# 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 7, 2026
Latest version

ABSTRACT. A regulated price breaks the link between the value a technology creates and the return to adopting it. We ask what finances adoption when the price does not reward it, and how a budget-constrained payer should intervene, using transcatheter aortic valve replacement (TAVR) under Medicare. For three years Medicare paid TAVR the rate of the procedure it replaced although it cost more, yet half of eventual adopters entered. A hospital-choice model with travel time as the price shows that adoption paid through patients. A TAVR program raises the hospital's value to patients across the risk distribution, including those who will not receive TAVR, and more to those likely to receive it; as the indication widened, its value to the average valve-replacement patient rose from 15 to 86 minutes of travel. An earlier rate update does little for diffusion: three years earlier, it adds 0.4 percent more hospitals, at $51 million each. Targeting does better: because adoption at the low rate reveals which hospitals earn this return, lump sums for predicted non-adopters and travel vouchers reach 634 hospitals instead of 619 at Medicare's outlay, or treat 35 percent more patients.

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

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

it on a schedule that lags the technology, a new technology can be paid the incumbent's rate 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 Medicare paid TAVR the rate of the open-heart procedure it replaced although it cost more,[1] and entry did not jump when Medicare raised the payment. Second, neither the price nor learning by doing explains this pattern; the evidence points to patient steering. Third, the return to adoption accrues on the hospital's whole program: a program draws patients across the risk distribution, most of whom do not receive TAVR. Fourth, the rate changes the

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

procedure, not the hospital, so the timing of the rate update is a weak lever for diffusion and targeting is what works.

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. The payment changes the procedure rather than the hospital: at the revision, programs already operating shifted patients from SAVR to TAVR and their TAVR volume rose by about 8 percent, while adoption did not jump. Learning by doing is present, but it operates on length of stay rather than on the device, and hospitals with no rival in their HRR enter no earlier than others, although hospitals that treat most of their own patient catchment enter earlier. 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). The patient's chance of receiving TAVR separates the two. It varies with the risk score, with clinical conditions the heart team weighs, and over time as coverage expanded by risk tier. A patient with little chance of TAVR who still prefers an adopter values the hospital rather than the procedure, which we read 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); for a patient likely to receive TAVR, the value of an adopter combines this preference for the hospital with a preference for the procedure. A mixed logit with hospital fixed effects and random travel coefficients (Berry et al., 1995) gives travel costs in line with hospital choice for planned cardiac care (Chandra et al., 2016; Gaynor et al., 2012): each ten minutes of driving lowers the odds of choosing a hospital by about 20 percent, and the nearest hospital carries a large premium. The value of an adopter rises with the chance of TAVR and extends across the risk distribution: patients with about a 20 percent chance value an adopter at 17 to 23 minutes of travel, although four in five of them receive SAVR, and patients with a 5 to 8 percent chance still value it at 9 to 14 minutes. A program therefore steers to the hospital patients most of whom receive the procedure it replaces: before the rate revision, 70 percent of patients who passed over a nearest hospital without a program to be treated at an adopter received SAVR, and adoption raises the hospital's SAVR volume as well as its TAVR volume. The advantage persists, because a program's value grows with its age, so an early entrant keeps its draw ahead of later rivals. As the indication widened, the value of an adopter to the average patient receiving a valve replacement rose from 15 minutes of travel in 2012 to 86 minutes in 2019.

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: \$4.9M 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 \$8.9M 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 shifts patients from SAVR to TAVR within 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 $467M, the observed schedule leaves 619 hospitals offering TAVR by 2019 and 18.1 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 24.4 thousand patients a year with the rise delayed to 2017. 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 $51M per additional hospital, against about $9.2M 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 23.9 thousand patients a year for a third 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). Two papers in this literature study TAVR. Yang (2024) estimates a dynamic equilibrium model of learning by doing, hospital choice and adoption, and finds that Medicare's volume requirements for TAVR programs attain the consumer welfare of unrestricted entry at lower aggregate fixed cost, because the hospitals they exclude are those patients value least. Hoagland (2024) shows that as cardiologists adopt TAVR they perform fewer coronary interventions, a loss concentrated among disadvantaged patients. In both, the return to a program and its spillovers run through the procedure, through experience with it in the first and through competition for operator capacity in the second. A clinical literature documents volume and outcomes, physician incentives, geographic access and hospital margins in TAVR (Vemulapalli et al., 2019; 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, adoption while the regulated price paid TAVR no more than the procedure it replaces, and show that the return runs through the hospital: adoption raises demand among patients who cannot receive TAVR, while learning, which shortens stays but does

not lower the cost of the device, does not account for early entry.

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). The closest paper in method is Schmidt-Dengler (2025), who estimates a dynamic game of MRI adoption among US hospitals and decomposes the effect of competition on adoption timing into business stealing and preemption, of which business stealing is the larger. There business stealing is recovered from adoption timing and market structure. We measure it from patient choice, linking 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), and show that under a regulated price business stealing is not only a cost competition imposes but the return that finances early adoption. The demand response extends to patients with little chance of receiving 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.

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

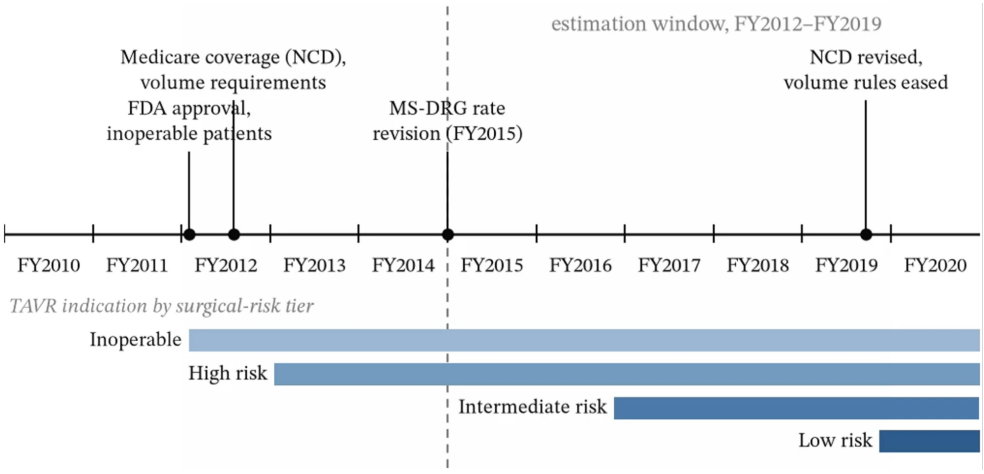



Figure 1: Coverage, coding and payment for TAVR under Medicare, fiscal years 2011–2019. The procedure was billable from the coverage decision under the payment groups for cardiac valve procedures; separate groups for endovascular valve replacement (MS-DRG 266/267) with a higher payment took effect in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

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 the chance of receiving TAVR differed widely across patients at every date: in 2012–2014 most AVR patients had little chance of it, and as coverage widened the chance rose, most for higher-risk patients; Section V uses this variation 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

[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.

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.

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 2.70 million patient-years, of which 14 percent receive a valve replacement in the year, a share that rises from 11 percent in 2012 to 15 percent in 2019. Table 2 describes the population.

Coverage, and hence the chance that a patient receives TAVR, was defined by surgical risk, so we recover the STS predicted operative mortality as it is computed in clinical

[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: The diagnosed population and treated patients

| | Untreated | SAVR | TAVR |
|---|---|---|---|
| Patient-years | 2,325,816 | 208,505 | 161,787 |
| Age, mean | 81.8 | 74.0 | 81.8 |
| Female (percent) | 53.4 | 37.0 | 46.9 |
| *STS predicted mortality (percent)* | | | |
| Mean | 3.24 | 1.72 | 3.31 |
| Median | 2.45 | 1.44 | 2.73 |
| Interquartile range | 1.46–4.26 | 0.97–2.08 | 1.82–4.35 |
| Below 2 (percent) | 41.2 | 72.8 | 30.6 |
| 2–4 (percent) | 30.8 | 22.4 | 39.5 |
| 4–8 (percent) | 22.6 | 4.6 | 26.5 |
| 8 or above (percent) | 5.4 | 0.1 | 3.3 |
| *Conditions in the prior twelve months (percent)* | | | |
| Balloon aortic valvuloplasty | 0.5 | 0.4 | 5.2 |
| Coronary bypass surgery | 15.8 | 6.8 | 20.4 |
| Home oxygen or chronic respiratory failure | 16.7 | 5.1 | 13.2 |
| Peripheral artery disease | 17.3 | 11.0 | 21.6 |
| Aortic atherosclerosis | 4.9 | 7.9 | 17.0 |
| Skilled-nursing or home-health use | 50.8 | 10.7 | 28.1 |
| Dialysis or CKD stage 4–5 | 10.5 | 4.7 | 8.9 |
| Prior valve surgery | 6.9 | 5.2 | 8.4 |
| Any of the eight | 70.1 | 35.9 | 65.4 |
| Two or more | 35.1 | 11.7 | 35.6 |
| *Travel time (minutes)* | | | |
| Nearest hospital, median | 14.2 | 15.5 | 13.8 |
| Nearest TAVR hospital, median | 20.9 | 23.5 | 17.6 |
| Treating hospital, median | | 25.1 | 27.0 |
| Treating hospital, mean | | 46.8 | 51.6 |
| Treated at the nearest hospital (percent) | | 42.7 | 33.1 |
| Treated at a TAVR hospital (percent) | | 75.4 | 100.0 |

*Notes:* Demand estimation sample, fiscal years 2012–2019: fee-for-service beneficiaries aged 65 or older diagnosed with aortic valve disease, and every beneficiary treated with TAVR or SAVR, one observation per patient-year. STS predicted mortality from the age–sex–comorbidity grid. Conditions are read from inpatient, outpatient, DME, home-health and skilled-nursing claims in the 365 days before the AVR admission; for untreated patients the window ends on 30 September of the fiscal year, so it also covers care after the diagnosis. TAVR hospital: a hospital operating a TAVR program in that year; the nearest one is among the patient's 50 nearest hospitals. Travel time is driving time from the residential ZIP code.

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 score alone does not determine who receives TAVR. The heart team assesses operability on conditions the score omits, and the pivotal trials defined inoperability partly by them: a calcified (porcelain) aorta, prior sternotomy, frailty and chest radiation (Leon et al., 2010). We therefore record eight conditions from each patient's claims in the twelve months before the admission, or before the end of the fiscal year for untreated patients: prior balloon aortic valvuloplasty, prior coronary bypass surgery, home oxygen or chronic respiratory failure, peripheral artery disease, aortic atherosclerosis, the claims counterpart of a porcelain aorta, skilled-nursing or home-health use, a marker of frailty, dialysis or advanced chronic kidney disease, and prior valve surgery; the last two were exclusions in the first trials.[4]

Table 3 shows that these conditions carry information the score does not. Among treated patients, year effects and STS tiers explain between 19 and 25 percent of the variation in modality within each policy period; adding the eight conditions raises this to between 29 and 35 percent and leaves the STS coefficients almost unchanged, so the conditions are close to orthogonal to the score. They enter with the signs the clinical criteria imply. Over 2012–2019, prior valvuloplasty and prior bypass surgery each raise the probability of TAVR by 26 percentage points, home oxygen by 13, aortic atherosclerosis by 10, skilled-nursing use by 9 and peripheral artery disease by 8, while dialysis lowers it by 6. Prior valve surgery has no effect before fiscal 2015 and raises the probability by 2 points afterwards, when TAVR inside a failed surgical valve became available. Within the score, the gradient flattens above 4 percent as the indication widens: relative to patients at 4 to 8 percent, those above 8 percent are 28 points more likely to receive TAVR in 2012–2014 and 5 points in 2017–2019. Table 4 shows the same pattern without functional form: among patients treated in 2012–2014 with a score below 2 percent, 3 percent of those with none of the conditions

[4]Codes, ICD-9 / ICD-10: valvuloplasty 35.96 / 027F, CPT 92986; coronary bypass V45.81, 414.02–414.05, 36.1 / Z95.1, I25.7, 021; home oxygen or respiratory failure V46.2, 518.83 / Z99.81, J96.1, and oxygen-equipment HCPCS codes; peripheral artery disease 440.2, 443.8–443.9, 447.1 / I70.2, I73.9, I77.1; aortic atherosclerosis 440.0 / I70.0; skilled-nursing or home-health use, any claim of either type; dialysis or chronic kidney disease stage 4–5 V45.11, V56, 585.4–585.6 / Z99.2, Z49, N18.4–N18.6; prior valve surgery V43.3, 35.1, 35.2 / Z95.2–Z95.4, 02RF–02RJ.

received TAVR, against 29 percent of those with two or more; by 2017–2019 the two shares are 28 and 72 percent.

TABLE 3: Who receives TAVR: risk score and conditions

| | 2012–2014 | | 2015–2016 | | 2017–2019 | | 2012–2019 | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| STS below 2 | −45.9 | −39.4 | −61.8 | −56.7 | −53.7 | −51.1 | −53.4 | −49.2 |
| | (0.4) | (0.4) | (0.4) | (0.4) | (0.2) | (0.2) | (0.2) | (0.2) |
| STS 2–4 | −27.5 | −24.5 | −27.2 | −26.3 | −15.7 | −15.7 | −22.2 | −21.2 |
| | (0.5) | (0.5) | (0.4) | (0.4) | (0.2) | (0.2) | (0.2) | (0.2) |
| STS 8 or above | 28.3 | 24.5 | 12.8 | 11.3 | 4.7 | 4.7 | 11.1 | 10.0 |
| | (1.2) | (1.2) | (0.6) | (0.7) | (0.3) | (0.3) | (0.3) | (0.4) |
| Valvuloplasty | | 38.4 | | 26.3 | | 17.1 | | 26.5 |
| | | (0.7) | | (0.7) | | (0.5) | | (0.4) |
| Coronary bypass | | 22.9 | | 30.6 | | 25.2 | | 26.0 |
| | | (0.4) | | (0.4) | | (0.3) | | (0.2) |
| Home oxygen | | 13.1 | | 13.9 | | 11.7 | | 12.8 |
| | | (0.5) | | (0.5) | | (0.4) | | (0.2) |
| Peripheral artery disease | | 7.6 | | 8.5 | | 7.6 | | 7.9 |
| | | (0.3) | | (0.4) | | (0.3) | | (0.2) |
| Aortic atherosclerosis | | 11.6 | | 10.9 | | 9.0 | | 10.2 |
| | | (0.4) | | (0.4) | | (0.3) | | (0.2) |
| Skilled-nursing use | | 9.8 | | 10.3 | | 7.6 | | 9.0 |
| | | (0.3) | | (0.4) | | (0.3) | | (0.2) |
| Dialysis or CKD 4–5 | | −4.0 | | −6.5 | | −7.6 | | −6.3 |
| | | (0.5) | | (0.6) | | (0.4) | | (0.3) |
| Prior valve surgery | | −0.7 | | 2.2 | | 2.3 | | 1.6 |
| | | (0.5) | | (0.6) | | (0.4) | | (0.3) |
| Year effects | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| $R^2$ | 0.194 | 0.319 | 0.249 | 0.347 | 0.224 | 0.286 | 0.319 | 0.392 |
| Patients | 111,778 | | 95,543 | | 162,971 | | 370,292 | |

*Notes:* Linear probability models of receiving TAVR among patients treated with TAVR or SAVR, in percentage points. The omitted STS tier is 4–8 percent. Heteroskedasticity-robust standard errors in parentheses. Conditions are defined in Table 2.

We therefore describe a patient by the STS score and the eight conditions wherever risk enters: in the demand model of Section V, where both shift the value of a TAVR hospital, and in the patient types of Section VII. The modality rule estimated there gives each patient's probability of receiving TAVR at an adopter, and it is this probability, rather than the STS threshold in force, that we use to identify patients with a negligible chance of the procedure.

TABLE 4: Share of treated patients receiving TAVR, by risk score and number of conditions (percent)

| | 2012–2014 | | | 2015–2016 | | | 2017–2019 | | |
|---|---|---|---|---|---|---|---|---|---|
| STS (percent) | None | One | Two+ | None | One | Two+ | None | One | Two+ |
| Below 2 | 2.9 | 8.8 | 29.4 | 10.5 | 26.3 | 53.4 | 27.8 | 48.4 | 71.5 |
| 2–4 | 13.1 | 26.6 | 52.0 | 41.1 | 57.9 | 75.9 | 70.5 | 79.4 | 87.9 |
| 4–8 | 40.0 | 54.0 | 71.0 | 78.0 | 82.5 | 89.1 | 93.3 | 93.8 | 94.4 |
| 8 or above | 75.3 | 82.3 | 88.5 | 96.6 | 95.7 | 96.4 | $> 98$ | $> 98$ | 98.2 |
| All | 8.5 | 20.8 | 46.7 | 25.8 | 46.9 | 71.1 | 47.6 | 68.0 | 83.9 |

*Notes:* Patients treated with TAVR or SAVR. Columns count how many of the eight conditions of Table 2 the patient carries. Cells with fewer than 11 SAVR patients are reported as bounds.

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 diagnosed population and the treated. TAVR patients are older than SAVR patients ($81.8$ against $74.0$ years), more often women, and riskier, with a median STS score of $2.7$ percent against $1.4$, as the risk-tiered indication implies; two thirds carry at least one of the eight conditions, against a third of SAVR patients. The untreated are as old as TAVR patients and have a similar distribution of the score. 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 (Appendix Table 15), which is the payment delay described in Section II; TAVR patients stay half as long ($4.9$ against $10.0$ days) and die in hospital half as often. Patients travel for aortic valve replacement: the mean trip to the treating hospital is about 50 minutes, against a median of 14 minutes to the nearest hospital, and two thirds of TAVR patients and more than half of SAVR patients are treated elsewhere than at their nearest hospital. Three quarters of SAVR patients are treated at a hospital that offers TAVR. The question for the demand model is how much of that travel adoption redirects.

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.[5] Per-case reimbursement is the payment amount on each claim.

[5]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

Figures 2 and 3 plot, year by year, the distribution of charges and of payments by procedure. 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 13), 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 to about $10 thousand in 2015, and the gap is largest in the years right after it. The densities show that the revision moved the whole TAVR payment distribution, not only its median: from 2015 the TAVR density lies to the right of the SAVR density and is more concentrated, and by 2019 the two overlap. These are distributions across episodes; Section VII builds the hospital-, year- and risk-specific margins the adoption model uses.

measure fully allocated average cost, which can exceed the cost of an additional admission.

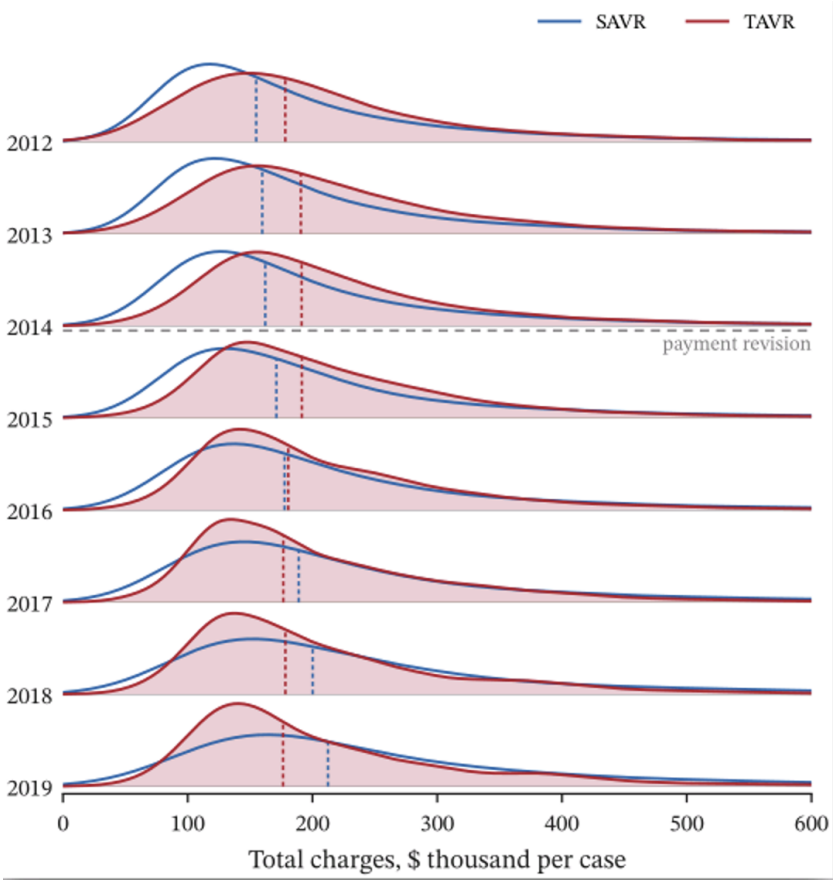



Figure 2: Total charges per case for TAVR and SAVR by fiscal year: kernel densities, one density scale for all years. TAVR shaded, SAVR as a line; dotted ticks mark the medians; the dashed rule marks the fiscal 2015 payment revision. Source: Medicare fee-for-service AVR episodes, FY2012–2019.

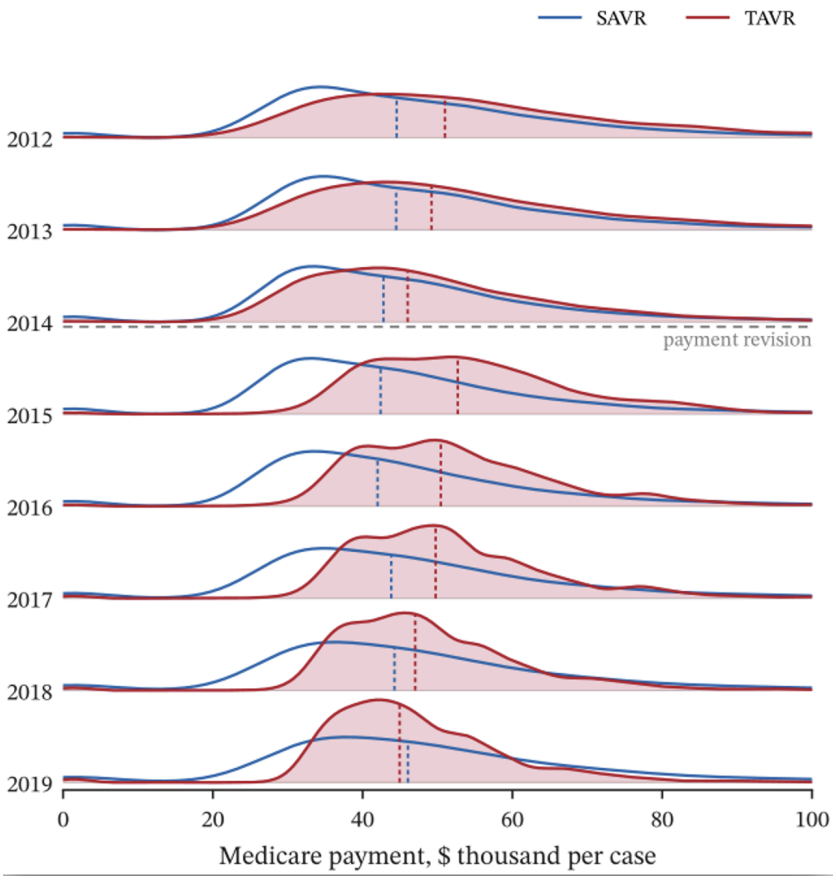



Figure 3: Medicare payment per case for TAVR and SAVR by fiscal year: kernel densities, one density scale for all years. TAVR shaded, SAVR as a line; dotted ticks mark the medians; the dashed rule marks the fiscal 2015 payment revision. Source: Medicare fee-for-service AVR episodes, FY2012–2019.

## 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. The rate changes the procedure rather than the hospital: at the revision, programs already operating shift patients from SAVR to TAVR and their TAVR volume rises by about 8 percent. Learning exists but shortens stays rather than lowering the device cost, and hospitals with no rival in their HRR enter no earlier than others, although hospitals that treat most of their own patient catchment enter earlier. Steering does appear: SAVR patients bypass their nearest hospital for one that offers TAVR, and after the revision a hospital's adoption hazard falls by about 15 percent for every ten percentage points of its patient catchment already served by rival adopters. 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 4) does not follow a smooth S-curve.[6] 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.[7] 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

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

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

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.

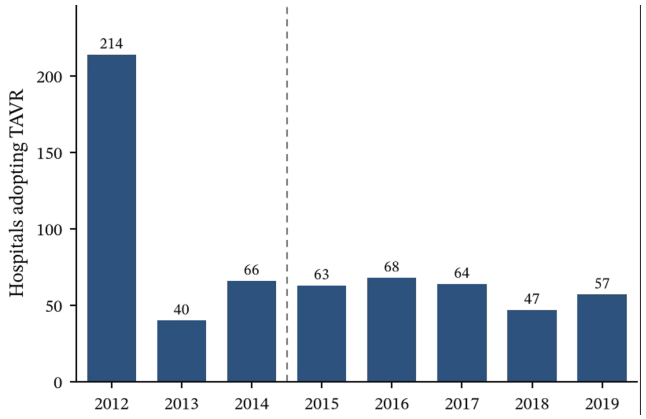



Figure 4: Distribution of hospitals by fiscal year of TAVR adoption.

### 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 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.[8]

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, $\mu_s$, 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

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

hospital that offers TAVR, whether or not they receive it,[9] 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,

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

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 $V_i^L \geq 0$,

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

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

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

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.[10] Figure 5 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 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.

[10] Heterogeneous beliefs about the revised rate add little. A hospital that is optimistic about $\mu_2$ 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 $F_i$.

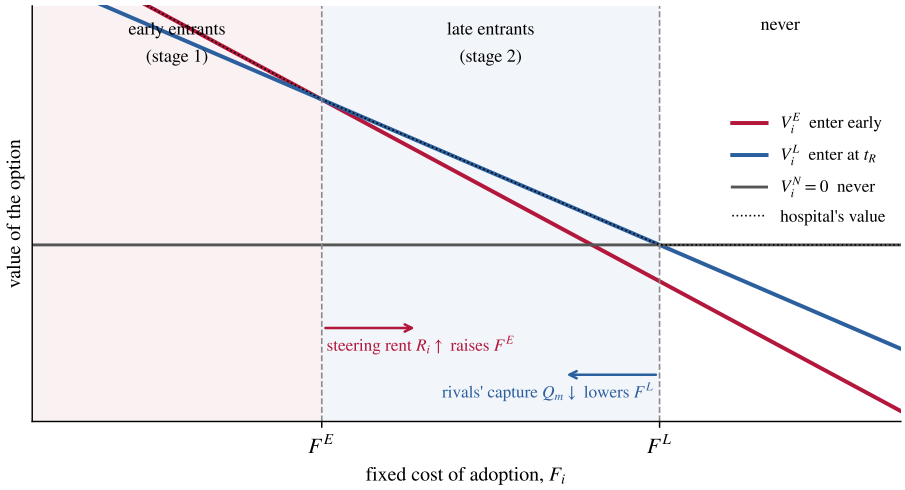



Figure 5: Entry cutoffs in the two-stage model. The lines are the values of entering early, $V_i^E$, entering at the revision, $V_i^L$, and never entering, against the fixed cost $F_i$; the dotted line is the hospital's value, the upper envelope. Hospitals below $F^E$ enter in stage 1, those between $F^E$ 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.

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,[11] and test $\Delta_{2015} \equiv \tau + 3\delta = 0$. From Table 5, $\widehat{\Delta}_{2015} = 0.471 + 3(-0.131) \approx 0.078$, an increase of roughly $8\%$ in TAVR volume at the cutoff.[12] Hospitals already offering TAVR respond to the rate, so the payment enters their decisions.

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

[12] 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.

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.

TABLE 5: TAVR volume at 2012 adopters around the 2015 rate revision

| | $\log$(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) |
| Hospital fixed effects | Yes (190) |
| Observations | 994 |
| $R^2$ | 0.805 |

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

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

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

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 $\delta_h$ absorb time-invariant hospital heterogeneity and DRG-based severity fixed effects $\lambda_s$ 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 6 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; $100 \cdot (e^{-0.03} - 1) \approx -2.95\%$).

TABLE 6: Learning: length of stay and charges by years since adoption

| | Length of stay (days) (1) | Log(total charges) (2) |
|---|---|---|
| Years since adoption | $-0.996^{***}$ | $-0.030^{***}$ |
| | (0.074) | (0.005) |
| Hospital and severity-category effects | Yes | Yes |
| Observations | 40,025 | 40,025 |
| $R^2$ | 0.346 | 0.683 |

*Notes:* Patient-level controls: Charlson Comorbidity Index and indicators for hypertension, congestive heart failure, chronic pulmonary disease, diabetes, renal disease and prior cardiac surgery. Standard errors in parentheses, clustered at the hospital level. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Learning exists, but it operates on the intensive margin: the prosthetic valve is a device cost that experience cannot erode, which caps $\eta_i$.

A second check uses market structure. If learning carried early entry, a hospital that keeps more of its local volume would enter earlier than others, since it keeps more of the volume over which experience pays. Using hospitals with AVR claims in fiscal 2012 and no earlier program, we estimate the linear probability model

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

for two measures of local market power $M_h$. The first, $\text{Monopoly}_h$, equals one if $h$ is the only hospital in its HRR, in HRRs with fewer than ten hospitals. Conditional on size and ownership type, monopoly status is not associated with earlier adoption (Table 7, Column 1): the estimate is $-2.1$ percentage points (s.e. 4.7), and the 95 percent interval excludes a monopoly advantage larger than 7 points. HRR boundaries are crossed by patients (Section III), so the second measure does not rely on them. It is the share of the patients of the hospital's own catchment that the hospital itself treated, $\sum_z w_{hz} f_{hz}$, where $w_{hz}$ is the share of $h$'s fiscal 2009–2011 AVR patients who live in ZIP code $z$ and $f_{hz}$ the share of $z$'s AVR patients that $h$ treated. On this measure local dominance predicts early entry (Column 2): ten percentage points more of its own catchment raise the probability of entering before the revision by 2.7 points (s.e. 0.9), against a mean of 34 percent. This is the sign learning predicts. It is also the sign of the defensive side of steering: a hospital that holds most of its catchment has the most patients to lose to a rival program, and entering first protects them. The market-structure test therefore does not separate the two channels;

the demand model of Section V measures the patient channel directly.

TABLE 7: Local market power and adoption before the rate revision

| | (1) | (2) |
|---|---|---|
| Monopoly (only hospital in HRR) | $-0.021$ | |
| | (0.047) | |
| Own share of catchment | | $0.272^{***}$ |
| | | (0.091) |
| Beds (hundreds) | $0.095^{***}$ | $0.087^{***}$ |
| | (0.007) | (0.006) |
| 95% interval, market-power term | $[-0.113, 0.070]$ | [0.094, 0.451] |
| Ownership-type FE | Yes | Yes |
| Hospitals | 766 | 935 |
| Mean of outcome | 0.351 | 0.341 |

*Notes:* Linear probability model of adoption in fiscal 2012–2014; hospitals with AVR claims in fiscal 2012 and no earlier programme. (1): HRRs with fewer than ten hospitals. (2): all HRRs; own share of catchment is the share of the fiscal 2009–2011 AVR patients of the hospital's catchment ZIP codes that it treated itself, weighted by its own patients. Standard errors clustered by HRR. $^{*}$ $p < 0.10$; $^{**}$ $p < 0.05$; $^{***}$ $p < 0.01$.

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 6 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. Among all treated patients whose nearest hospital had no program but who were treated at an adopter, 70 percent received SAVR in 2012–2014, about 2.4 SAVR cases for every TAVR case, and 50 percent over 2012–2019. 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.

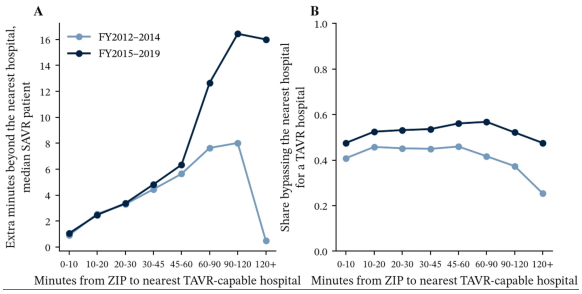



Figure 6: SAVR patients' travel toward TAVR hospitals, by driving time from the patient's ZIP code to the nearest TAVR-capable hospital, fiscal 2012–2014 and 2015–2019. Panel A: median extra minutes beyond the nearest hospital. Panel B: share bypassing the nearest hospital for a TAVR hospital.

### 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 on the hospitals at risk after the revision. For each hospital $h$ that has not adopted by the end of fiscal year $t-1$, $t = 2015, \ldots, 2019$, we model adoption in $t$ as a discrete-time hazard,

$$\Pr\big(\tau_h = t \mid \tau_h \geq t\big) \;=\; G\big(\beta\, R_{h,t-1} + \mathbf{X}_h'\boldsymbol{\gamma} + \lambda\big), \tag{5}$$

with $G$ complementary log-log or linear, and $\lambda$ year or HRR-by-year effects. $R_{h,t-1}$ measures rival adoption at the end of $t-1$ in two ways. The HRR competitor rate is the share of the other hospitals in $h$'s HRR that had adopted. It takes the same value for every hospital at risk in an HRR-year, so it is identified only from variation across HRRs or within an HRR over time, and it cancels from any model stratified by HRR and year. Catchment exposure is the share of $h$'s own pre-TAVR patient catchment already served by rival adopters, $\sum_z w_{hz} \sum_{k \neq h,\, \tau_k \leq t-1} f_{kz}$, with $w_{hz}$ and $f_{kz}$ the fiscal 2009–2011 patient shares defined above. It differs across hospitals of the same HRR and year, so it is identified with HRR-by-year effects, which absorb any common local shock to the value of entry. $\mathbf{X}_h$ collects log beds, teaching status and ownership type; standard errors are clustered by HRR. The conventional confound, that rival entry proxies a common positive local shock, would bias $\beta$ upward.

Table 8 reports the estimates. Across HRRs, with year effects only (Column 1), the

competitor rate is negative and imprecise, $-0.13$ (s.e. 0.23), consistent with common positive local shocks offsetting exhaustion. Within HRRs (Column 2) it is $-0.40$ (s.e. 0.04): ten percentage points more rival adoption lower the annual hazard by 4 points, against an at-risk adoption rate of 11.6 percent. The catchment measure compares hospitals of the same HRR in the same year: ten percentage points more exposure lower the hazard by 1.7 points (s.e. 0.8), about 15 percent (Column 3), and by 19 percent in the complementary log-log model (Column 4). Earlier entry by the rivals that serve a hospital's own patients lowers the value of entering for those that remain.

Two qualifications bound this reading. A hospital still at risk despite heavy rival entry around it is more likely to have a high setup cost, so dynamic selection alone biases $\beta$ downward; the hospital controls absorb part of this selection, not all of it. And a negative $\beta$ does not single out the exhaustion of patients from other forms of rival crowding, such as consolidated referrals or competition for cardiac staff; Section V measures the patient channel directly.

TABLE 8: Adoption hazard after the rate revision and prior rival adoption

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| HRR competitor rate$_{t-1}$ | $-0.126$ | $-0.402^{***}$ | | |
| | (0.230) | (0.044) | | |
| Catchment exposure$_{t-1}$ | | | $-0.170^{**}$ | $-2.129^{***}$ |
| | | | (0.082) | (0.412) |
| log beds | $0.866^{***}$ | $0.100^{***}$ | $0.095^{***}$ | $0.984^{***}$ |
| | (0.168) | (0.022) | (0.021) | (0.169) |
| Teaching (GME) | $0.313^{**}$ | $0.036^{*}$ | $0.035^{*}$ | $0.353^{***}$ |
| | (0.129) | (0.019) | (0.019) | (0.127) |
| Estimator | cloglog | LPM | LPM | cloglog |
| Fixed effects | Year | HRR, year | HRR $\times$ year | Year |
| Hospital-years at risk | 2,292 | 2,289 | 1,953 | 2,292 |
| Adoptions | 265 | 265 | 265 | 265 |
| Hospitals | 568 | 568 | 568 | 568 |
| HRRs | 204 | 204 | 204 | 204 |

*Notes:* Discrete-time hazard of TAVR adoption, fiscal 2015–2019; a hospital is at risk in year $t$ if it has not adopted by the end of $t-1$. HRR competitor rate: share of the other hospitals in the hospital's HRR that adopted by the end of $t-1$. Catchment exposure: share of the hospital's fiscal 2009–2011 AVR patient catchment, by ZIP code, served by hospitals that adopted by the end of $t-1$; it varies across hospitals within an HRR-year. Columns (1) and (4) are complementary log-log models with year effects; (2) and (3) linear probability models. All columns control for ownership type. Column (3) drops HRR-years with a single hospital at risk. Standard errors clustered by HRR. $^{*}$ $p<0.10$; $^{**}$ $p<0.05$; $^{***}$ $p<0.01$.

Larger and teaching hospitals adopt sooner in every column of Table 8, consistent with heterogeneity in $F_i$. 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.

Learning cannot fully account for early entry: a hospital with no rival in its HRR, which keeps all of its local volume, enters no earlier than others, and the learning gain does not depend on rivals, whereas entry by rivals into a hospital's own catchment lowers its hazard. Dominance of its own catchment predicts early entry, but that sign fits steering as well as learning. The reduced forms therefore point to competition for patients as the return that finances early adoption, which Section V measures directly.

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, seen in patient sorting and in the negative effect of cumulative rival entry on later hazards, is consistent with 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 across the risk distribution, more to those likely to receive the procedure but also to those with little chance of receiving it, and that this value rose as the indication widened; 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 and clinical conditions. 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 patient's chance of receiving TAVR, which varies with risk, with the conditions and over time as coverage expanded by risk tier, is what lets us separate a preference for the procedure from a preference for the hospital that offers it. 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 $\chi_j$ 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 for each condition, 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.[13]

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

[13] Appendix Figure 11 reports an event study of hospital AVR admissions around adoption (Sun and Abraham, 2021). Admissions grow slowly before adoption, by about 0.05 log points a year, as Medicare's coverage decision implies, since a hospital could open a program only after reaching minimum annual volumes of valve replacements and catheter procedures. They then jump by 0.38 log points in the adoption year, about seven years of the prior trend, and the risk mix of admissions, flat before adoption, shifts upward at the same date. The change is a discontinuity at adoption, not the continuation of a trend. With the tenure terms in the utility, the value of a program at adoption is identified from this discontinuity, and a pre-existing trend of this size could bias only the tenure coefficients, which are small.

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.[14] 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

$$\begin{aligned} U_{ijt} = {} & \delta_j + \gamma_t + \beta_{1,i}\,\mathrm{tt}_{ij} + \beta_2\,\mathrm{tt}_{ij}^2 + \beta_c\,\mathbf{1}\{j = c(i)\} \\ & + \kappa_t\, r_i + \textstyle\sum_k \kappa_{kt}\, s_{ik} + r_i\left(\omega_1\,\mathrm{size}_j + \omega_2\,\mathrm{teach}_j\right) \\ & + a_{jt}\left(\alpha_{0t} + \alpha_{1t}\, r_i + \textstyle\sum_k \alpha_{kt}\, s_{ik} + (\psi_0 + \psi_1 r_i)\, E_{jt}\right) + \sigma_0\,\zeta_i + \varepsilon_{ijt}, \qquad U_{i0t} = \varepsilon_{i0t}, \end{aligned} \tag{7}$$

where $\mathrm{tt}_{ij}$ is the driving time in minutes from $i$'s ZIP code to hospital $j$, $c(i)$ is the hospital nearest to $i$, $s_{ik} \in \{0, 1\}$ indicates condition $k$ of the eight defined in Section III, $E_{jt}$ is the number of years since $j$ opened its program, capped at seven, $\mathrm{size}_j$ and $\mathrm{teach}_j$ are the hospital's standardized log beds and teaching status, and $\varepsilon_{ijt}$ is i.i.d. type-I extreme value. The outside option is normalized to zero. The fixed effect $\delta_j$ absorbs hospital attributes common to all patients and years; the year effect $\gamma_t$ shifts the value of treatment at any hospital relative to no treatment, with $\gamma_{2012}$ 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 $\eta_i$ and $\zeta_i$ 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 plain logit would imply. Medicare regulates prices and standardizes cost sharing, and no price term enters.

Risk and the conditions enter twice. The terms $\kappa_t r_i$ and $\kappa_{kt} s_{ik}$ shift the value of being treated anywhere: sicker patients are, in the data, less likely to undergo a valve replacement at all. The TAVR term is the additional value of a hospital that offers TAVR. It varies with the patient's risk and conditions and with the year, because who can receive TAVR depends on all three. The indication expanded by risk tier over the sample (Section II), and within a

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

tier the heart team weighs frailty, anatomy and prior surgery, which the STS score omits; Section III shows that the eight conditions shift the probability of TAVR as much as the score does (Tables 3 and 4). The TAVR coefficients are common to all patients who consider the hospital, including those who go on to receive SAVR. A program's value may also grow as it matures: the patients a program draws build its referrals, outcomes and reputation, so a hospital that steers patients early keeps an advantage over rivals that enter later. The tenure term $(\psi_0 + \psi_1 r_i)E_{jt}$ allows this. The interactions of risk with hospital size and teaching status let sicker patients sort toward large and teaching hospitals whether or not they offer TAVR, so that this sorting is not attributed to the program.

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 6). Each patient's set consists of the forty hospitals nearest the patient's residential ZIP code by driving time and the outside option; treated patients whose chosen hospital lies outside that set are excluded. The rule excludes 4.8 percent of treated patients over 2012–2019, between 4.4 and 5.7 percent in each year, and affects TAVR and SAVR patients alike (4.7 and 4.9 percent). The median treated patient travels about 25 minutes, far inside the reach of the fortieth-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: $\delta_j$ 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 patient's chance of receiving TAVR does: a positive value of an adopter to a

patient who will almost surely receive SAVR is a preference for the hospital. That chance is set by a combination of factors, not by the risk threshold alone: coverage expanded by risk tier, and within a tier the heart team weighs the conditions, which shift the probability of TAVR as much as the score does (Tables 3 and 4). We therefore measure it in the data, as the share of treated patients of the same year, risk tier and number of conditions who receive TAVR, and Section VI reads the value of an adopter against it; 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 ($\beta_1$, $\beta_2$, $\sigma$), the preference for the nearest hospital ($\beta_c$), the year effects, the risk and condition gradients of treatment ($\kappa_t$, $\kappa_{kt}$), the risk interactions with hospital size and teaching ($\omega_1$, $\omega_2$), the dispersion of the preference for treatment ($\sigma_0$), and the technology terms ($\alpha_{0t}$, $\alpha_{1t}$, $\alpha_{kt}$, $\psi_0$, $\psi_1$). 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 $\gamma_t$ given the hospital effects, and the way that share falls with the risk score pins down $\kappa_t$. 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 $\sigma_0$ is identified from how the treated share responds when a patient's choice set changes: when a hospital in the set adopts, a large $\sigma_0$ 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 $\delta_j$ 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; $\alpha_{1t}$ and $\alpha_{kt}$ are identified from that difference, net of $\kappa_t$ and $\kappa_{kt}$, and the tenure terms from differences across programs

of different ages within a year. 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 $\alpha_{0t}$. 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 $\alpha_{1t}$ 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, 2.70 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: $\delta$ 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 $\delta$, 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 9). 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. The value of a TAVR program grows with the patient's chance of receiving the procedure, and it extends to patients with a low chance, who are almost always treated with SAVR, so a program pulls patients across the risk distribution. And the value a program adds rose over the sample, from 15 minutes of travel for the average patient receiving a valve replacement in 2012 to 86 minutes in 2019.

TABLE 9: Patient demand for hospitals: estimated parameters

| *A. Travel and the preference for treatment* | | |
|---|---|---|
| Travel time (minutes), $\beta_1$ | | $-0.0223$ |
| Travel time squared ($\times 10^6$), $\beta_2$ | | 3.94 |
| Nearest hospital, $\beta_c$ | | 1.307 |
| Std. dev. of the travel coefficient, $\sigma$ | | 0.0051 |
| Std. dev. of the preference for treatment, $\sigma_0$ | | 0.0000 |
| *B. Value of a TAVR hospital* | | |
| Fiscal year | Level, $\alpha_{0t}$ | $\times$ risk, $\alpha_{1t}$ |
| 2012 | $-0.103$ | 0.100 |
| 2013 | $-0.336$ | 0.229 |
| 2014 | $-0.413$ | 0.269 |
| 2015 | $-0.382$ | 0.300 |
| 2016 | $-0.422$ | 0.382 |
| 2017 | $-0.329$ | 0.455 |
| 2018 | $-0.541$ | 0.597 |
| 2019 | $-0.299$ | 0.609 |
| Tenure, $\psi_0$ | | 0.0293 |
| Tenure $\times$ risk, $\psi_1$ | | 0.0066 |
| Log likelihood | | $-1{,}674{,}377.8$ |
| Patient-years | | 2,696,108 |
| Hospitals | | 933 |

*Notes:* Mixed logit with hospital fixed effects, concentrated out, and an outside option normalized to zero. Panel B: the terms of a hospital's TAVR program, $a_{jt}(\alpha_{0t}+\alpha_{1t}r_i+(\psi_0+\psi_1 r_i)E_{jt})$, with $r_i$ the STS predicted mortality in percentage points and $E_{jt}$ years since the program opened, capped at seven. The program terms for the eight conditions are in Table 10; the year effects, the risk and condition terms of the value of treatment, and the risk interactions with hospital size and teaching status are in Appendix Table 14. Standard errors to be added.

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 20 percent.[15] The quadratic term is positive, so the disutility of travel flattens at long distances: it offsets about 2 percent of the linear term at 100 minutes and 4 percent at 200. The option of the closest hospital is worth about 59 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 standard logit gives $-0.121$ per minute (Raval and Rosenbaum, 2021) and the distance

[15]The travel coefficient is $-0.0223$ per minute, so ten minutes change the log-odds by $-0.223$, $e^{-0.223} \approx 0.80$. With price fixed by Medicare, travel time is the price patients pay, and all other hospital attributes are measured against it in minutes.

TABLE 10: Value of a TAVR hospital by patient condition

| | 2012 | 2013 | 2014 | 2015 | 2016 | 2017 | 2018 | 2019 |
|---|---|---|---|---|---|---|---|---|
| Balloon valvuloplasty | 2.095 | 2.271 | 3.288 | 2.585 | 2.595 | 3.065 | 1.989 | 1.994 |
| Coronary bypass | 0.409 | 0.664 | 0.708 | 0.844 | 1.076 | 1.310 | 1.362 | 1.390 |
| Home oxygen or respiratory failure | 0.289 | 0.345 | 0.354 | 0.408 | 0.504 | 0.416 | 0.409 | 0.594 |
| Peripheral artery disease | 0.153 | 0.251 | 0.294 | 0.272 | 0.262 | 0.270 | 0.461 | 0.685 |
| Aortic atherosclerosis | 0.393 | 0.535 | 0.529 | 0.638 | 0.750 | 0.679 | 0.846 | 0.520 |
| Skilled-nursing or home-health use | 0.137 | 0.054 | 0.181 | 0.204 | 0.181 | 0.229 | 0.133 | 0.055 |
| Dialysis or CKD stage 4–5 | −0.062 | −0.049 | −0.056 | −0.148 | 0.038 | −0.280 | −0.183 | −0.495 |
| Prior valve surgery | 0.343 | 0.452 | 0.533 | 0.644 | 0.715 | 0.743 | 0.871 | 0.728 |

*Notes:* Coefficients $\alpha_{kt}$ on the indicator of condition $k$ interacted with the hospital's TAVR program, by fiscal year, from the demand model of Table 9. A coefficient of 0.0223 is worth one minute of travel. Conditions are recorded from the patient's claims in the twelve months before the admission (Section III). Standard errors to be added.

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, 1.31, 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 0.0051, about a quarter of the mean. Under the normal distribution fewer than 0.001 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 including the cash; welfare net of the cash can fall (Section VIII.C). The dispersion of the preference for treatment is estimated at zero.[16] Because this term is common to every hospital in a patient's set, it would govern 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 $\kappa_t$ is negative throughout (Appendix Table 14), $-0.32$ per point of the risk score in 2012 and between $-0.40$ and $-0.70$ after: holding

[16]The inside-option baseline already varies by year with the patient's risk score and with each of the eight conditions, and these observed characteristics account for the heterogeneity in the propensity to be treated; no unobserved dispersion remains. We keep the term so that the specification nests the random-coefficient model; at $\sigma_0 = 0$ the inside–outside margin is a logit in the observed characteristics.

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.

Figure 7 expresses the TAVR terms in minutes of travel. Panel A draws the value of an adopter against the STS score in three years, for a patient with none of the eight conditions and for a patient with one, at the average tenure of programs in that year. In every year the value rises with risk. A single condition raises the value at every level of risk: the line for patients with one condition is positive everywhere, and the line for patients with none crosses zero near a risk score of 1 percent. Table 10 reports the condition terms: every condition other than dialysis raises the value of an adopter in every year, by as much as an hour of travel for prior bypass surgery in 2019.

The value of an adopter combines a preference for the procedure and a preference for the hospital, and the patient's chance of receiving TAVR separates them. Panel B lists every year, risk tier and number of conditions in which fewer than 10 percent of treated patients received TAVR, with the value of an adopter to those patients. Patients with a 5 to 8 percent chance value an adopter at 9 to 14 minutes; in 2012, patients at 2 to 4 percent risk received TAVR in 5 percent of cases and value an adopter at 9 minutes. The healthiest patients, below 2 percent risk, whose chance of TAVR was 1 to 8 percent, are close to indifferent, at 0 to 4 minutes: for them proximity dominates the choice of hospital, and the nearest-hospital premium alone is worth 59 minutes of travel. Computed the same way for the groups with about a 20 percent chance, four in five of whom receive SAVR, the value of an adopter is 17 to 23 minutes of travel.

A condition that shifts the value of an adopter without shifting the chance of the procedure gives a check on this reading (Appendix Figure 12). Prior bypass surgery, home oxygen, peripheral artery disease and aortic atherosclerosis raise both the probability of TAVR and the value of an adopter, and dialysis lowers both. A prior valve operation changes the probability of TAVR by between $-2.4$ and $2.8$ percentage points in every year, yet raises the value of an adopter by 15 to 39 minutes of travel, a preference for the hospital that does not run through the procedure.

Over time, as coverage widened by risk tier, the value of an adopter to the average patient receiving a valve replacement, TAVR or SAVR, rose from 15 minutes of travel in 2012 to 37 in 2015 and 86 in 2019, and its slope in risk steepened, from $\alpha_{1t} = 0.10$ per point in 2012 to $0.61$ in 2019. Within a year, older programs are worth more: each year of operation adds 1.3 minutes of travel for a patient at zero risk and 0.3 more per point of risk, about 2 minutes a year at 2 percent risk and 13 minutes after seven years. An early entrant therefore keeps the advantage its draw builds, the demand counterpart of the captured pool of Section IV.B, and the reason a rival's early entry lowers a late entrant's return.

The late adopters therefore entered a market in which the technology mattered to a larger share of patients than it had for the 2012 cohort, the demand-side counterpart of the diffusion documented in Section IV.A. 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 370,000 patient-years in 2019 of which 56,000 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),

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

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 $\text{tt}_{ij}$ 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

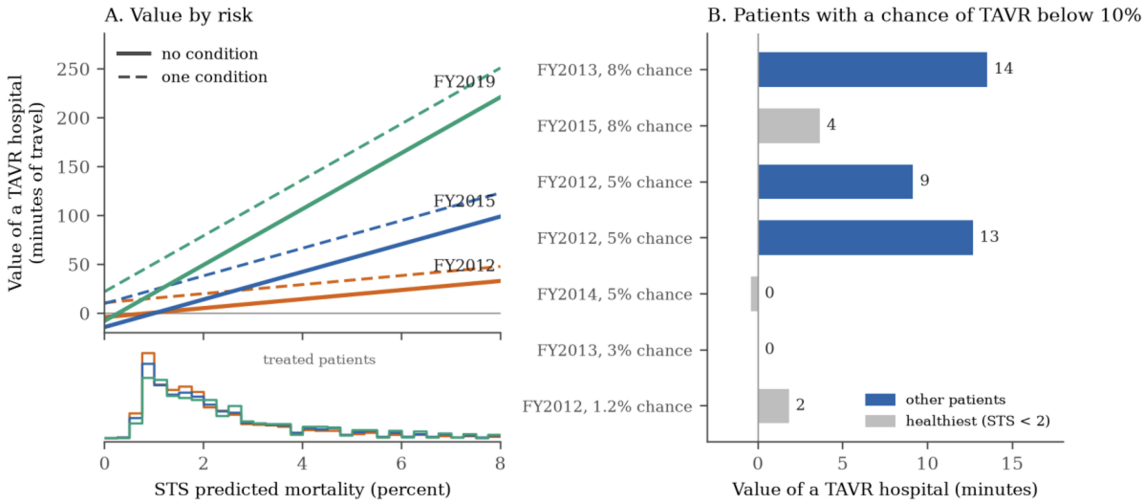



Figure 7: The value of a TAVR hospital. Panel A: value of an adopter relative to a non-adopter, in minutes of travel, against the STS risk score, for a patient with none of the eight conditions (solid) and with one condition (dashed, average of the conditions other than dialysis, weighted by their prevalence among patients treated at an adopter), at the average tenure of programs in each year; the lower strip shows the risk distribution of treated patients. Panel B: patient groups in which fewer than 10 percent of treated patients received TAVR, by fiscal year; the bar is the value of an adopter to the group, and the label gives the year and the group's chance of TAVR, the share of its treated patients who received it. Groups are defined by STS tier and the number of the eight conditions; grey bars are the healthiest patients, below 2 percent risk with no condition.

and location of adoption rather than hospital quality. The technology terms $\alpha_{0t}, \alpha_{1t}$ 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 \$4.9M for adopters against \$8.9M 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 $\theta$ denote a patient's type: the predicted operative mortality $r$ and the conditions $\mathbf{s}$ of Section III. The market in year $t$, the diagnosed beneficiaries of Section V, has size $M_t$ and type distribution $F_t(\theta)$. The demand system delivers hospital choice probabilities by type,

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

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,

$$p_{h,t}(\theta) = \Lambda\Big(\phi_h + \beta_{\theta,t}\, r + \mathbf{s}'\boldsymbol{\rho}_{q(t)} + g\,(t - 2012) + d\,\mathbf{1}\{t \geq 2015\} + \gamma\,\Delta_{h,t}(\theta)\Big), \qquad (9)$$

where $\Lambda$ is the logistic function, $\phi_h$ a hospital effect, $\beta_{\theta,t}$ a risk slope that differs by year, $\boldsymbol{\rho}_{q(t)}$ the effects of the eight conditions, with dialysis and prior valve surgery combined, which differ by period $q(t)$, and $\Delta_{h,t}(\theta)$ the per-case margin gap between the two procedures defined below. The common time term is a linear trend $g$ and a step $d$ from fiscal 2015 rather than a set of year effects, for the reason given in Section VII.C. 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),

$$\text{cost}^m_{h,t}(\theta) \;=\; c_{h,t} \cdot ch^m_{h,t}(\theta), \qquad m \in \{T, S\}, \tag{10}$$

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.[17] 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 type distribution of the patients hospital $h$ admits under state $a_h$, the market distribution reweighted by $M_t\, s_{h,t}(\theta; a_h, A_{-h,t})$,

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

and let the per-admission profit be

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

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

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

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

$$\begin{aligned}\Delta\pi_{h,t}(A_{-h,t}) &\equiv \pi^{op}_{h,t}(1) - \pi^{op}_{h,t}(0) \\ &= \underbrace{\big[Y_{h,t}(1) - Y_{h,t}(0)\big]\;\mu^{0}_{h,t}}_{\text{admissions at pre-adoption mix}} \;+\; \underbrace{Y_{h,t}(1)\,\big[\mu^{1}_{h,t} - \mu^{0}_{h,t}\big]}_{\text{composition shift at adopted mix}}\,. \end{aligned} \tag{11}$$

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 $\Delta_{h,t}(\theta)$.

### VII.B. The timing problem

Each hospital chooses $\tau_h \in \mathcal{T} = \{t_0, t_0{+}1, \ldots, T, \infty\}$, where $\infty$ denotes never, $t_0 = 2012$ and $T = 2019$. The present value at $\tau$ of adopting at $\tau$, 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 $\tau$, specified below, and let $\varepsilon_{h\tau}$ be i.i.d. type-I extreme value with scale $\sigma_\varepsilon$. 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,

$$P_h(\tau) = \frac{\exp v_{h\tau}}{1 + \sum_{\tilde{\tau} \neq \infty} \exp v_{h\tilde{\tau}}}. \tag{12}$$

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 $\tau$. We estimate the scale $\sigma_\varepsilon$; 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

$$K_{h\tau} = \kappa_0 + Z_h\beta_Z + X_h\beta_X + \psi_{r(h)} + \lambda_\tau + \nu_h, \qquad \nu_h \sim N(0, \sigma_\nu^2), \tag{13}$$

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 $\lambda_\tau$ 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 $\nu_h$ is integrated out by Gauss–Hermite quadrature over the choice probabilities (12), 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 (13) 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 $\nu_h$ 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 11 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 common time term in (9) is a linear trend with a step at the 2015 revision rather than a set of year effects. The revision changed the margin gap at every hospital at once, so year effects would absorb it and leave $\gamma$ identified only from the small differences in the gap across hospitals and patients within a year; the trend and step leave the year-to-year movement of the gap around that path to identify it as well. The margin coefficient is $\hat{\gamma} = 0.51$ (s.e. 0.15): a larger TAVR margin over SAVR raises the probability that a program assigns an admitted patient to TAVR. The trend is 0.48 a year and the step 0.17 (s.e. 0.03). The counterfactuals move the step with the payment change, so a rate change shifts the TAVR share at programs already operating through both $\gamma$ and $d$; the step also absorbs other changes in fiscal 2015, such as the valve generations approved at its end, so this attribution is generous to the payment. The median adopter's net adoption cost is \$4.9M, with an interquartile range of \$2.8M to \$7.4M. The scale $\hat{\sigma}_\varepsilon$ is \$2.45M and the standard deviation of the permanent hospital effect $\hat{\sigma}_\nu$ is \$0.005M: 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 8). The median hospital that never adopts faces \$8.9M against \$4.9M for one that does. Net adoption cost is negative for

TABLE 11: Supply model estimates

| Stage 1: modality | | Stage 2: adoption timing | |
|---|---|---|---|
| Parameter | Estimate | Parameter | Estimate |
| $\gamma$ (margin gap) | 0.509 | $\sigma_\varepsilon$ | 24.495 |
| | (0.147) | $\sigma_\nu$ | 0.050 |
| $\beta_\theta$ (risk) | 1.046 | $\kappa_0$ | 58.986 |
| | (0.027) | size (log beds, standardized) | -30.569 |
| risk $\times$ 2013 | 0.294 | teaching, $X_1$ | -8.478 |
| | (0.036) | $\lambda$ 2013 | 19.483 |
| risk $\times$ 2014 | 0.422 | $\lambda$ 2014 | 20.714 |
| | (0.036) | $\lambda$ 2015 | 21.945 |
| risk $\times$ 2015 | 0.641 | $\lambda$ 2016 | 23.176 |
| | (0.035) | $\lambda$ 2017 | 24.407 |
| risk $\times$ 2016 | 0.728 | $\lambda$ 2018 | 25.638 |
| | (0.034) | $\lambda$ 2019 | 26.869 |
| risk $\times$ 2017 | 0.883 | hospitals | 945 |
| | (0.035) | adopters 2012–2019 | 521 |
| risk $\times$ 2018 | 1.122 | log likelihood | -1,468.1 |
| | (0.035) | | |
| risk $\times$ 2019 | 1.108 | | |
| | (0.036) | | |
| trend, per year | 0.483 | | |
| | (0.005) | | |
| step from 2015 | 0.174 | | |
| | (0.029) | | |
| claims | 315,918 | | |
| hospitals | 613 | | |
| log likelihood | -143,438.4 | | |

*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. The common time term in Stage 1 is a linear trend in years since 2012 plus a step from fiscal 2015. $\sigma_\varepsilon$, $\sigma_\nu$, $\kappa_0$ and the size, teaching and $\lambda$ 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.

9 percent of adopters, 46 of 521, and one of the 327 never-adopters.[18]

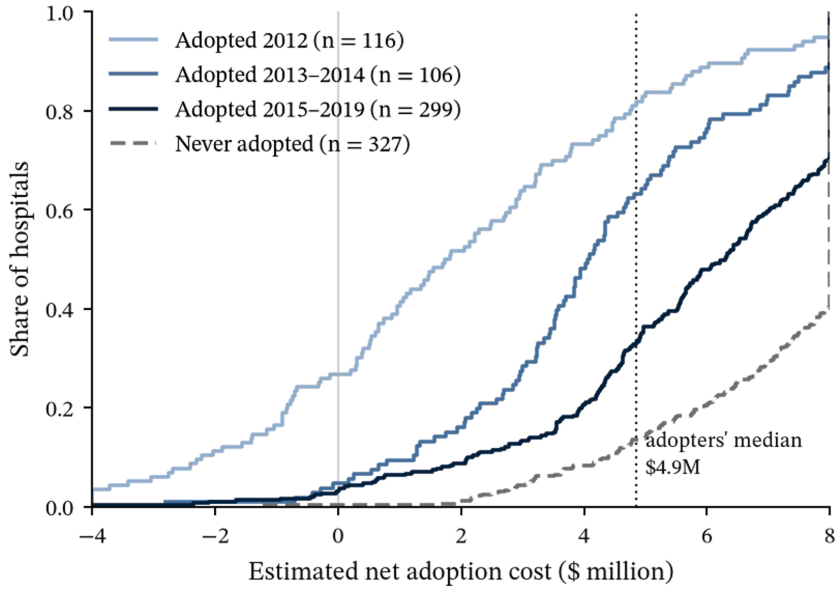



Figure 8: Estimated net adoption cost by adoption cohort. Each curve is the cumulative distribution of $K_{h\tau}$ across hospitals, in millions of dollars, read at the adopter's own adoption year or, for hospitals that never adopt, at the average across candidate years. The dotted line marks the adopters' median, \$4.9M. 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 \$3.1M, and teaching hospitals by a further \$0.8M. The adoption-year effects are positive, \$1.9M in 2013 relative to 2012 and \$2.7M 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.

Figure 9 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 path after 2012, over-predicting entry in 2013, under-predicting it in 2014 to 2017 and over-predicting it in 2018 and 2019; the size panel shows that the late over-prediction comes from large hospitals.

[18] $K_{h\tau}$ 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.

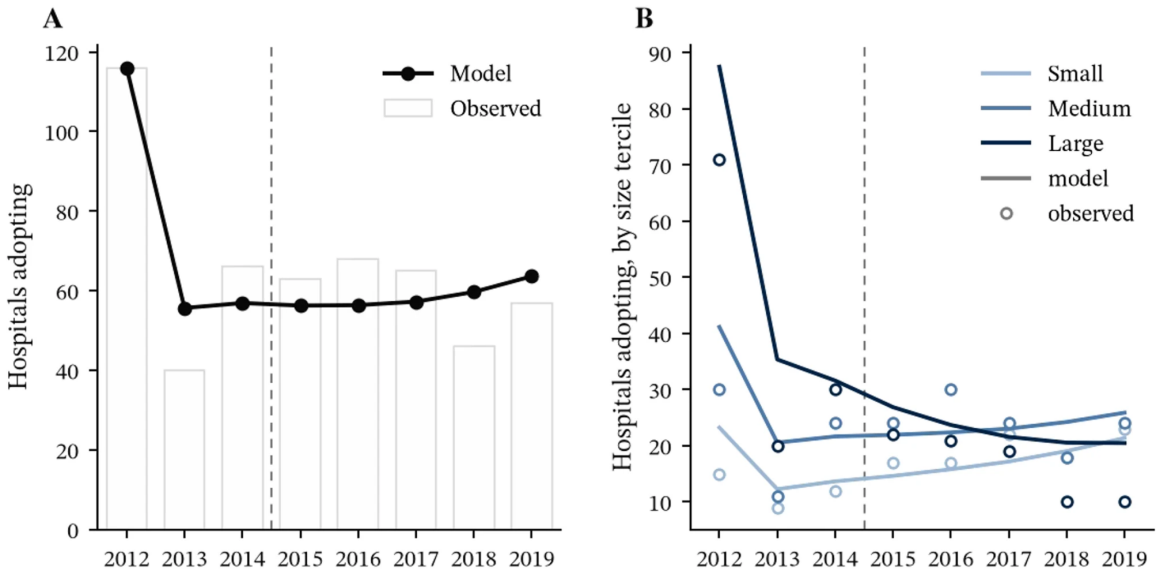



Figure 9: Model fit. Panel A shows hospitals adopting TAVR in each fiscal year, observed (bars) and implied by the estimated timing probabilities (line), for the 848 hospitals that choose when to enter, of which 521 adopt by 2019; the 97 trial sites in before the window are excluded. Panel B shows the 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, and it turns on targeting rather than timing: at the observed outlay of $467M the frontier reaches 634 hospitals against the observed 619, or 24.4 thousand patients a year against 18.1 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. Only with twice the observed outlay does the patient frontier advance the rise to 2012; the coverage frontier keeps it at or after 2015 at every budget. 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 $\nu_i$ 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 $\tau$ 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 \$0.005M (Table 11), 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 \$2.45M, 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 10 summarizes the policy sweep and Table 13 reports the best policy at four budgets. The observed policy, the payment rise of 2015, costs \$467 million in present value and leaves 619 hospitals offering TAVR by 2019 and 18.1 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 \$9.7 million for targeted hospitals, and the voucher at full strength; or 24.4 thousand patients a year, 6.3 thousand or 35 percent more, with the rise moved to 2017, a lump sum of \$0, and the voucher.

The frontier delays the rate rise for both objectives, and Table 12 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 \$117 million and

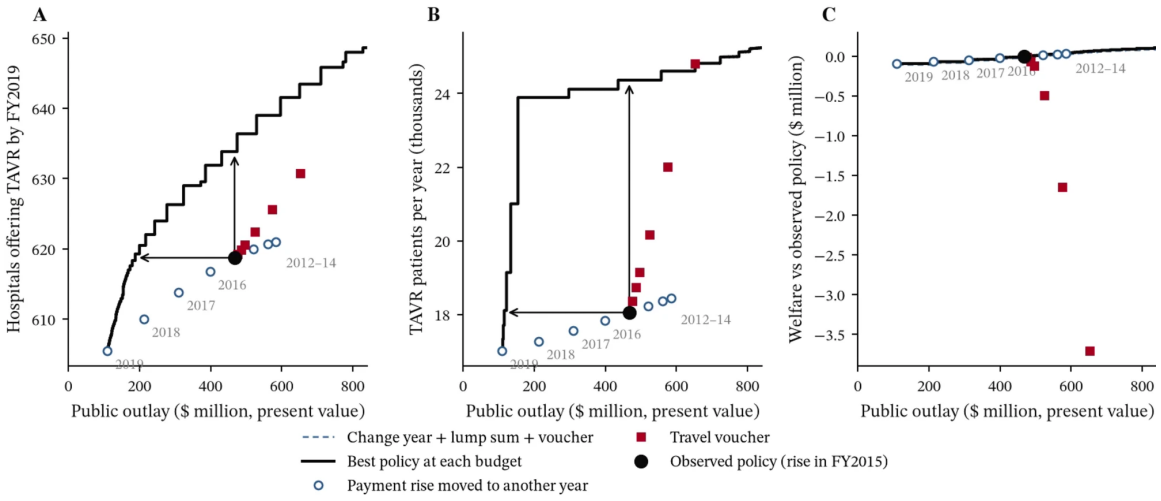



Figure 10: Policy frontier. Each policy combines a year in which the 2015 payment rise takes effect, a lump sum paid at adoption to hospitals the model predicts will stay out, and a travel voucher that pays part of the extra travel time to a TAVR hospital. Panel A reports hospitals offering TAVR by 2019, panel B TAVR patients per year, and panel C welfare $W$, patients' consumer surplus from hospital access relative to the observed policy, net of voucher cash, 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.

adds 2.3 hospitals, about $51 million per hospital; delaying it to 2019 saves $358 million and loses 13.3 hospitals, about $27 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 $212 million and leaves 610 hospitals; the coverage policy at the observed budget spends a further $218 million on lump sums and the voucher and reaches 634, about $9.2 million per additional hospital. The lump sum itself, $9.7 million, is above the median adoption cost of hospitals that never adopt, $8.9 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 2017 and spends little on lump sums. Patients treated in 2019 follow the same policies as patients per year, because most of the gain in either comes from the voucher, which moves patients to programs already open.

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 6.8 thousand and coverage by 12.0 hospitals, for $185 million. With the rise delayed to 2019 and no lump sum, it reaches 23.9 thousand patients a year, 5.8 thousand more than observed, for $154 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 12 hospitals between half and the full observed outlay, 9 more up to one and a half times it, and 7 more up to twice it: 622, 634, 643 and 650 hospitals. It never advances the rise: at the two larger budgets the best policy moves it to 2017 and 2016 and spends the money on lump sums of $11.3 and $15.4 million. The patient frontier reaches 23.9 thousand patients a year at half the observed outlay and gains little after, 24.4, 24.8 and 25.3 thousand at the three larger budgets; at one and a half times the outlay it keeps the rise in 2015 and adds a small lump sum, and at twice the outlay 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 $0.1 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 (Figure 10, panel C). On its own the voucher lowers welfare net of its cash, by up to $3.7 million at full strength, because it pays patients for travel they would not otherwise undertake; utility including the cash rises. 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, $121 million for the coverage policy at the observed budget. Because the modality rule responds to the payment, an earlier rise also raises the TAVR share at programs already operating, and the patient counts include that response.

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.

TABLE 12: 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 | 584.1 | -13.3 | 621.0 | 147.53 | 34.48 | +0.0 |
| 2013 | 560.2 | -10.5 | 620.7 | 146.87 | 34.44 | +0.0 |
| 2014 | 520.0 | -6.0 | 619.9 | 145.80 | 34.38 | +0.0 |
| 2015 (status quo) | 467.4 | +0.0 | 618.7 | 144.42 | 34.29 | +0.0 |
| 2016 | 397.9 | +7.6 | 616.8 | 142.62 | 34.17 | -0.0 |
| 2017 | 310.5 | +17.7 | 613.9 | 140.47 | 33.99 | -0.0 |
| 2018 | 212.1 | +29.1 | 610.0 | 138.22 | 33.79 | -0.1 |
| 2019 | 109.0 | +41.2 | 605.5 | 136.09 | 33.57 | -0.1 |

*Notes:* $ million, PV at 2012. 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.

TABLE 13: Counterfactual: best policy at each budget

| Objective | Budget (×SQ) | Best policy ($t_R$; $S$, $M; $k$) | Outlay ($M) | Value | vs SQ | Induced reimb. |
|---|---|---|---|---|---|---|
| *Infinite horizon (perpetuity after 2019)* | | | | | | |
| Hospitals offering TAVR by 2019 | 0.5 | 2019; 8.32; 1.00 | 217.3 | 622.0 | +3.3 | −301.0 |
| | 1 | 2018; 9.70; 1.00 | 430.6 | 633.8 | +15.1 | −309.7 |
| | 1.5 | 2017; 11.31; 1.00 | 650.5 | 643.4 | +24.7 | −318.5 |
| | 2 | 2016; 15.38; 1.00 | 894.7 | 650.4 | +31.6 | −327.7 |
| TAVR patients per year, 2012–19 (k) | 0.5 | 2019; 0.00; 1.00 | 153.7 | 23.89 | +5.84 | −301.9 |
| | 1 | 2017; 0.00; 1.00 | 436.1 | 24.36 | +6.31 | −319.7 |
| | 1.5 | 2015; 3.86; 1.00 | 694.4 | 24.82 | +6.77 | −333.0 |
| | 2 | 2012; 6.12; 1.00 | 923.1 | 25.30 | +7.25 | −348.4 |
| TAVR patients in 2019 (k) | 0.5 | 2019; 0.00; 1.00 | 153.7 | 45.11 | +10.82 | −301.9 |
| | 1 | 2017; 0.00; 1.00 | 436.1 | 45.18 | +10.89 | −319.7 |
| | 1.5 | 2015; 0.00; 1.00 | 652.2 | 45.28 | +10.99 | −333.0 |
| | 2 | 2012; 6.12; 1.00 | 923.1 | 45.41 | +11.11 | −348.4 |
| Consumer surplus vs SQ ($M) | 0.5 | 2018; 0.00; — | 212.1 | −0.1 | −0.07 | +29.1 |
| | 1 | 2015; 0.00; — | 467.4 | 0.0 | +0.00 | +0.0 |
| | 1.5 | 2012; 6.12; — | 684.3 | 0.1 | +0.07 | −26.3 |
| | 2 | 2012; 13.19; — | 891.9 | 0.1 | +0.11 | −38.7 |

*Notes:* $ million, PV at 2012; SQ is the observed policy, outlay $467.4M. 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.

## IX. Conclusion

Under a regulated price that paid TAVR no more than the procedure it replaced, adoption was financed through patients. A TAVR program raised a hospital's value to patients across the risk distribution, and most of the patients it drew received the incumbent procedure: before the rate revision, 70 percent of patients who passed over a nearest hospital without a program to be treated at an adopter received SAVR. The return to adoption therefore accrued on the hospital's whole program, and a program's value grew with its age, which is why more than two fifths of the hospitals that chose to adopt entered before the rate paid and why their entry lowered the return to those that waited.

The timing of the rate update is a weak lever for diffusion, and targeting is what works. 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. At the observed outlay both objectives favor a later rise, to 2018 for coverage and 2017 for patients treated; coverage keeps the rise at or after 2015 at every budget we consider, and only with twice the observed outlay does the patient objective advance it 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

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

TABLE 14: Patient demand: value of treatment at any hospital

| | 2012 | 2013 | 2014 | 2015 | 2016 | 2017 | 2018 | 2019 |
|---|---|---|---|---|---|---|---|---|
| *A. Year effects and risk* | | | | | | | | |
| Year effect, $\gamma_t$ | 0 | 0.134 | 0.175 | 0.098 | 0.138 | −0.022 | 0.088 | −0.237 |
| Risk, $\kappa_t$ | −0.317 | −0.396 | −0.425 | −0.430 | −0.489 | −0.559 | −0.688 | −0.701 |
| *B. Conditions, $\kappa_{kt}$* | | | | | | | | |
| Balloon valvuloplasty | 0.027 | 0.164 | −0.784 | −0.008 | −0.017 | −0.702 | 0.239 | 0.098 |
| Coronary bypass | −0.691 | −0.839 | −0.842 | −1.010 | −1.183 | −1.442 | −1.446 | −1.490 |
| Home oxygen or respiratory failure | −0.530 | −0.622 | −0.655 | −0.680 | −0.764 | −0.638 | −0.662 | −0.812 |
| Peripheral artery disease | −0.170 | −0.150 | −0.208 | −0.158 | −0.128 | −0.090 | −0.253 | −0.525 |
| Aortic atherosclerosis | 0.776 | 0.752 | 0.728 | 0.635 | 0.414 | 0.452 | 0.239 | 0.563 |
| Skilled-nursing or home-health use | −1.740 | −1.695 | −1.732 | −1.750 | −1.714 | −1.721 | −1.591 | −1.445 |
| Dialysis or CKD stage 4–5 | −0.049 | −0.097 | −0.102 | −0.042 | −0.209 | 0.062 | 0.027 | 0.370 |
| Prior valve surgery | −0.797 | −0.865 | −0.917 | −0.838 | −0.695 | −0.647 | −0.659 | −0.463 |

*Notes:* Terms that shift the value of treatment at any hospital relative to no treatment, from the demand model of Table 9; the 2012 year effect is normalized to zero. Not shown: the risk interactions with hospital size and teaching status, $\omega_1 = -0.0010$ and $\omega_2 = 0.0005$. Standard errors to be added.

TABLE 15: AVR episodes by procedure: resource use and payment

| | 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.

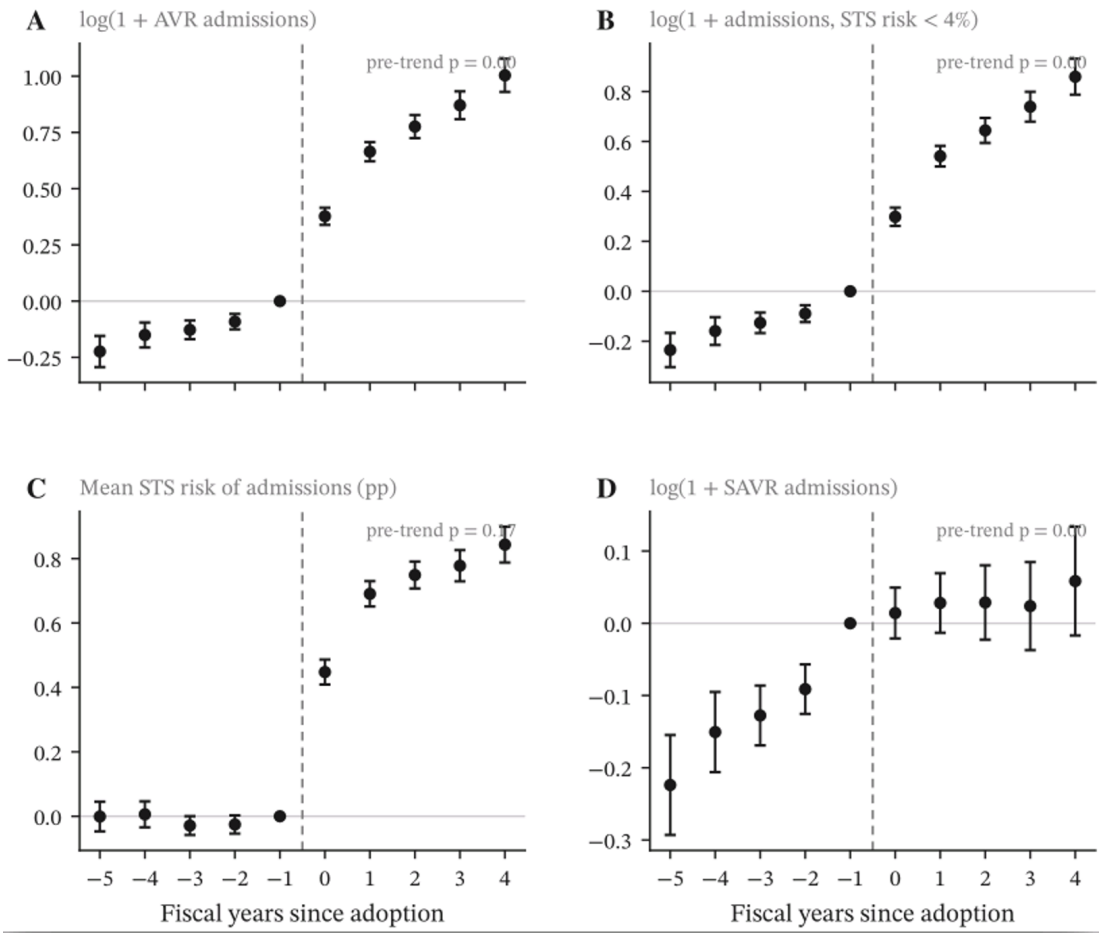



Figure 11: AVR admissions around a hospital's TAVR adoption. Interaction-weighted event-study coefficients (Sun and Abraham, 2021) with hospital and year effects, never-adopters as the control cohort and adoption cohorts 2012–2019 (pivotal-trial sites excluded); relative year $-1$ omitted; 95 percent confidence intervals clustered by hospital. Panels: (A) $\log(1 + \text{AVR admissions})$; (B) $\log(1 + \text{admissions with STS risk below 4 percent})$; (C) mean STS risk of admissions, in percentage points; (D) $\log(1 + \text{SAVR admissions})$. The pre-trend $p$-value is the joint test that the coefficients for relative years $-5$ to $-2$ are zero.

### 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

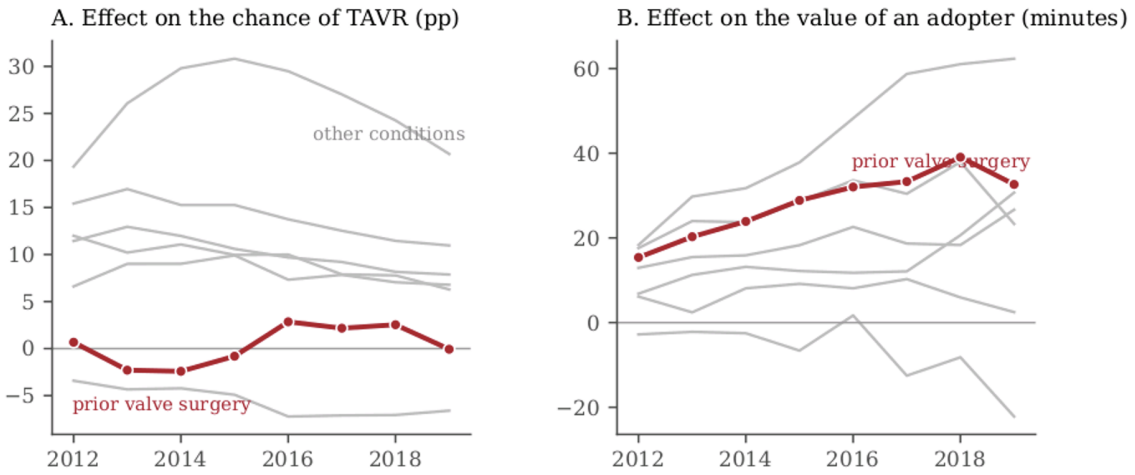



Figure 12: Conditions and the value of an adopter. Panel A: effect of each condition on the probability of TAVR, in percentage points, from a linear probability model on patients treated at an adopter with STS-tier effects, by fiscal year. Panel B: the condition's term in the value of an adopter, in minutes of travel. Red: prior valve surgery; grey: prior coronary bypass, home oxygen, peripheral artery disease, aortic atherosclerosis, skilled-nursing or home-health use, dialysis.

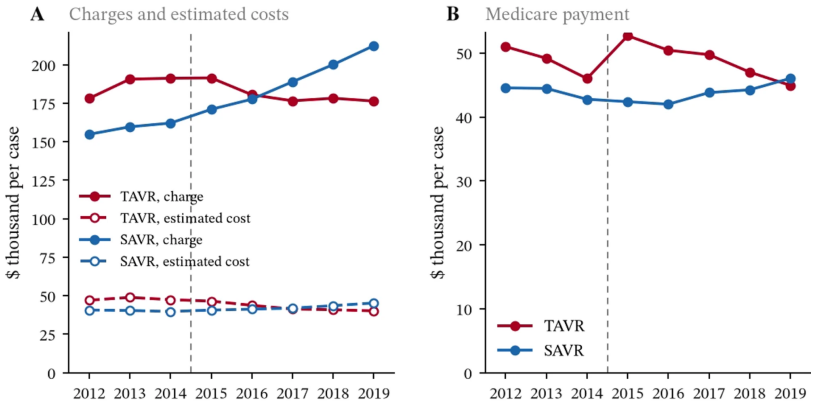



Figure 13: Charges, estimated costs and Medicare reimbursement per case for TAVR and SAVR by fiscal year. Medians across episodes. Panel A shows charges and costs estimated with hospital-year cost-to-charge ratios; Panel B the Medicare payment.

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

$$\log(Y_{iht}) = \beta_0 + \beta_1 \cdot \text{Early Adopter}_h + \beta_2 \cdot \text{Post Reimb}_t + \beta_3 \cdot (\text{Early Adopter}_h \cdot \text{Post Reimb}_t) + \mathbf{Z}'_{ht}\theta + \delta_s + \epsilon_{iht} \tag{14}$$

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 $\delta_s$ are city fixed effects. Selection of high-benefit patients by early adopters predicts $\beta_3 < 0$.

TABLE 16: SAVR patient characteristics and TAVR adoption timing: difference in differences

| | 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) |
| City fixed effects | Yes | Yes |
| Observations | 15,232 | 15,232 |
| $R^2$ | 0.082 | 0.073 |

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

In Table 16, SAVR DRG weight falls by about $7\%$ system-wide after 2015, consistent with higher-risk patients shifting to TAVR. The Early Adopter × 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 14 shows the within-hospital TAVR share shifting right after 2014. We estimate

$$\log(\text{SAVR}_{h,t}) \;=\; \alpha \;+\; \beta_1\,\text{Adopt}_{h,t} \;+\; \beta_2\,\text{Adopt}_{h,t}\times\mathbf{1}\{t \geq 2015\} \;+\; \gamma_h \;+\; \delta_{t\times\text{HRR}} \;+\; \varepsilon_{h,t}, \tag{15}$$

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

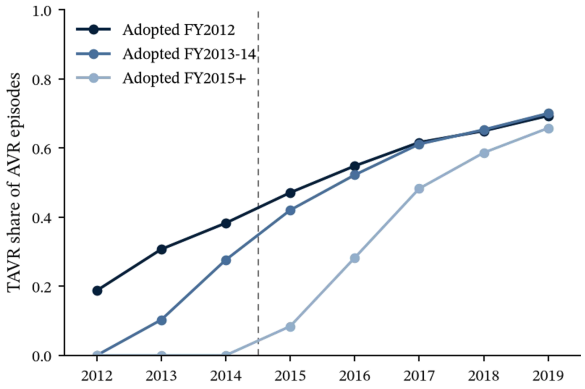



Figure 14: Within-hospital TAVR shares by fiscal year (early adopters)

TABLE 17: Effect of TAVR adoption on SAVR volumes

| | $\log(\text{SAVR volume})$ |
|---|---|
| Adopt | 0.032 |
| | (0.045) |
| Adopt $\times$ $\mathbf{1}$(fiscal year $\geq$ 2015) | $-0.141^{**}$ |
| | (0.060) |
| Hospital fixed effects | Yes |
| Fiscal year $\times$ HRR fixed effects | 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 17 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

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

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 $F_i^E \leq F_i^L$ 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 $F_i \leq F_i^E$ ($\mathcal{E}$), at $t_R$ if $F_i^E < F_i \leq F_i^L$ ($\mathcal{M}$), 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 $V_i^E \geq V_i^L$ and $V_i^E \geq 0$, which hold for $F_i$ below the smaller of the two roots, $F_i^E$. Late entry requires $V_i^L \geq 0$, i.e. $F_i \leq L_i$, and $V_i^L \geq V_i^E$. For $F_i > F_i^E$ either $V_i^E < V_i^L$ or $V_i^E < 0$, and in the latter case $V_i^L \geq 0$ implies $V_i^L > V_i^E$; so the late set is $(F_i^E, L_i]$, empty if $L_i \leq F_i^E$, 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 $\ell_i$ by a gain $v_{i,1}\mu_2$ and leaves $L_i$ unchanged. It moves hospitals in $\mathcal{M}$ with $F_i \leq (R_i + v_{i,1}\mu_2)/(1-\beta)$ from $t_R$ to stage 1; it recruits from $\mathcal{S}$ only hospitals with $L_i < F_i \leq \Delta\pi_{i,1} + v_{i,1}(\mu_2 - \mu_1) + \beta\Delta\pi_{i,2}^E$, 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 $\mathcal{R} = \{L_i < F_i \leq L_i + s\} \subset \mathcal{S}$ are recruited and enter at $t_R$;*

*(ii) hospitals in $\mathcal{D} = \{F_i^E - \frac{\beta}{1-\beta}s < F_i \leq F_i^E\} \subset \mathcal{E}$ 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 $v_{i,1}\,g(F_i^E)\,\beta \geq v_{i,2}\,g(L_i)\,(1-\beta)$.*

*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 $\nu_i$ independent of $x_i$, mean zero and scale $\sigma_\nu$, 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.

SUN, L. AND S. ABRAHAM (2021): "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 225, 175–199.

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.