. The technology terms α_{0t}, α_{1t} are held at their estimated year-specific values, so a counterfactual that moves adoption earlier applies the pull estimated for that calendar year. And because the market includes the untreated, counterfactual policies change not only where patients are treated but how many are treated; the fiscal accounting in Section VIII counts the payments for newly treated patients separately from the payment increment on existing cases.

VII. Adoption timing

The demand estimates give each hospital's admissions with and without a program in every year, given rivals' adoption. Combined with reimbursement and costs from the claims, they give the net present value of adopting at each candidate date, and observed adoption dates then identify the remaining unknown, the net cost of adoption. The model recovers a distribution of net adoption costs in dollars that separates adopters from non-adopters in the direction it requires, a median of \$5.0M for adopters against \$9.4M for hospitals that never adopt, and reproduces the 2012 wave and the absence of a second one in 2015 (Section VII.C). Each hospital solves a perfect-foresight timing problem given the observed adoption paths of its rivals: it knows the reimbursement path and chooses an adoption year once, irreversibly. Time t is discrete. At t_0 the hospital draws a shock for each candidate year and chooses the year that maximizes the present value of adopting net of the shock. In each year after adoption, patients choose hospitals, the hospital assigns each admitted patient to TAVR or SAVR, and profits are realized. Because all shocks are known at t_0, waiting has no option value: a hospital delays only because a later date is worth more.

VII.A. Demand, modality and operating profit

Let θ denote a patient's predicted operative mortality. The market in year t, the diagnosed beneficiaries of Section V, has size M_t and risk distribution $F_t(\theta)$. The demand system

delivers hospital choice probabilities by risk,

$$s_{h,t}(\theta; a_h, A_{-h,t}) \equiv \Pr(\text{treated at } h \mid \theta, t; a_h, A_{-h,t}),$$

where $a_h \in \{0, 1\}$ indicates whether h has adopted and $A_{-h,t}$ stacks rivals' adoption states, held at their observed values. Expected admissions are

$$Y_{h,t}(a_h) = M_t \int s_{h,t}(\theta; a_h, A_{-h,t})\, dF_t(\theta).$$

For adopters, let $p_{h,t}(\theta) = \Pr(\text{TAVR} \mid h, \theta, t)$ be the probability that an admitted patient receives TAVR; for non-adopters $p_{h,t}(\theta) = 0$. We estimate p as a logit on the claims of adopting hospitals, with hospital and year effects, the patient's risk, risk interacted with year, and the per-case margin gap between the two procedures; patients above the 99th percentile of SAVR patients' risk in that year receive TAVR wherever a program exists; at a hospital without one they earn it no margin, which affects about 1 percent of SAVR patients. Per-case margins are

$$m^m_{h,t}(\theta) = \text{reimb}^m_{h,t}(\theta) - \text{cost}^m_{h,t}(\theta), \qquad \Delta_{h,t}(\theta) = m^T_{h,t}(\theta) - m^S_{h,t}(\theta), \qquad m \in \{T, S\},$$

and period-t operating profit is

$$\pi^{op}_{h,t} = M_t \int s_{h,t}(\theta; a_h, A_{-h,t}) \Big[m^S_{h,t}(\theta) + p_{h,t}(\theta)\, \Delta_{h,t}(\theta) \Big]\, dF_t(\theta).$$

Charges and reimbursements per case are predicted from regressions on the claims, by modality, with year effects, risk-bucket effects and a quadratic in risk. Hospital effects are estimated on SAVR claims and imposed on TAVR claims, so a hospital's TAVR level is its SAVR level plus the common TAVR terms. TAVR charges also depend on the hospital's years of experience with the procedure, $E_{h,t} = \max\{0,\, t - \tau_h\}$, entered as indicators. Because the hospital effect is held at its SAVR value, the experience profile is identified from differences in experience across hospitals within a year. The assumption is that a hospital's TAVR charges relative to its SAVR charges do not vary with its adoption date except through experience. The hospital observes a margin shock when it assigns the modality, with the variance of the regression residuals, and it is integrated out. For hospitals without TAVR claims, the hospital effects in the modality, charge and reimbursement regressions are not identified; we set them to their fitted values from a regression of the identified effects on

standardized log beds and region. Costs follow the cost-to-charge approach of Reddy et al. (2025),

$$\text{cost}^{m}_{h,t}(\theta) \;=\; c_{h,t} \cdot ch^{m}_{h,t}(\theta), \qquad m \in \{T, S\}, \tag{9}$$

where $ch^{m}_{h,t}(\theta)$ is charges per case and $c_{h,t}$ is the hospital's ratio of total costs to total charges in fiscal year t from its Medicare cost report. The ratios average 0.25 and fall from 0.28 in 2012 to 0.22 in 2019; missing hospital-years take the year's median.[15] Experience enters the charge predictions and the risk composition of admitted patients.

Because adoption changes which patients choose h, both admissions and the composition of admitted patients depend on a_h. Let $H^{a_h}_{h,t}$ denote the risk distribution of the patients hospital h admits under state a_h, the market distribution reweighted by $M_t\, s_{h,t}(\theta; a_h, A_{-h,t})$, and let the per-admission profit be

$$\pi^{\text{adm}}_{h,t}(\theta; a_h) \;=\; m^{S}_{h,t}(\theta) \;+\; p_{h,t}(\theta; a_h)\, \Delta_{h,t}(\theta), \qquad p_{h,t}(\theta; 0) \equiv 0.$$

Operating profit under a_h is admissions times the mean per-admission profit over the admitted mix,

$$\pi^{op}_{h,t}(a_h) \;=\; Y_{h,t}(a_h) \cdot \underbrace{\mathbb{E}_{\theta \sim H^{a_h}_{h,t}}\left[\pi^{\text{adm}}_{h,t}(\theta; a_h)\right]}_{\mu^{a_h}_{h,t}},$$

so the period-t incremental profit from adoption decomposes into a volume effect and a composition effect,

$$\begin{aligned} \Delta\pi_{h,t}(A_{-h,t}) &\equiv \pi^{op}_{h,t}(1) - \pi^{op}_{h,t}(0) \\ &= \underbrace{\left[Y_{h,t}(1) - Y_{h,t}(0)\right]\, \mu^{0}_{h,t}}_{\text{admissions at pre-adoption mix}} \;+\; \underbrace{Y_{h,t}(1)\, \left[\mu^{1}_{h,t} - \mu^{0}_{h,t}\right]}_{\text{composition shift at adopted mix}} . \end{aligned} \tag{10}$$

The first term is the steering return of Section IV.B in dollars. The second records that after adoption h admits a different risk mix and earns different per-admission dollars through the modality rule $p_{h,t}(\theta; 1)$ and the margins $\Delta_{h,t}(\theta)$.

[15] The ratio is hospital-wide, so the implied TAVR cost applies the hospital's average markup to a charge that includes the device.

VII.B. The timing problem

Each hospital chooses $\tau_h \in \mathcal{T} = \{t_0, t_0{+}1, \ldots, T, \infty\}$, where ∞ denotes never, $t_0 = 2012$ and $T = 2019$. The present value at τ of adopting at τ, with annual discount factor $\delta = 0.92$, is

$$\Delta\Pi_h(\tau) = \sum_{t=\tau}^{T} \delta^{\,t-\tau} \Delta\pi_{h,t}(A_{-h,t}) \;+\; \delta^{\,T+1-\tau}\, \frac{\max\{0, \bar{\Delta\pi}_h\}}{1-\delta}, \qquad \Delta\Pi_h(\infty) = 0,$$

where $\bar{\Delta\pi}_h$ is the average incremental flow over the last three sample years, so that a program is valued as a perpetuity from 2020 on. Let $K_{h,\tau}$ be the net one-time cost of adopting in year τ, specified below, and let $\varepsilon_{h\tau}$ be i.i.d. type-I extreme value with scale σ_ε. The adoption year solves

$$\tau_h \in \arg\max_{\tau \in \mathcal{T}} \Big\{ \delta^{\,\tau - t_0} \big(\Delta\Pi_h(\tau) - K_{h,\tau} \big) + \varepsilon_{h\tau} \Big\},$$

so that with $v_{h\tau} = \delta^{\,\tau-t_0}\big(\Delta\Pi_h(\tau) - K_{h,\tau}\big)/\sigma_\varepsilon$ and $v_{h,\infty} = 0$ the timing probabilities are multinomial logit over the eight candidate years and the never-adopt option,

$$P_h(\tau) = \frac{\exp v_{h\tau}}{1 + \sum_{\tilde{\tau} \neq \infty} \exp v_{h\tilde{\tau}}}. \tag{11}$$

Profits and the adoption cost are both discounted to t_0, so the bracket is the net present value at t_0 of a cost paid at τ. We estimate the scale σ_ε; it converts net adoption cost into dollars and is identified by how sharply adoption responds to cross-sectional variation in $\Delta\Pi - K$.

Estimation proceeds in two stages. The modality rule is estimated by maximum likelihood on the claims of adopting hospitals, and the profits $\Delta\pi_{h,t}$ are built from it, the margin predictions, and the demand shares. The timing parameters are then estimated by maximum likelihood on observed adoption years, with the profits held fixed.

VII.C. Net adoption cost

We parameterize

$$K_{h\tau} = \kappa_0 + Z_h \beta_Z + X_h \beta_X + \psi_{r(h)} + \lambda_\tau + \nu_h, \qquad \nu_h \sim N(0, \sigma_\nu^2), \tag{12}$$

where Z_h is standardized log beds, X_h is an indicator for teaching status, $\psi_{r(h)}$ are region effects (states, with small states grouped into contiguous regions so that every region has both adopters and hospitals that never adopt), and λ_τ are adoption-year effects. The year

terms are a step after 2012 and a linear drift, $\lambda_\tau = c + g\,(\tau - t_0)$ for $\tau > t_0$ and zero in 2012; the size slope is the same in every year. The permanent effect ν_h is integrated out by Gauss–Hermite quadrature over the choice probabilities (11), so the timing model is a mixed logit. The 97 pivotal-trial sites, which ran TAVR programs in the clinical trials before Medicare coverage, are treated as already in: they enter the demand side and the counterfactuals as adopters from the first year but make no timing choice, so the timing likelihood runs over the remaining 848 hospitals, of which 521 adopt by 2019.

$K_{h\tau}$ enters the index linearly, with free sign. The object $\Delta\Pi_{h\tau}$ measures only the incremental Medicare TAVR–SAVR margin, so what (12) identifies is $K_{h\tau} = F_{h\tau} - B_h$, the gross setup cost net of every return to operating a program that accrues outside that margin: commercial and Medicare Advantage volume, referral capture into the rest of the cardiac service line, physician recruitment, and the non-pecuniary return to being first. Only the difference is identified, since $\Delta\Pi$ and K enter the index solely as $\Delta\Pi - K$, and $K < 0$ records $B > F$. This is the convention of the entry literature (**???**), in which the estimated object is a threshold with free sign; in the dynamic-investment literature (Ryan, 2012; Igami, 2017) a physical construction cost is held positive by functional form. $K_{h\tau}$ is the empirical counterpart of the fixed cost F_i of the two-stage model, and the counterfactuals of Section VIII use it in that role. We impose no hospital fixed effect. Each hospital makes one adoption decision, so a hospital-specific intercept would be one free parameter per observation, and ν_h is the identified random-effects counterpart. Hospital heterogeneity enters chiefly through $\Delta\Pi$, where hospital, year and risk effects in charges, reimbursements and patient risk are identified from hundreds of thousands of claims; K carries one decision per hospital. The 97 trial sites, whose adoption cost was incurred before the observation window, remain in the demand and modality stages.

Table 8 reports both stages. The modality rule is estimated on 315,918 claims at the 613 adopting hospitals whose post-adoption claims include both procedures; at the others every admitted patient receives the same procedure, so the hospital effect is not identified and is set to its projection on size and region, as for hospitals that never adopt. Assignment to TAVR rises steeply with risk, and more steeply each year as the indication widened. The margin gap has no detectable effect, $\hat{\gamma} = -0.02$ (0.17). The 2015 revision changed the gap at every hospital at once and is absorbed by the year effects, so γ is identified only from differences in the gap across hospitals and patients within a year, which are small. The median adopter's net adoption cost is \$5.0M, with an interquartile range of \$2.6M to \$7.8M. The scale $\hat{\sigma}_\varepsilon$ is \$2.58M and the standard deviation of the permanent hospital effect $\hat{\sigma}_\nu$ is

TABLE 8: Supply model estimates

Stage 1: modality		Stage 2: adoption timing	
Parameter	Estimate	Parameter	Estimate
γ (margin gap)	-0.022	σ_ε	25.832
	(0.168)	σ_ν	0.060
β_θ (risk)	1.086	κ_0	58.705
	(0.024)	size (log beds, standardized)	-32.560
risk $\times$ 2013	0.176	experience, X_1	-9.033
	(0.032)	λ 2013	22.195
risk $\times$ 2014	0.324	λ 2014	24.253
	(0.032)	λ 2015	26.311
risk $\times$ 2015	0.539	λ 2016	28.370
	(0.032)	λ 2017	30.428
risk $\times$ 2016	0.664	λ 2018	32.486
	(0.032)	λ 2019	34.545
risk $\times$ 2017	0.921	hospitals	945
	(0.033)	adopters 2012–2019	521
risk $\times$ 2018	1.092	log likelihood	-1,475.8
	(0.034)		
risk $\times$ 2019	0.984		
	(0.034)		
year 2013	0.624		
year 2014	1.046		
year 2015	1.516		
year 2016	1.977		
year 2017	2.519		
year 2018	2.820		
year 2019	3.088		
claims	315,918		
hospitals	613		
log likelihood	-162,715.3		

Notes: Standard errors (Stage 1) in parentheses, conditional on the first-step charge, payment and risk estimates and the fixed costs; none for Stage 2. σ_ε, σ_ν, κ_0 and the size, experience and λ terms are in units of \$100,000, so 8.10 is \$0.81 million. Size enters with one slope, the same in every year. The 97 pivotal-trial sites that performed TAVR before fiscal 2012 are treated as already in and make no timing choice; the adopter count excludes them.

\$0.01M: once size, teaching and region are accounted for, there is almost no permanent unobserved heterogeneity in adoption cost, and the remaining dispersion is the date-specific shock $\varepsilon_{h\tau}$, which varies across a hospital's candidate years. The distribution of K separates by adoption status in the direction the model requires (Figure 7). The median hospital that never adopts faces \$9.4M against \$5.0M for one that does, and one hospital that fails to adopt has a negative estimate, at $-$\$1.3M. Net adoption cost is negative for 10 percent of adopters, 53 of 521, and one of the 327 never-adopters.[16]

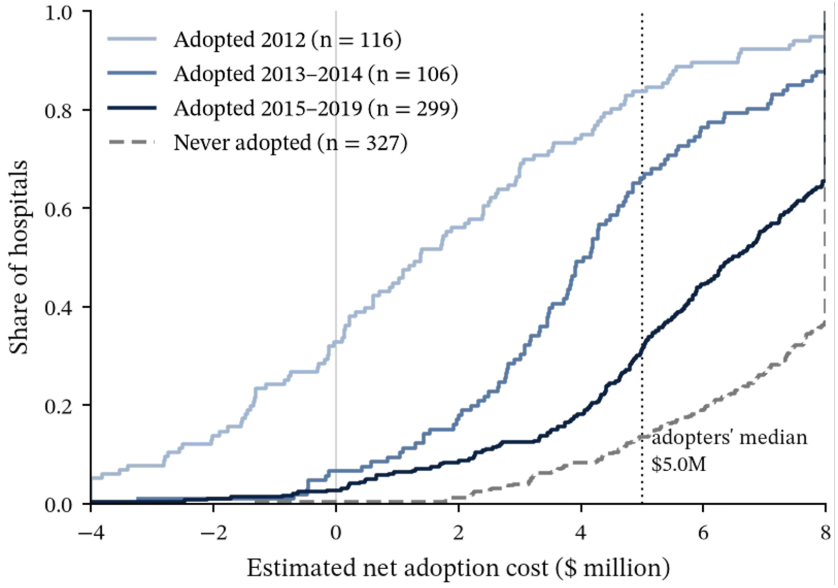


Figure 7: Estimated net adoption cost by adoption cohort. Each curve is the cumulative distribution of $K_{h\tau}$ across hospitals, in millions of dollars, read at the adopter's own adoption year or, for hospitals that never adopt, at the average across candidate years. The dotted line marks the adopters' median, \$5.0M. The 2012 cohort excludes the 97 trial sites treated as in before the window.

Size carries most of the observable variation in adoption cost. A hospital one standard deviation larger in log beds faces a cost lower by \$3.3M, and teaching hospitals by a further \$0.9M. The adoption-year effects are positive, \$2.2M in 2013 relative to 2012 and \$3.5M by 2019. We treat the year profile as a nuisance term: it absorbs whatever made 2012 attractive beyond the margins in the model, including the release at coverage of hospitals that were equipped and willing but could not bill.

[16] $K_{h\tau}$ is read at each adopter's own adoption year and, for hospitals that never adopt, at the average across candidate years, the measure the counterfactuals use. For hospitals with a negative estimate, the Medicare fee-for-service margin cannot rationalize adoption: their objective includes returns outside that margin, which the model does not measure separately. These hospitals are predicted to adopt, so they are never offered the lump sum, which targets predicted non-adopters, and they enter the counterfactuals mainly through the rate increment paid on their TAVR cases.

Figure 8 compares observed and model-implied entry by fiscal year. The model matches the 2012 wave and the total, 521 adopters observed against 521.7 predicted, and reproduces the absence of a second wave in 2015. It smooths the post-2015 path, under-predicting entry in 2015 to 2017 and over-predicting it in 2018 and 2019; the size panel shows that the late over-prediction comes from large hospitals.

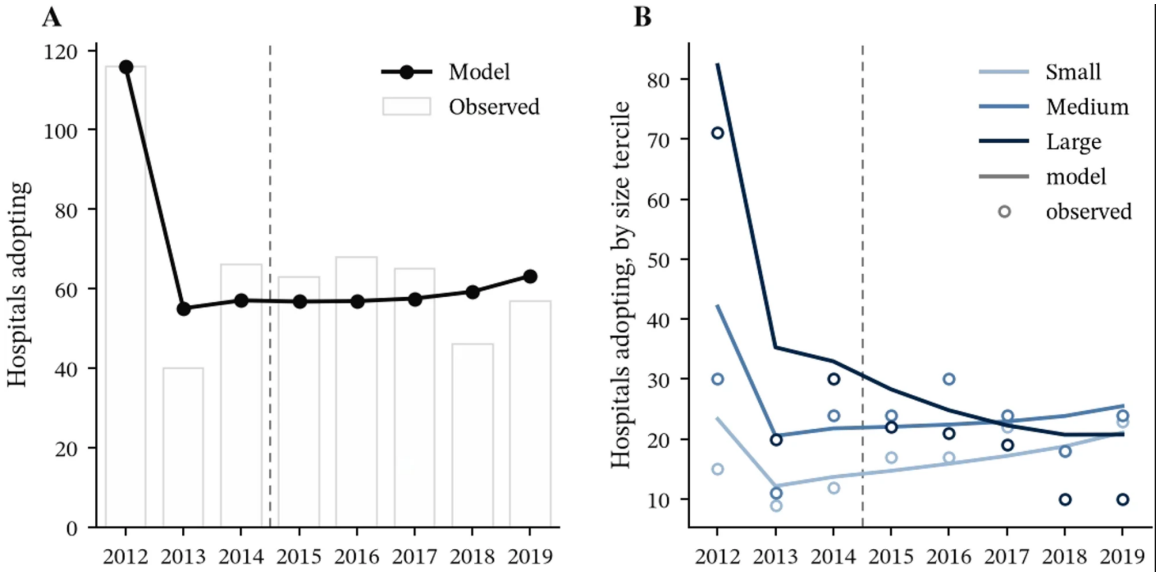


Figure 8: Model fit. Panel A shows hospitals adopting TAVR in each fiscal year, observed (bars) and implied by the estimated timing probabilities (line), for the 848 hospitals that choose when to enter, of which 521 adopt by 2019; the 97 trial sites in before the window are excluded. Panel B shows the same by tercile of hospital size, model (lines) and observed (circles). The dashed line marks the fiscal 2015 payment revision.

Section VIII takes three objects from these estimates: the timing probabilities $P_{h\tau}$, the net adoption cost $K_{h\tau}$, and each hospital's predicted probability of never adopting, which the targeted transfer conditions on.

VIII. Counterfactual policy design

The adoption-timing model of Section VII delivers, for every hospital, the value of adopting at each candidate date and a net adoption cost in dollars; the demand model of Section V delivers how patients move when access changes and what that movement is worth to them in travel time. Together they allow us to re-solve every hospital's entry decision and every patient's hospital choice under any payment rule, and so to compute who adopts, when, who is treated where, what patients gain, and what the payer spends. We use them to ask

whether a planner with the budget Medicare spent could have reached higher diffusion. The answer is yes: at the observed outlay of $502M the frontier reaches 634 hospitals against the observed 619, or 22.1 thousand patients a year against 19.6 thousand. At the observed outlay and below, the best policy for both objectives delays the rate rise and spends the saving on the travel voucher at full strength and on a lump sum targeted on predicted non-adopters, large for coverage and small or zero for patients. With more money than Medicare spent, the patient frontier advances the rise to 2012, while the coverage frontier keeps it delayed. The planner has three instruments, the year in which the rate rises, a lump sum paid at adoption to hospitals that have not entered, and a travel voucher paid to patients who receive TAVR, and three objectives, the number of hospitals that offer TAVR, the number of patients treated, and patient welfare measured in travel time, each evaluated at a given outlay.

VIII.A. What the two-stage model says about each instrument

The two entry conditions of Section IV.B, (1) and (2), are enough to say what each instrument does and where its cost falls; Appendix B states and proves the results. The planner observes adoption dates and hospital observables but not the fixed cost F_i, which is private information. The three instruments act on different terms of the conditions: an earlier rate rise pays the revised rate in stage 1 and acts on the stage-1 loss; a lump sum paid at the update acts on the fixed cost of late entrants; a voucher enlarges the contestable pool and so acts on the steering rent of early entrants and the residual return of late entrants at once.

The first result is sorting. Under the observed policy the adoption date is nondecreasing in F_i: hospitals with $F_i \leq F_i^E$ enter at the unrevised rate, those with $F_i^E < F_i \leq F_i^L$ enter at the update, and the rest never enter (Lemma 1). A hospital that enters at the unrevised rate reveals that its rent to early entry covers its stage-1 loss and the cost of paying F_i early. The delayed rate therefore separates three groups that need three different things: early entrants need nothing, hospitals that enter at the update need the revised rate, and hospitals that never enter need something the rate does not provide. An earlier rise reaches almost none of the last group. It replaces the stage-1 loss by a gain for hospitals that already enter, moves forward the entry of some hospitals that would have entered at the update, and recruits from the never-adopters only those whose shortfall at the revised rate is smaller than one stage of profit at that rate (Corollary 1). Most of its outlay goes to hospitals that had already entered.

The second result concerns a transfer conditioned on the adoption date. Reaching the hospitals that never enter requires a payment that early entrants do not collect, and the natural mechanism is a menu: enter early without a transfer, or enter at the update and

receive s. Offered to every hospital that enters at the update, the transfer recruits hospitals with $L_i < F_i \leq L_i + s$ but also delays the early entrants nearest their cutoff, which give up little by waiting; to first order in s the mass delayed exceeds the mass recruited by the factor $\beta/(1-\beta)$ times the ratio of the fixed-cost densities at the two cutoffs when a late cohort exists (Proposition 1). Because the payoff to waiting is discounted only once while the cost of paying F_i early is $(1-\beta)F_i$, a dollar of transfer moves $\beta/(1-\beta)$ dollars of fixed cost across the early cutoff but only one dollar across the late cutoff, so an untargeted transfer can lower the number of patients treated while raising the number of hospitals. Targeting removes the problem. If $F_i = \varphi(x_i) + \nu_i$ and the transfer is offered only to hospitals with $\varphi(x_i)$ above their late cutoff, those predicted never to enter, a hospital that would have entered unaided is never offered the transfer and has nothing to wait for; as the dispersion of ν_i vanishes the delayed set vanishes and a transfer equal to each hospital's shortfall recruits every never-adopter with no rent (Proposition 2). The delayed rate is what makes targeting feasible: entry at the unrevised rate reveals which hospitals have a return to early entry, and the estimated model of Section VII is the mapping from observables to adoption cost recovered from that sorting. The counterfactuals take that mapping as known, so they price the cost of delaying payment and not the value of the information the delay produced.

The third instrument is the voucher. A voucher paid to patients who travel to an adopter enlarges the pool an adopter can attract and so raises the return to entry at both dates. Unlike the transfer, it is not conditioned on the adoption date, so it gives an early entrant no reason to wait. It also moves patients between adopters, so a hospital whose patients are drawn to a subsidized rival can lose. Its cost is of a different kind. Every patient who travels collects it, whether or not the voucher changed her choice, and the cash is a transfer that does not enter welfare. Whether a voucher dominates a targeted transfer at a given budget is therefore an empirical question about the share of voucher cash that is inframarginal against the share of the transfer that is rent, and the frontier answers it.

VIII.B. Implementation

In the model of Section VII the counterpart of $L_i - F_i$ at adoption date τ is $\Delta\Pi_h(\tau) - K_{h\tau}$, and its systematic part plays the role of $\varphi(x_h)$. The permanent unobservable is negligible, $\hat{\sigma}_\nu$ is \$0.01M (Table 8), so the permanent part of adoption cost is a function of observables, as in the limit of Proposition 2. The date-specific shock is not: with $\hat{\sigma}_\varepsilon$ at \$2.58M, adoption dates are probabilistic, and a lump sum targeted on hospitals predicted not to adopt still reaches some that would have entered unaided and gives some of them a reason to wait.

We do not size transfers to an incentive constraint; every counterfactual re-solves adoption timing, so the entry a transfer delays is counted in its outcomes and priced on the frontier.

Each policy is a triple (t_R, S, k). The 2015 payment rise keeps its observed size and takes effect in year $t_R \in \{2012, \ldots, 2019\}$; after 2019 it is in force under every t_R. From t_R on, a lump sum S is paid at adoption to hospitals that the model predicts, with probability above one half, will not adopt by 2019 under (t_R, k), and that would sustain at least thirty TAVR cases a year after adoption. The voucher pays a share k of the extra minutes a TAVR patient travels when her nearest hospital has no program, at \$0.50 per minute, and enters patients' utility through the travel coefficient of Section V; it is paid to TAVR recipients only. Patients choose a hospital before the heart team assigns the procedure, so in hospital choice the voucher enters at its expected value, the subsidized minutes times the share of patients admitted to adopters that year who receive TAVR, and the same share scales the cash in the outlay. The demand model has no price variation to value a minute in dollars. The observed policy is $(2015, 0, 0)$. The 97 trial sites operate under every policy and are never offered the lump sum. Outlay is the present value at 2012 of the payment increment on TAVR cases from t_R, the lump sums and the voucher cash. Outlay counts spending through 2019. After 2019 the revised rate is in force under every policy, so later spending differs across policies only through the programs a policy adds, which the outlay omits. The change in base payments Medicare makes because a policy changes who receives TAVR is reported separately as induced reimbursement and enters neither the outlay nor welfare. We sweep t_R, a grid of lump sums and $k \in \{0, 0.1, 0.2, 0.3, 0.5, 0.75, 1\}$, and for each budget report the best policy for each objective. Each experiment re-solves the adoption-timing problem of Section VII under the modified profit stream at the infinite-horizon estimates, holding fixed the cost parameters, the modality model, market sizes and rivals' adoption paths. TAVR counts are net of patients a new program takes from existing TAVR hospitals; otherwise the exercise is partial equilibrium.

VIII.C. Results

Figure 9 summarizes the policy sweep and Table 10 reports the best policy at four budgets. The observed policy, the payment rise of 2015, costs \$502 million in present value and leaves 619 hospitals offering TAVR by 2019 and 19.6 thousand TAVR patients a year. At the same outlay the frontier reaches 634 hospitals, 15 more or 2.4 percent of observed coverage, with the rise moved to 2018, a lump sum of \$10.4 million for targeted hospitals, and the voucher at full strength; or 22.1 thousand patients a year, 2.5 thousand or 13 percent more,

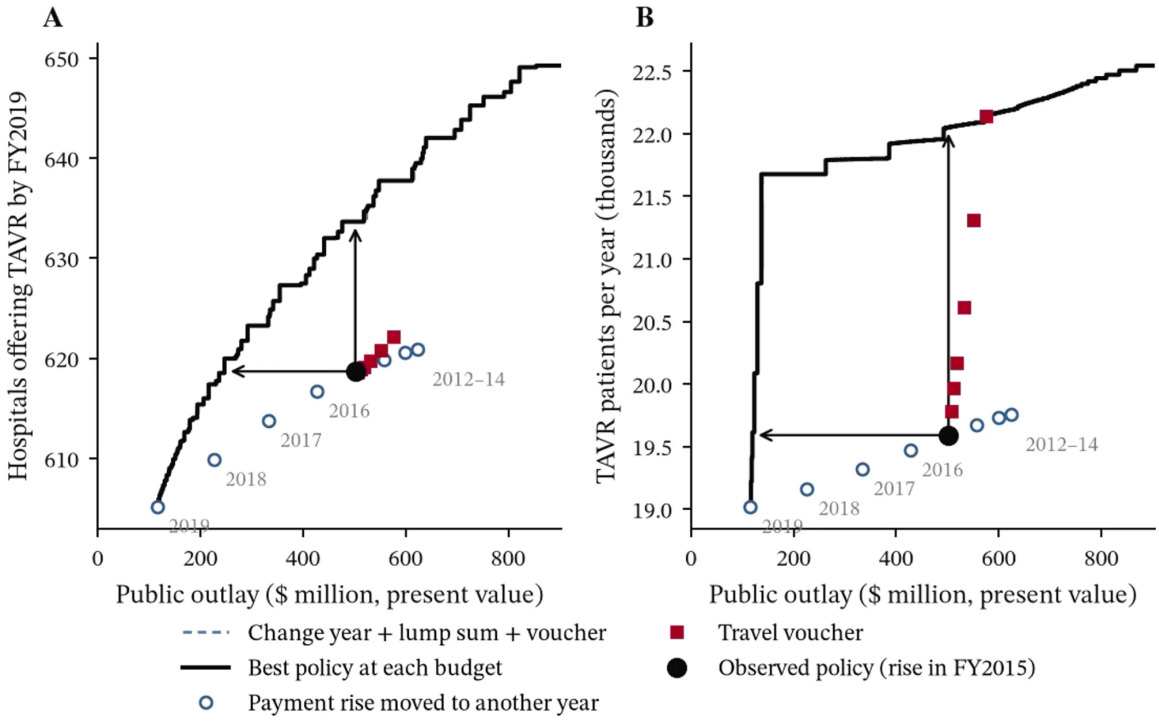


Figure 9: Policy frontier. Each policy combines a year in which the 2015 payment rise takes effect, a lump sum paid at adoption to hospitals the model predicts will stay out, and a travel voucher that pays part of the extra travel time to a TAVR hospital. Panel A reports hospitals offering TAVR by 2019 and panel B TAVR patients per year, each against the public outlay in present value. The solid line is the best policy at each budget, the circles move the payment rise alone, the squares add the voucher alone at rising strength, and the dot is the observed policy. The arrows show the coverage the observed budget buys on the frontier and the budget at which the frontier reaches the observed coverage.

with the rise moved to 2016, a lump sum of $1.0 million, and the voucher. Read the other way, the frontier reaches the observed coverage for about $247 million, half the observed outlay, and the observed number of patients for about $119 million, a quarter of it.

The frontier delays the rate rise for both objectives, and Table 9 shows why. The rise is paid on every TAVR case at every adopter, most of which enter without it, so moving it changes outlay far more than entry. Advancing it from 2015 to 2012 costs $122 million and adds 2.2 hospitals, about $55 million per hospital; delaying it to 2019 saves $387 million and loses 13.5 hospitals, about $29 million per hospital. A lump sum targeted on predicted non-adopters is far cheaper per hospital. With the rise moved to 2018 the schedule costs $226 million and leaves 610 hospitals; the coverage policy at the observed budget spends the remaining $250 million on lump sums and the voucher and reaches 634, about $10.5 million per additional hospital. The lump sum itself, $10.4 million, is close to the median adoption cost of hospitals that never adopt, $9.4 million (Section VII.C). Each year of delay therefore releases money that the lump sum converts into more coverage than the rise bought. This is the sorting result in numbers: hospitals with a return to early entry enter at the unrevised rate, and the rise pays them without changing their decision.

How far the rise is delayed depends on the objective. The lump sum is paid only for adoption from t_R, so t_R also dates the entry it buys. Coverage counts a hospital that enters in 2019 as fully as one that enters in 2012, so the coverage frontier delays the rise to 2018 and spends most of the saving on lump sums. Patients treated per year count every year a program operates, so they value entry that comes earlier, which an earlier rise brings forward; the patient frontier delays the rise only to 2016 and spends little on lump sums. Patients treated in 2019 follow coverage, because volume in the last year is set by the programs open in that year.

The voucher is the instrument for patients. Every policy on the coverage and patient frontiers pays it at full strength. On top of the observed policy it raises patients a year by up to 2.6 thousand and coverage by 3.5 hospitals, for $74 million. With the rise delayed to 2019 and no lump sum, it reaches 21.7 thousand patients a year, 2.1 thousand more than observed, for $136 million: most of the gain in patients treated comes from moving patients to existing programs rather than from new programs. Coverage responds to the voucher only through the larger pool it gives an adopter, so the coverage frontier relies on the lump sum.

The two objectives also respond differently to the budget. The coverage frontier gains 14 hospitals between half and the full observed outlay, 12 more up to one and a half times it, and 7 more up to twice it: 620, 634, 646 and 653 hospitals. It never advances the rise: at

the two larger budgets the best policy moves it to 2018 and 2016 and spends the money on lump sums of $16.4 and $14.1 million. The patient frontier reaches 21.7 thousand patients a year at half the observed outlay and gains little after, 22.1, 22.4 and 22.6 thousand at the three larger budgets; at the two largest it advances the rise to 2012: once the voucher is at full strength, the remaining margin for patients is earlier entry, which an earlier rise buys.

Patient welfare in travel time barely moves: across budgets the best policy changes consumer surplus by at most $0.3 million against outlays in the hundreds of millions, and it never includes the voucher, whose cash is a transfer and does not count in welfare. Travel time is the only welfare the demand model measures; the clinical return to TAVR over SAVR does not enter it, which is why patients treated is the objective that carries the health gain. Induced reimbursement is negative for every frontier policy that raises TAVR: at the schedule without the increment, a TAVR case is paid less than the SAVR case it replaces. Total fiscal cost, outlay plus induced reimbursement, is therefore below the outlay, $370 million for the coverage policy at the observed budget. One channel is shut in these numbers. The estimated modality rule does not respond to the margin gap (Section VII.C), so a rate change moves TAVR volume only through entry and its timing, not through the share of admitted patients a program assigns to TAVR. The 8 percent rise in volume at existing programs at the 2015 revision (Section IV.C) suggests that the counterfactuals understate what an earlier rise buys in patients treated, and so overstate the case for delay when the objective is patients.

The introduction asked whether a planner with the observed budget can reach more hospitals, or treat more patients, than the observed policy. The frontier answers yes for each objective, by a different policy, and the choice between them is the planner's.

IX. Conclusion

In this setting, a payer with the budget Medicare spent could have reached more hospitals and more patients by delaying the rate rise and spending the saving on a lump sum at adoption targeted on predicted non-adopters and on a travel voucher for TAVR patients. The rate rise is almost entirely inframarginal for entry: it pays every case at every adopter, and most adopters enter without it. The delayed rate is also what makes targeting possible, because entry at the unrevised rate reveals which hospitals have a private return to early adoption. How far to delay depends on the objective. Coverage favors delay to 2018 and keeps the rise delayed at every budget we consider; patients treated favor an earlier rise,

TABLE 9: Counterfactual: the 2015 payment change moved to another year

Change year	Outlay	Induced reimb.	Hospitals by 2019	TAVR 2012–19 (k)	TAVR 2019 (k)	CS vs SQ
Infinite horizon (perpetuity after 2019)						
2012	623.4	-6.7	620.9	158.06	36.34	+0.1
2013	599.5	-5.7	620.6	157.86	36.32	+0.1
2014	557.0	-3.4	619.9	157.40	36.28	+0.0
2015 (status quo)	501.9	+0.0	618.7	156.73	36.22	+0.0
2016	427.1	+4.9	616.7	155.77	36.12	-0.0
2017	333.1	+11.2	613.8	154.58	35.98	-0.1
2018	225.9	+18.3	609.9	153.30	35.80	-0.2
2019	115.1	+24.9	605.2	152.15	35.58	-0.2

Notes: $ million, PV at 2012. Outlay: payment increment on TAVR cases, lump sums and voucher cash ($0.50 per minute). Induce reimb.: extra base payment from the change in who gets TAVR, vs status quo; not in the outlay.

TABLE 10: Counterfactual: best policy at each budget

Objective	Budget (×SQ)	Best policy (t_R; S, $M; k)	Outlay ($M)	Value	vs SQ	Induced reimb.
Infinite horizon (perpetuity after 2019)						
Hospitals offering TAVR by 2019	0.5	2019; 8.89; 1.00	246.8	620.0	$+1.3$	-97.9
	1	2018; 10.36; 1.00	475.8	633.7	$+15.0$	-105.6
	1.5	2018; 16.43; 1.00	751.5	646.1	$+27.4$	-108.1
	2	2016; 14.09; 1.00	950.3	652.6	$+33.9$	-125.3
TAVR patients per year, 2012–19 (k)	0.5	2019; 0.00; 1.00	136.4	21.68	$+2.09$	-97.5
	1	2016; 1.03; 1.00	501.5	22.05	$+2.46$	-113.7
	1.5	2012; 2.23; 1.00	751.1	22.38	$+2.79$	-126.9
	2	2012; 7.62; 1.00	963.1	22.64	$+3.05$	-138.4
TAVR patients in 2019 (k)	0.5	2019; 8.89; 1.00	246.8	40.55	$+4.33$	-97.9
	1	2018; 10.36; 1.00	475.8	40.86	$+4.64$	-105.6
	1.5	2018; 16.43; 1.00	751.5	41.15	$+4.93$	-108.1
	2	2016; 14.09; 1.00	950.3	41.33	$+5.11$	-125.3
Consumer surplus vs SQ ($M)	0.5	2018; 3.03; —	250.7	-0.2	-0.16	$+17.2$
	1	2015; 0.00; —	501.9	0.0	$+0.00$	$+0.0$
	1.5	2012; 4.81; —	748.1	0.2	$+0.18$	-20.5
	2	2012; 8.89; —	937.7	0.3	$+0.28$	-32.1

Notes: $ million, PV at 2012; SQ is the observed policy, outlay $501.9M. Outlay: payment increment on TAVR cases, lump sums and voucher cash ($0.50 per minute). Induced reimb.: extra base payment from the change in who gets TAVR, vs status quo; not in the outlay. S: lump sum per adopting hospital ($M); k: share of the extra travel paid.

and, with more money than Medicare spent, a rise advanced to 2012. Most of the gain in patients comes from the travel voucher, which moves patients to existing programs. Two limits deserve statement alongside the recommendation. The cost of the voucher scales with the value of travel time, which the demand model does not identify. And how far this carries beyond TAVR depends on whether adoption elsewhere also earns a return outside the regulated price, which is the primitive the argument rests on.

A. Additional descriptive evidence

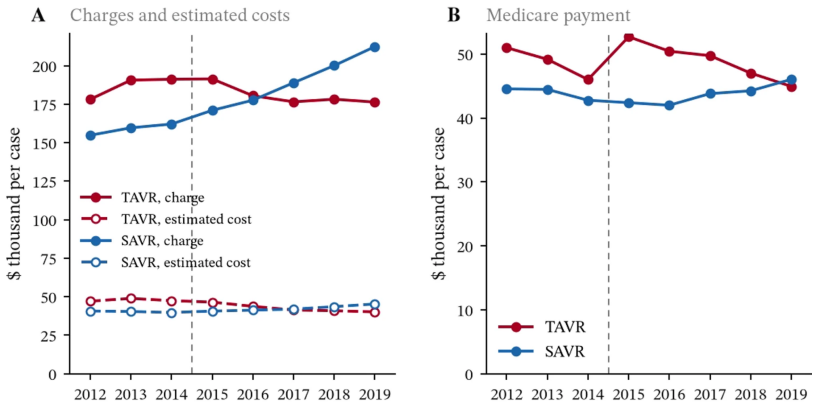


Figure 10: Charges, estimated costs and Medicare reimbursement per case for TAVR and SAVR by fiscal year. Panel A shows charges and costs estimated with hospital-year cost-to-charge ratios; Panel B the Medicare payment.

This appendix reports the two further tests of the hospital response to margins referred to in Section IV.C.

I.A. Selection of high-benefit patients

If early adopters placed greater weight on patient benefit than later entrants, they would concentrate TAVR on patients with higher expected benefit (older or more severe cases). A direct comparison of TAVR patients across hospitals is not feasible because non-adopters have no TAVR patients. We use SAVR, the main substitute, as the benchmark outcome: if early adopters channel high-benefit patients into TAVR more aggressively, SAVR severity and age should fall more sharply at early adopters after reimbursement improves. We control for hospital-year case mix using a leave-one-out mean: for each SAVR patient i, the control variable is the mean age and severity across all *other* AVR patients at the same hospital-year.

Outcomes are in logs. We estimate

$$\log(Y_{iht}) = \beta_0 + \beta_1 \cdot \text{Early Adopter}_h + \beta_2 \cdot \text{Post Reimb}_t + \beta_3 \cdot (\text{Early Adopter}_h \cdot \text{Post Reimb}_t) + \mathbf{Z}'_{ht}\theta + \delta_s + \epsilon_{iht} \tag{13}$$

where Y_{iht} is DRG weight or age for SAVR patient i at hospital h in year t; Early Adopter$_h$ indicates pre-2015 TAVR adoption; Post Reimb$_t$ indicates fiscal years after the reimbursement change; $\mathbf{Z}_{ht}$ collects hospital-level case-mix and structural controls; and δ_s are city fixed effects. Selection of high-benefit patients by early adopters predicts $\beta_3 < 0$.

	log(DRG Weight) (1)	log(Age) (2)
Early Adopter	-0.032^{*}	-0.011
	(0.017)	(0.009)
Post Reimbursement	-0.072^{***}	-0.034^{***}
	(0.015)	(0.008)
Early Adopter · Post Reimbursement	0.046^{**}	-0.001
	(0.015)	(0.008)
Constant	1.415^{***}	3.449^{***}
	(0.308)	(1.544)
City fixed effects	Yes	Yes
Observations	15,232	15,232
R^2	0.082	0.073

Note: Standard errors in parentheses, clustered at the city level.
$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

TABLE 11: Difference-in-Differences: SAVR Patient Characteristics and TAVR Adoption Timing

In Table 11, SAVR DRG weight falls by about 7% system-wide after 2015, consistent with higher-risk patients shifting to TAVR. The Early Adopter $\times$ Post Reimbursement interaction is $+0.046$ and significant: the severity decline at early adopters is roughly 4.7 percentage points *smaller*, not larger, than at later adopters; the age interaction is small and insignificant. The sign goes against the prediction. We treat this evidence as suggestive rather than dispositive, since SAVR severity is an indirect proxy for the unobserved TAVR allocation decision.

I.B. Within-hospital modality reallocation

A purely clinical decision rule should not reallocate the modality mix at a given hospital when relative reimbursement changes, holding the patient pool fixed. Figure 11 shows the within-hospital TAVR share shifting right after 2014. We estimate

$$\log(\text{SAVR}_{h,t}) \;=\; \alpha \;+\; \beta_1\,\text{Adopt}_{h,t} \;+\; \beta_2\,\text{Adopt}_{h,t}\times\mathbf{1}\{t \geq 2015\} \;+\; \gamma_h \;+\; \delta_{t\times\text{HRR}} \;+\; \varepsilon_{h,t}, \tag{14}$$

where $\text{Adopt}_{h,t} = \mathbf{1}\{t \geq T_h\}$ indicates that hospital h is offering TAVR by year t, γ_h are hospital fixed effects, and $\delta_{t\times\text{HRR}}$ are fiscal-year by HRR fixed effects.

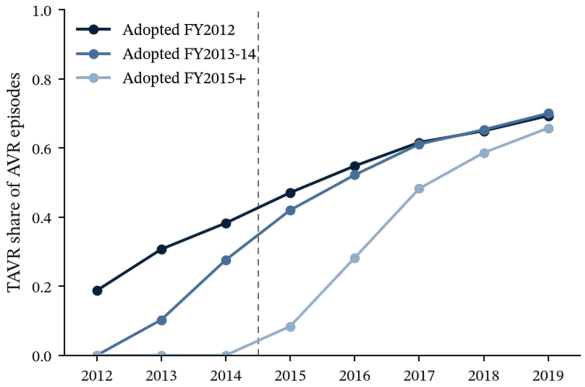


Figure 11: Within-hospital TAVR shares by fiscal year (early adopters)

	$\log(\text{SAVR volume})$
Adopt	0.032 (0.045)
Adopt $\times$ $\mathbf{1}$(fiscal year $\geq$ 2015)	-0.141^{*} (0.060)
Hospital FE	Yes
Fiscal Year $\times$ HRR FE	Yes
Observations	4,900
R^2	0.849

Notes: Standard errors in parentheses, clustered at the hospital level. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

TABLE 12: Effect of TAVR adoption on SAVR volumes

Table 12 shows an asymmetric response. SAVR volume at adopters is essentially unchanged when TAVR margins are weak ($\hat{\beta}_1$ small and insignificant). After 2015, when the relative TAVR margin improves, SAVR falls significantly among adopters ($\hat{\beta}_2 < 0$): the

same hospitals reallocate within their AVR mix toward the higher-margin procedure. This is the empirical basis for the modality policy in Section VII.

B. The planner's problem in the two-stage model

This appendix states and proves the results summarized in Section VIII.A. Hospitals, markets and payoffs are those of Section IV.B. For hospital i in market m write

$$R_i := \chi_i S_m \mu_1^S + \beta \chi_i (S_m - Q_m) \mu_2^S + \beta v_{i,2} \eta_i, \qquad \ell_i := -v_{i,1} \mu_1 > 0, \qquad L_i := \chi_i Q_m \mu_2^S + v_{i,2} \mu_2,$$

for the rent to early entry (steering plus learning), the stage-1 loss, and the stage-2 return of a late entrant. The values of entering early, late, and never are

$$V_i^E = V_i^L + R_i - \ell_i - (1-\beta) F_i, \qquad V_i^L = \beta (L_i - F_i), \qquad V_i^N = 0.$$

Assumption 1. *Within a market, hospitals differ in the fixed cost F_i, which is private information with continuous density g on $[0, \bar{F}]$, and in observables x_i. The planner observes adoption dates and x_i, not F_i.*

Assumption 2. *$(R_i - \ell_i)/(1-\beta) \leq L_i$, and g is positive at F_i^E and L_i.*

Under Assumption 2, $F_i^E = (R_i - \ell_i)/(1-\beta)$ and $F_i^L = L_i$ in Lemma 1, so the late cohort $\mathcal{M} = (F_i^E, L_i]$ is an interval that can be nonempty. This is the case in which the delayed rate sorts hospitals into three groups, and Propositions 1 and 2 are stated under it.

The planner's instruments are: a stage-1 rate increment $\rho \in \{0, \mu_2 - \mu_1\}$, paid on every stage-1 TAVR case at every operating hospital, since a regulated price cannot differ across hospitals; a lump sum $s \geq 0$ paid at t_R to a hospital that enters in stage 2, possibly conditioned on x_i; and a voucher $\theta \geq 0$ per minute of travel, paid to a patient who travels to an adopter other than her nearest hospital. She evaluates a policy on coverage N, the number of hospitals that ever adopt; access A, the discounted number of TAVR cases; welfare W, patients' consumer surplus from hospital access in travel-time units, net of voucher cash; and outlay C, the present value of rate increments, lump sums and voucher cash. For a budget B the frontier is the maximum of each outcome over policies with $C \leq B$. The observed policy is the delayed rate with no transfers, $\rho = s = \theta = 0$.

II.A. Sorting under a delayed rate

Lemma 1 (Sorting)**.** *Under the observed policy the adoption date is nondecreasing in F_i: there are cutoffs $F_i^E \le F_i^L$ with*

$$F_i^E = \min\Big\{\tfrac{R_i-\ell_i}{1-\beta},\ \Delta\pi_{i,1} + \beta\Delta\pi_{i,2}^E\Big\}, \qquad F_i^L = \max\{L_i,\ F_i^E\},$$

such that hospital i enters in stage 1 if $F_i \le F_i^E$ ($\mathcal{E}$), at t_R if $F_i^E < F_i \le F_i^L$ ($\mathcal{M}$), and never if $F_i > F_i^L$ ($\mathcal{S}$).

Proof. $V_i^E - V_i^L$, V_i^L and V_i^E are linear in F_i with slopes $-(1-\beta)$, $-\beta$ and -1. Early entry requires $V_i^E \ge V_i^L$ and $V_i^E \ge 0$, which hold for F_i below the smaller of the two roots, F_i^E. Late entry requires $V_i^L \ge 0$, i.e. $F_i \le L_i$, and $V_i^L \ge V_i^E$. For $F_i > F_i^E$ either $V_i^E < V_i^L$ or $V_i^E < 0$, and in the latter case $V_i^L \ge 0$ implies $V_i^L > V_i^E$; so the late set is $(F_i^E, L_i]$, empty if $L_i \le F_i^E$, which is $(F_i^E, F_i^L]$. The rest never enter. □

Corollary 1 (Earlier update)**.** *Paying the revised rate in stage 1 ($\rho = \mu_2 - \mu_1$) replaces the stage-1 loss ℓ_i by a gain $v_{i,1}\mu_2$ and leaves L_i unchanged. It moves hospitals in $\mathcal{M}$ with $F_i \le (R_i + v_{i,1}\mu_2)/(1-\beta)$ from t_R to stage 1; it recruits from $\mathcal{S}$ only hospitals with $L_i < F_i \le \Delta\pi_{i,1} + v_{i,1}(\mu_2 - \mu_1) + \beta\Delta\pi_{i,2}^E$, those whose shortfall at the revised rate is smaller than one stage of profit at that rate plus the early-entry rent; and it pays $v_{i,1}(\mu_2 - \mu_1)$ to every hospital in $\mathcal{E}$, which changes no decision. The payment to $\mathcal{E}$ is increasing in stage-1 volume.*

II.B. Screening on the adoption date

Proposition 1 (Gate transfer)**.** *Under Assumption 2, offer $s > 0$ to every hospital that enters at t_R. Then:*

(i) hospitals in $\mathcal{R} = \{L_i < F_i \le L_i + s\} \subset \mathcal{S}$ are recruited and enter at t_R;

(ii) hospitals in $\mathcal{D} = \{F_i^E - \frac{\beta}{1-\beta}s < F_i \le F_i^E\} \subset \mathcal{E}$ delay entry from stage 1 to t_R;

(iii) $\mathcal{D}$ is non-empty for every $s > 0$, so no positive transfer is exactly incentive compatible;

(iv) to first order in s, the mass delayed exceeds the mass recruited by the factor $\frac{\beta}{1-\beta}\,\frac{g(F_i^E)}{g(L_i)}$, and the TAVR cases lost in stage 1 from $\mathcal{D}$ exceed those gained in stage 2 from $\mathcal{R}$ when $v_{i,1}\, g(F_i^E)\, \beta \ge v_{i,2}\, g(L_i)\, (1-\beta)$.

Proof. With the transfer, $V_i^L = \beta(L_i - F_i + s)$. A hospital in $\mathcal{S}$ enters iff $F_i \leq L_i + s$, which is (i). A hospital in $\mathcal{E}$ keeps entering early iff $V_i^E \geq V_i^L$, i.e. $R_i - \ell_i - (1-\beta)F_i \geq \beta s$, i.e. $F_i \leq F_i^E - \beta s/(1-\beta)$, which is (ii) and, since $g > 0$, (iii). For (iv), the masses are $g(L_i)\, s$ and $g(F_i^E)\, \beta s/(1-\beta)$ to first order; each recruit adds $\beta v_{i,2}$ discounted cases and each delayed hospital loses $v_{i,1}$ stage-1 cases. □

Suppose $F_i = \varphi(x_i) + \nu_i$, with ν_i independent of x_i, mean zero and scale σ_ν, and $\varphi(x_i)$ continuously distributed. A *targeted* transfer is offered only to hospitals with $\varphi(x_i) > L_i$, those predicted never to enter.

Proposition 2 (Targeting)**.** *Under Assumption 2 and the targeted transfer, the delayed set is $\mathcal{D} \cap \{\varphi(x_i) > L_i\}$. Its mass converges to zero as $\sigma_\nu \to 0$, for any s. In the limit, the personalized transfer $s(x_i) = \varphi(x_i) - L_i$ recruits all of $\mathcal{S}$ with no rent and no delay.*

Proof. A delayed hospital has $F_i \leq F_i^E \leq L_i$ and $\varphi(x_i) > L_i$, so $\nu_i < F_i^E - \varphi(x_i) < 0$ with $|\nu_i| > \varphi(x_i) - L_i > 0$. As $\sigma_\nu \to 0$ this event has vanishing probability because $\varphi(x_i)$ has no atom at L_i. At $\sigma_\nu = 0$, $s(x_i)$ leaves every hospital in $\mathcal{S}$ exactly indifferent and is offered to no hospital outside it. □

References

ACEMOGLU, D. AND A. FINKELSTEIN (2008): "Input and Technology Choices in Regulated Industries: Evidence from the Health Care Sector," *Journal of Political Economy*, 116, 837–880.

ACKERBERG, D. A. AND G. GOWRISANKARAN (2006): "Quantifying equilibrium network externalities in the ACH banking industry," *RAND Journal of Economics*, 37, 738–761.

AGGARWAL, A., D. LEWIS, M. MASON, R. SULLIVAN, AND J. VAN DER MEULEN (2017): "Effect of patient choice and hospital competition on service configuration and technology adoption within cancer surgery: a national, population-based study," *The Lancet Oncology*, 18, 1445–1453.

ALEXANDER, D. AND M. SCHNELL (2024): "The Impacts of Physician Payments on Patient Access, Use, and Health," *American Economic Journal: Applied Economics*, 16, 142–177.

BACH, D. S., D. SIAO, S. E. GIRARD, C. DUVERNOY, B. D. MCCALLISTER, AND S. K. GUALANO (2009): "Evaluation of Patients with Severe Symptomatic Aortic Stenosis Who Do Not Undergo Aortic Valve Replacement," *Circulation: Cardiovascular Quality and Outcomes*, 2, 533–539.

BARON, D. P. AND R. B. MYERSON (1982): "Regulating a Monopolist with Unknown Costs," *Econometrica*, 50, 911–930.

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63.

BESLEY, T. AND S. COATE (1992): "Workfare versus Welfare: Incentive Arguments for Work Requirements in Poverty-Alleviation Programs," *American Economic Review*, 82, 249–261.

BLOOM, N., C. PROPPER, S. SEILER, AND J. VAN REENEN (2015): "The Impact of Competition on Management Quality: Evidence from Public Hospitals," *Review of Economic Studies*, 82, 457–489.

CABRAL, M., C. M. CAREY, AND S. MILLER (2025): "The Impact of Provider Payments on Health Care Utilization of Low-Income Individuals: Evidence from Medicare and Medicaid," *American Economic Journal: Economic Policy*, 17, 106–143.

CAPPS, C., D. DRANOVE, AND M. SATTERTHWAITE (2003): "Competition and Market Power in Option Demand Markets," *RAND Journal of Economics*, 34, 737–763.

CARROLL, J. D., M. J. MACK, S. VEMULAPALLI, H. C. HERRMANN, T. G. GLEASON, G. HANZEL, G. M. DEEB, V. H. THOURANI, D. J. COHEN, N. DESAI, A. J. KIRTANE, S. FITZGERALD, J. MICHAELS, C. KROHN, F. A. MASOUDI, R. G. BRINDIS, AND J. E. BAVARIA (2020): "STS-ACC TVT Registry of Transcatheter Aortic Valve Replacement," *Journal of the American College of Cardiology*, 76, 2492–2516.

CENTERS FOR MEDICARE & MEDICAID SERVICES (2014): "Medicare Program; Hospital Inpatient Prospective Payment Systems for Acute Care Hospitals and the Long-Term Care Hospital Prospective Payment System and Fiscal Year 2015 Rates; Final Rule," Federal Register 79(163): 49854–50449, CMS-1607-F.

CHANDRA, A., A. FINKELSTEIN, A. SACARNY, AND C. SYVERSON (2016): "Health Care Exceptionalism? Performance and Allocation in the US Health Care Sector," *American Economic Review*, 106, 2110–2144.

CHANDRA, A. AND J. SKINNER (2012): "Technology Growth and Expenditure Growth in Health Care," *Journal of Economic Literature*, 50, 645–680.

CLEMENS, J. AND J. D. GOTTLIEB (2014): "Do Physicians' Financial Incentives Affect Medical Treatment and Patient Health?" *American Economic Review*, 104, 1320–1349.

COOPER, Z., S. GIBBONS, S. JONES, AND A. MCGUIRE (2011): "Does Hospital Competition Save Lives? Evidence from the English NHS Patient Choice Reforms," *The Economic Journal*, 121, F228–F260.

DAFNY, L. S. (2005): "How Do Hospitals Respond to Price Changes?" *American Economic Review*, 95, 1525–1547.

DAMLUJI, A. A., M. FABBRO, R. H. EPSTEIN, M. RIVERA, A. NAYAK, C. E. ALFONSO, M. G. COHEN, AND J. D. CARROLL (2020): "Transcatheter Aortic Valve Replacement in Low-Population Density Areas: Assessing Healthcare Access for Older Adults With Severe Aortic Stenosis," *Circulation: Cardiovascular Quality and Outcomes*, 13, e006245.

DE GROOTE, O. AND F. VERBOVEN (2019): "Subsidies and Time Discounting in New Technology Adoption: Evidence from Solar Photovoltaic Systems," *American Economic Review*, 109, 2137–2172.

DEVERS, K. J., L. R. BREWSTER, AND L. P. CASALINO (2003): "Changes in Hospital Competitive Strategy: A New Medical Arms Race?" *Health Services Research*, 38, 447–469.

DRANOVE, D., C. GARTHWAITE, C. HEARD, AND B. WU (2022): "The economics of medical procedure innovation," *Journal of Health Economics*, 81.

DRANOVE, D., M. SHANLEY, AND C. SIMON (1992): "Is Hospital Competition Wasteful?" *RAND Journal of Economics*, 23, 247–262.

DUNNE, T., S. D. KLIMEK, M. J. ROBERTS, AND D. Y. XU (2013): "Entry, exit, and the determinants of market structure," *RAND Journal of Economics*, 44, 462–487.

ECRI INSTITUTE (2013): "ECRI Institute Unveils Costs and Configurations for Top Hybrid OR Imaging Systems," Press release, accessed October 30, 2026.

EINAV, L., A. FINKELSTEIN, AND N. MAHONEY (2018): "Provider Incentives and Healthcare Costs: Evidence from Long-Term Care Hospitals," *Econometrica*, 86, 2161–2219.

ELIASON, P. J., P. L. E. GRIECO, R. C. MCDEVITT, AND J. W. ROBERTS (2018): "Strategic Patient Discharge: The Case of Long-Term Care Hospitals," *American Economic Review*, 108, 3232–3265.

FINKELSTEIN, A. (2007): "The Aggregate Effects of Health Insurance: Evidence from the Introduction of Medicare," *Quarterly Journal of Economics*, 122, 1–37.

FORNELL, D. (2011): "How to Plan for a Hybrid OR," Imaging Technology News, accessed October 30, 2026.

FUDENBERG, D. AND J. TIROLE (1985): "Preemption and Rent Equalization in the Adoption of New Technology," *Review of Economic Studies*, 52, 383–401.

GAYNOR, M., K. HO, AND R. J. TOWN (2015): "The Industrial Organization of Health-Care Markets," *Journal of Economic Literature*, 53, 235–284.

GAYNOR, M., R. MORENO-SERRA, AND C. PROPPER (2013): "Death by Market Power: Reform, Competition, and Patient Outcomes in the National Health Service," *American Economic Journal: Economic Policy*, 5, 134–166.

GAYNOR, M., C. PROPPER, AND S. SEILER (2012): "Free to Choose? Reform and Demand Response in the English National Health Service," Working Paper 18574, National Bureau of Economic Research.

——— (2016): "Free to Choose? Reform, Choice, and Consideration Sets in the English National Health Service," *American Economic Review*, 106, 3521–3557.

GOOLSBEE, A. AND A. PETRIN (2004): "The Consumer Gains from Direct Broadcast Satellites and the Competition with Cable TV," *Econometrica*, 72, 351–381.

GRILICHES, Z. (1957): "Hybrid Corn: An Exploration in the Economics of Technological Change," *Econometrica*, 25, 501–522.

GROENEVELD, P. W., L. YANG, A. SEGAL, P. KARACA-MANDIC, AND G. P. KANTER (2020): "The Effects of Market Competition on Cardiologists' Adoption of Transcatheter Aortic Valve Replacement," *Medical Care*, 58, 996–1003.

HO, K. AND A. PAKES (2014): "Hospital Choices, Hospital Prices, and Financial Incentives to Physicians," *American Economic Review*, 104, 3841–3884.

HOAGLAND, A. (2024): "Innovations and Inequities in Access to Medical Services," Working Paper.

HORN, D., A. SACARNY, AND A. ZHOU (2022): "Technology adoption and market allocation: The case of robotic surgery," *Journal of Health Economics*, 86, 102672, december 2022.

HUGHES, J. E. AND M. PODOLEFSKY (2015): "Getting Green with Solar Subsidies: Evidence from the California Solar Initiative," *Journal of the Association of Environmental and Resource Economists*, 2, 235–275.

IGAMI, M. (2017): "Estimating the Innovator's Dilemma: Structural Analysis of Creative Destruction in the Hard Disk Drive Industry, 1981–1998," *Journal of Political Economy*, 125, 798–847.

IUNG, B., A. CACHIER, G. BARON, ET AL. (2005): "Decision-Making in Elderly Patients with Severe Aortic Stenosis: Why Are So Many Denied Surgery?" *European Heart Journal*, 26, 2714–2720.

KESSLER, D. P. AND M. B. MCCLELLAN (2000): "Is Hospital Competition Socially Wasteful?" *Quarterly Journal of Economics*, 115, 577–615.

LAFFONT, J.-J. AND J. TIROLE (1993): *A Theory of Incentives in Procurement and Regulation*, Cambridge, MA: MIT Press.

LANGER, A. AND D. LEMOINE (2022): "Designing Dynamic Subsidies to Spur Adoption of New Technologies," *Journal of the Association of Environmental and Resource Economists*, 9, 1197–1234.

LEAGUE, R. J. (2023): "Regulation and Diffusion of Innovation Under Information Spillovers: The Case of New Medical Procedures," Tech. rep.

LEON, M. B., C. R. SMITH, M. MACK, AND ET AL. (2010): "Transcatheter Aortic-Valve Implantation for Aortic Stenosis in Patients Who Cannot Undergo Surgery," *New England Journal of Medicine*, 363, 1597–1607.

LUFT, H. S., J. C. ROBINSON, D. W. GARNICK, S. C. MAERKI, AND S. J. MCPHEE (1986): "The Role of Specialized Clinical Services in Competition among Hospitals," *Inquiry*, 23, 83–94.

MACK, M. J., J. M. BRENNAN, R. BRINDIS, J. CARROLL, F. EDWARDS, F. GROVER, D. SHAHIAN, E. M. TUZCU, E. D. PETERSON, J. S. RUMSFELD, ET AL. (2013): "Outcomes Following Transcatheter Aortic Valve Replacement in the United States," *JAMA*, 310, 2069–2077.

MCCARTHY, F. H., D. C. SAVINO, C. R. BROWN, J. E. BAVARIA, V. KINI, D. D. SPRAGAN, T. R. DIBBLE, H. C. HERRMANN, S. ANWARUDDIN, J. GIRI, P. W. GROENEVELD, AND N. D. DESAI (2017): "Cost and contribution margin of transcatheter versus surgical aortic valve replacement," *Journal of Thoracic and Cardiovascular Surgery*, 154, 1872–1880.

NATHAN, A. S., L. YANG, N. YANG, L. A. EBERLY, S. A. M. KHATANA, E. J. DAYOUB, S. VEMULAPALLI, H. JULIEN, D. J. COHEN, B. K. NALLAMOTHU, S. J. BARON, A. C. FANAROFF, R. K. WADHERA, N. R. DESAI, J. GIRI, AND P. W. GROENEVELD (2021): "Socioeconomic and Geographic Characteristics of Hospitals Establishing Transcatheter Aortic Valve Replacement Programs, 2012–2018," *Circulation: Cardiovascular Quality and Outcomes*, 14, e008260.

NICHOLS, A. L. AND R. J. ZECKHAUSER (1982): "Targeting Transfers through Restrictions on Recipients," *American Economic Review Papers and Proceedings*, 72, 372–377.

RAVAL, D. AND T. ROSENBAUM (2021): "Why is Distance Important for Hospital Choice? Separating Home Bias from Transport Costs," *Journal of Industrial Economics*, 69, 338–368.

REARDON, M. J., N. M. V. MIEGHEM, J. J. POPMA, AND ET AL. (2017): "Surgical or Transcatheter Aortic-Valve Replacement in Intermediate-Risk Patients," *New England Journal of Medicine*, 376, 1321–1331.

REDDY, K. P., P. W. GROENEVELD, J. GIRI, A. C. FANAROFF, AND A. S. NATHAN (2023): "Economic Considerations in Access to Transcatheter Aortic Valve Replacement," *Contemporary Reviews in Interventional Cardiology*, forthcoming or check for exact volume/issue if available.

REDDY, K. P., K. SHULTZ, L. A. EBERLY, S. A. M. KHATANA, A. C. FANAROFF, D. J. KUMBHANI, S. ELMARIAH, P. FIORILLI, H. HERRMANN, N. D. DESAI, P. ATLURI, W. Y. SZETO, F. MCCARTHY, D. J. COHEN, P. W. GROENEVELD, J. GIRI, AND A. S. NATHAN (2025): "Contribution margins and utilization of transcatheter aortic valve replacement versus surgical aortic valve replacement in the Medicare population," *American Heart Journal*, 281, 59–70.

REINGANUM, J. F. (1981): "On the Diffusion of New Technology: A Game Theoretic Approach," *Review of Economic Studies*, 48, 395–405.

ROBINSON, J. C. AND H. S. LUFT (1985): "The impact of hospital market structure on patient volume, average length of stay, and the cost of care," *Journal of Health Economics*, 4, 333–356.

RYAN, S. P. (2012): "The Costs of Environmental Regulation in a Concentrated Industry," *Econometrica*, 80, 1019–1061.

SAKOWITZ, S., S. S. BAKHTIYAR, J. CURRY, K. ALI, N. CHERVU, AND P. BENHARASH (2025): "Market Competition and Use of Transcatheter Aortic Valve Replacement in the United States," *JACC: Advances*, 5, 102313.

SCHMIDT-DENGLER, P. (2025): "The Timing of New Technology Adoption: The Case of MRI," *RAND Journal of Economics*, 56, 511–527.

SHELDON, T. L. AND R. DUA (2019): "Measuring the cost-effectiveness of electric vehicle subsidies," *Energy Economics*, 84, 104545.

SMITH, C. R., M. B. LEON, M. J. MACK, AND ET AL. (2011): "Transcatheter versus Surgical Aortic-Valve Replacement in High-Risk Patients," *New England Journal of Medicine*, 364, 2187–2198.

SPRINGEL, K. (2021): "Network Externality and Subsidy Structure in Two-Sided Markets: Evidence from Electric Vehicle Incentives," *American Economic Journal: Economic Policy*, 13, 393–432.

STROBEL, R. J., D. S. LIKOSKY, A. A. BRESCIA, K. M. KIM, X. WU, H. J. PATEL, G. M. DEEB, AND M. P. THOMPSON (2019): "The Effect of Hospital Market Competition on the Adoption of Transcatheter Aortic Valve Replacement," *Annals of Thoracic Surgery*, 109, 473–479.

TANGORRA, L. (2015): "ROI From Transcatheter Aortic Valve Replacement (TAVR)," Diagnostic and Interventional Cardiology (DAIC), accessed October 30, 2026.

VEMULAPALLI, S., J. D. CARROLL, M. J. MACK, Z. LI, D. DAI, A. S. KOSINSKI, D. J. KUMBHANI, C. E. RUIZ, V. H. THOURANI, G. HANZEL, T. G. GLEASON, H. C. HERRMANN, R. G. BRINDIS, AND J. E. BAVARIA (2019): "Procedural Volume and Outcomes for Transcatheter Aortic-Valve Replacement," *New England Journal of Medicine*, 380, 2541–2550.

XING, J., B. LEARD, AND S. LI (2021): "What does an electric vehicle replace?" *Journal of Environmental Economics and Management*, 107, 102432.

YANG, K. K. (2024): "Experience Effects and Technology Adoption: Evidence from Aortic Valve Replacement," Working paper, Duke University.