Diffusing Innovations Under Market Competition: Evidence from Drug-Eluting Stents

Ginger Zhe Jin[1], Hsienming Lien[2], and Xuezhen Tao*[3]

[1]University of Maryland & NBER
[2]National Chengchi University & CIER
[3]Shanghai University of Finance and Economics

May 29, 2026

Abstract

This paper examines how hospital competition and insurance reimbursement policies shape the diffusion of medical innovations. Using patient-level data from Taiwan's drug-eluting stent (DES) market in the Taipei area, we estimate a structural model that integrates patient demand with hospitals' endogenous portfolio and pricing decisions. We document a fundamental trade-off: intensified competition reduces prices but weakens hospitals' incentives to adopt new technologies. Counterfactual simulations that vary market structure and incorporate selective contracting or targeted patient coupons show that the effectiveness of reimbursement policies depends critically on hospitals' strategic responses and the surrounding competitive environment.

JEL: D4, I18, L13, O33.

Keywords: drug-eluting stents (DES), technology diffusion, hospital competition, selective contracting, reimbursement policy.

*We thank Leila Agha, Jeffrey Macher, Pierre Bodéré, Andrew Ching, Lucy Wang, and seminar and conference participants at the National Taiwan University, Research Center for Humanities and Social Sciences at Academic Sinica, Shanghai University of Finance and Economics, the 2024 Annual Health Econometrics Workshop, the 2024 Asian-Pacific Industrial Organization Conference, the 2025 International Industrial Organization Conference, the 2025 DC-IO day, Georgetown University, University of Illinois at Urbana-Champaign, University of Pennsylvania, Purdue University and the 2026 ASSA meetings for constructive comments. Our data was generously provided by the National Health Insurance Database of Taiwan. Lien appreciates financial support from the Ministry of Science and Technology (Grant no: MOST 108-2410-H-004-027-MY2). All rights reserved. All errors are ours.

1 Introduction

Rising health expenditure poses a critical challenge for economies worldwide, with the diffusion of new medical technologies as a primary driver of both improved outcomes and rising costs (Chernew and Newhouse, 2011; Chandra and Skinner, 2012; Skinner and Staiger, 2015; Dunn et al., 2023, 2024). This dual impact creates a fundamental policy tradeoff: promoting timely access to valuable innovations while maintaining the sustainability of health spending. While governments employ reimbursement design and market structure regulation as the principal levers to navigate this tradeoff, the effectiveness of these policies is not automatic (Nexon and Ubl, 2010; McClellan, 2011; Chandra et al., 2024). Their ultimate impact is mediated by the strategic responses of healthcare providers, most notably hospitals and doctors, who act as the crucial intermediaries between innovation and patient access (Chandra et al., 2011). Understanding the behavior of these downstream agents is therefore paramount to designing effective health policy.

As gatekeepers, these intermediaries translate the potential of a new technology into realized market outcomes. Their strategic choices on whether, when, and at what price to offer new products are shaped by a complex interplay of economic forces. Operating in local markets, they face competitive pressures from rival providers, heterogeneous demand from patients, and the specific constraints and opportunities embedded in reimbursement rules. These forces govern their core decisions on technology adoption and pricing, which in turn can accelerate or delay diffusion and ultimately determine the allocation of surplus among patients, providers, manufacturers, and the public payer.

To investigate the role of intermediaries, we analyze the cardiac stent market in Taipei, Taiwan. Before bare metal stents (BMS), balloon angioplasty was used to compress arterial plaque, but this method lacked durable effectiveness. BMS improved outcomes by keeping arteries open, yet caused restenosis within the stent in 20 to 30% of patients due to scar tissue formation. Drug eluting stents (DES) addressed this issue by releasing medication to inhibit scar growth, offering superior clinical benefits at a higher cost (Tu et al., 2007; Kandzari et al., 2017). Further advancements in DES technology, such as the use of stronger alloys (e.g cobalt chromium), ultrathin struts, and more biocompatible drugs (e.g. Zotarolimus), have significantly reduced the risk of stent thrombosis, making newer generations of DES a major improvement over earlier stent models (Navarese et al., 2014; Kandzari et al., 2017; Bangalore et al., 2018).[1]

[1]For instance, Navarese et al. (2014) shows that the new generation DES demonstrated a 22% reduction in the odds of myocardial infarction compared with the old generation devices, without increasing mortality. Among these, everolimus eluting stents (EES) significantly lowered the risk of stent thrombosis compared to paclitaxel eluting stents (PES). While both generations showed comparable efficacy against restenosis,

This market of analysis operates within the framework of Taiwan's National Health Insurance (NHI), a universal coverage system that provides an ideal empirical context for our research. Three features of this setting are particularly salient. First, the Taipei area is a decentralized, competitive hospital market where numerous providers make portfolio choices. This allows us to observe a rich set of strategic decisions as hospitals decide whether to *adopt* a new manufacturer's brand or *upgrade* an existing one. Second, Taiwan NHI's "top up" reimbursement policy creates a stark financial tradeoff for patients: BMS are fully covered and effectively free, while DES require large out of pocket payments in addition to NHI's base stent reimbursement to hospitals. This payment policy directly shapes patient demand, constrains hospital pricing and, more importantly, influences hospitals' decision to adopt new DES. Third, two systemwide, exogenously timed cuts to the base stent reimbursement provide valuable shocks to hospital margins, aiding the identification of supply side responses.

This rich institutional setting enables us to formalize and investigate our central research question: how do market structure and reimbursement design jointly shape hospitals' technology adoption and pricing decisions, and what are the resulting equilibrium effects on patient welfare, provider profits, and public spending? Answering this question poses a significant analytical challenge. The observed market outcomes, including prices, technology availability, and patient choices, are equilibria resulting from the interaction of latent patient demand and strategic hospital supply. A hospital's incentive to invest in a new technology, for instance, depends on its expectations of patient demand and the competitive responses of its rivals. To disentangle these interdependent forces and recover the underlying behavioral and cost parameters, a structural approach is required.

Our analysis proceeds by developing and estimating a structural model of the Taipei DES market. We specify a discrete choice model for patient demand that accommodates rich heterogeneity in preferences for clinical novelty, price, and hospital characteristics, using a mixed logit formulation. On the supply side, we model hospitals as strategic agents in a two stage quarterly game. In the first stage, they make a portfolio decision among adopting a new brand, upgrading an existing one, or maintaining the status quo, weighing expected profits against the fixed costs of adjustment. In the second stage, conditional on the marketwide portfolio configuration, hospitals engage in Bertrand Nash pricing competition.

EES stands out as the safest option. Bangalore et al. (2018) shows that newer ultrathin strut drug eluting stents (DES) reduce target lesion failure (TLF) by approximately 16% at one year.Kandzari et al. (2017) employs randomized 1,334 patients and found that the third generation stent (ultrathin struts, bioresorbable polymer) was superior to the second generation Xience stent (thin struts, durable polymer) for the primary endpoint of target lesion failure at 12 months (6.2% vs. 9.6%).

To identify the parameters of this model, we combine a unique set of datasets. The core of our analysis is built on individual level inpatient claims, which provide granular information on patient demographics, clinical conditions, and the specific stent choices made. Crucially, we supplement these administrative records with a manually collected hospital, model, and time panel of the exact out of pocket prices faced by patients. To account for socioeconomic heterogeneity, we merge these files with township level data on average household reported taxable income. This fusion of clinical, pricing, and income data provides the rich empirical foundation necessary to estimate our model.

The model is estimated via simulated maximum likelihood, with a control function approach to address the endogeneity of prices. Our demand side estimates reveal the core tension facing providers: patients are price sensitive, particularly those with lower income or greater comorbidity, but simultaneously place a substantial premium on access to the "newest generation" technologies. On the supply side, the estimates for the fixed cost parameters provide direct evidence on the economics of diffusion, confirming that upgrading an existing product line is significantly less costly for a hospital than adopting an entirely new brand, which rationalizes the observed path of technology diffusion.

The estimated model provides a rich framework for quantitatively evaluating alternative policy regimes. We conduct two sets of counterfactual simulations. First, we analyze the interplay of market structure and patient demand. Simulating equilibria under varying numbers of competing hospitals reveals a fundamental tradeoff: while more intense competition disciplines prices and benefits consumers, it also erodes hospital profit margins, thereby weakening the incentive to invest in frontier technologies. Counterfactual simulations show that, as the number of hospitals increases, intensified price competition dominates reduced adoption incentives, leading to a net gain in consumer surplus and total welfare. We also find that strong patient preferences for "newness" act as a powerful demand side pull for innovation, but the distribution of the resulting surplus gains between patients and hospitals depends critically on the intensity of competition.

Second, we evaluate the effectiveness of alternative reimbursement designs. A supply side policy of selective contracting, in which the NHI leverages its bargaining power to negotiate wholesale price discounts with specific manufacturers in exchange for exclusive reimbursement, can generate a "quadruple win" that raises consumer surplus, curbs the payer's expenditure, accelerates diffusion of the newest-generation devices, and sustains participating manufacturers' revenue, with hospitals also benefiting from the negotiated discount. However, we find that this outcome is most robust in concentrated markets, where exclusive contracts allow hospitals to retain sufficiently large rents to justify upgrading their DES

portfolios. In more competitive markets, by contrast, rent dissipation limits the policy's leverage.

Increasing the NHI's DES specific reimbursement rate (applied uniformly across all DES models) has the opposite pattern: it is most effective in boosting DES utilization in competitive markets, where high pass through rates lead to price reductions but the policy also imposes a substantial fiscal burden on the insurer. Alternatively, a demand side patient coupon targeting lower income individuals can improve equity with limited marketwide effects when the size of the program is not sufficiently large or the market is competitive enough so that hospitals have minimal incentives to adjust prices or product offerings. When more patients are eligible for patient coupons, hospitals, especially those facing less market competition, have a stronger incentive to raise the list price for all patients, which subsequently hurts unsubsidized patients and entails more government spending to fund patient coupons.

Taken together, these findings highlight the importance of accounting for hospitals' strategic responses and market structure when designing policies to promote the diffusion of medical innovations.

Our paper contributes to research on competition and innovation by shifting attention from upstream R&D incentives to downstream diffusion. Classic work studies how market structure shapes investment incentives (Schumpeter, 1942; Arrow, 1962; Aghion et al., 2005); in our setting, hospitals are intermediaries whose portfolio choices govern whether innovations reach patients. Modeling adoption and upgrading as discrete portfolio choices also connects our supply side to recent empirical IO work on product entry and portfolio design (Head and Mayer, 2019; Wollmann, 2018; Demirer and Olssen, 2024).

We also contribute to the literature on medical technology diffusion and provider responses to reimbursement policy. Prior work emphasizes intellectual property, physician learning, payment incentives, and technology-driven disparities (Kyle, 2007; Conley and Udry, 2010; Skinner and Staiger, 2015; Dafny, 2005; Dunn et al., 2023, 2024). Our contribution is to estimate a unified equilibrium model in which hospital portfolio choice, patient demand, prices, and public spending are jointly determined. This framework complements work on hospital-manufacturer bargaining in cardiac stents (Grennan, 2013, 2014; Grennan and Swanson, 2020) by focusing on the downstream patient-facing market.

Finally, our counterfactuals speak to selective contracting and targeted subsidies in healthcare. Selective contracting is widely studied in drug formularies, provider networks, and pharmacy benefit management (Duggan and Scott Morton, 2010; Sorensen, 2003; Pakes et al., 2015; Ho and Lee, 2019; Starc and Swanson, 2021). We analyze how such policies affect not only prices and access, but also hospitals' incentives to carry frontier technologies. The patient coupon simulations similarly connect to work on demand-side subsidies and

equity, emphasizing how hospitals' strategic pricing can mediate the incidence of targeted assistance (Decarolis, 2015; Tebaldi, 2025; Liu et al., 2021; Olateju et al., 2024).

The rest of the paper is organized as follows. Section 2 describes the background and summarizes the data in our analysis sample. Section 3 lays out our structural model of patient choice of hospital and stent type on the demand side and hospitals' pricing and portfolio management decisions on the supply side. Section 4 presents our estimates, and Section 5 conducts three counterfactual simulations to highlight the interplay of market competition, patient willingness to pay for newness, and government reimbursements for DES. A brief conclusion is offered in Section 6.

2 Data and Institutional Background

We begin by outlining the key institutional features of the cardiac stent market in Taiwan, from the underlying clinical trade-offs between stent technologies to the specific reimbursement policies that shape the economic incentives of patients and hospitals. We then describe the unique datasets we assemble to analyze behavior within this setting.

2.1 Procedure of Cardiac Stent

The treatment of coronary artery disease encompasses a spectrum of options, ranging from non-invasive medication to open-heart surgery. Between these extremes, percutaneous coronary intervention (PCI) with stent implantation has become a standard, less invasive procedure. The intervention involves inserting a stent—a small, expandable mesh tube—to scaffold a narrowed artery and restore blood flow. The first wave of these devices, bare-metal stents (BMS), were effective scaffolds but were limited by high rates of in-stent restenosis, a re-narrowing of the artery caused by tissue regrowth (Agostoni et al., 2005).

To address this limitation, drug-eluting stents (DES) were developed. These devices are coated with antiproliferative drugs that significantly reduce the risk of restenosis. This clinical benefit, however, comes with significant economic and medical trade-offs. The drug coating can delay arterial healing, which in early-generation models was associated with a higher risk of stent thrombosis (blood clot formation) that happened in 20-30% of patients, particularly for older patients or those with severe comorbidities. Consequently, DES require a longer course of dual antiplatelet therapy (DAPT) to mitigate this risk (Tu et al., 2007; Kaiser et al., 2010).

Even within drug-eluting stents (DES), notable advancements over time have significantly improved clinical safety and efficacy—an evolution central to the technological "newness"

incorporated into our demand model. The initial iteration of DES used strong polymers (like paclitaxel) to stop scar tissue from forming, but the thick coating irritated the artery, slowed healing, and slightly increased the risk of dangerous blood clots. Later iterations of DES improved by using thinner, stronger metals (like cobalt-chromium) and more natural, body-friendly drug coatings (like zotarolimus), giving the same benefits with much lower risks (Navarese et al., 2014; Bangalore et al., 2018; Kandzari et al., 2017).

2.2 Top-up Payments for Cardiac Stent

The healthcare environment in which these clinical decisions are made is defined by Taiwan's National Health Insurance (NHI) program.[2] During our study period, the NHI employed a "top-up" reimbursement model for stents: hospitals received a fixed base payment for any stent procedure, set at a level sufficient to cover the wholesale cost of a BMS. This structure created a powerful set of incentives. For patients, it meant that BMS were effectively free at the point of care, while the choice of a more advanced DES required a substantial out-of-pocket payment equal to the hospital's posted retail price net of the fixed NHI reimbursement. For hospitals, it created a margin management problem: their profit from a DES procedure depended on the wholesale price negotiated with manufacturers, the retail price set for patients, and the fixed NHI payment.

This system was subject to two system-wide reimbursement cuts during our sample period: a 26% reduction from 27,000 to 19,940 NTD (approximately $1 USD for 30 NTD) in 2009 and a further 18% reduction in 2012.[3] These centrally mandated cuts directly compressed hospital revenue per stent usage, providing exogenous variation that aids in the identification of supply-side responses. The combination of a fixed base payment and a patient-borne top-up cost makes patients highly sensitive to the posted price, rendering it a central strategic variable for hospitals competing for demand.

Finally, we assume that patients select the optimal combination of hospital and stent type to maximize their utility. Although physicians may influence the final stent choice, several factors suggest that patients play a substantial role in this decision. Stent implantation is

[2]Beyond NHI coverage, Taiwanese residents may also purchase private health insurance to offset out-of-pocket spending. However, despite its broad availability, private insurance is largely supplemental and typically offers limited reimbursement. As shown below, we observe price-elastic demand for DES across nearly all demographic groups, indicating that patients remain price-sensitive even when private insurance can partially defray the patient's out-of-pocket costs after NHI benefits are applied.

[3]These cuts were partly driven by declining wholesale costs of BMS. Though the NHI is keen to reduce health spending as a public insurer, it determines stent reimbursement rates through periodic market transaction price surveys, setting the payment equal to the volume-weighted average procurement price across hospitals plus a "reasonable zone" markup of approximately 20% to cover handling, storage, and administrative costs (Hsu and Lu, 2015).

typically an elective procedure; the cost of the procedure itself (except for the out-of-pocket payment required for DES) is fully covered by the NHI regardless of hospital choice; and patients can compare out-of-pocket expenses across hospitals and DES models through the NHI's website. Together, these conditions foster an environment conducive to price-sensitive selection of both hospital and stent type.

2.3 Hospital Procurement and Portfolio Choice

Our supply-side model treats hospitals as the primary decision makers over DES adoption and product mix, an assumption that matches Taiwan's institutional setting. In contrast to the U.S. market, where interventional cardiologists may practice independently and maintain direct relationships with manufacturers, cardiologists in Taiwan are typically salaried hospital employees who do not negotiate device contracts, receive volume-based rebates, or independently determine hospital formularies. DES portfolio and price decisions therefore operate through hospital procurement and device review processes rather than through individual physician contracting. Because we do not observe physician compensation or physician-patient discussions, we do not model physician-level selection separately and instead focus on hospital-level pricing and portfolio management.

This hospital-level autonomy is reinforced by the decentralized structure of procurement. After a DES receives approval from Taiwan's Food and Drug Administration, it can be sold through authorized distributors to qualified hospitals.[4] Hospitals then choose which products to review, negotiate wholesale prices with manufacturers or local distributors, and set the retail prices faced by patients.[5] This sequence supports modeling adoption and upgrading as hospital choices rather than as manufacturer-controlled entry.

The observed market data lend additional support to this interpretation. Taiwan's Ministry of Health and Welfare operates a "Self-Paid Medical Device Price Comparison Website" that enables patients to compare prices for identical devices across hospitals, with coronary DES serving as a featured example.[6] Disclosure rules further require hospitals to report the DES models they have adopted, together with prices and product characteristics. These

[4]The TFDA approval process is product-level: manufacturers submit applications, and upon approval, the device may be sold to qualified healthcare institutions in Taiwan. See Credevo (2024).

[5]Public procedural documents, for instance, Mackay Memorial Hospital's "New Medical Device Application and Review Process," outline a hospital-led workflow in which a clinical department submits a request, the procurement unit confirms regulatory approval and reimbursement eligibility, and a hospital committee evaluates the product prior to adoption. The document is available on the Purchasing Division's downloads page: `https://www.mmh.org.tw/departmain3.php?id=99`. Similar review procedures exist at other major hospitals

[6]Ministry of Health and Welfare, "Self-Paid Medical Device Price Comparison Website," available at `https://www.mohw.gov.tw/cp-4631-54371-1.html`.

records reveal substantial cross-hospital variation in both portfolios and prices, even among hospitals belonging to the same hospital system.[7] This stands in sharp contrast to the United States, where GPO contracting and physician preference have been shown to shape device adoption and pricing (Grennan, 2013; Grennan and Swanson, 2020; Zhu et al., 2023). Together with the decentralized procurement procedures described above, the observed within-system variation in portfolios and prices locates the relevant strategic margin at the hospital level, motivating the supply-side specification we develop below.

2.4 Data Sources and Sample Construction

Our study population is defined as all patients receiving any stent treatment in the competitive Taipei metropolitan area, which includes Taipei City, New Taipei City, Keelung City and Yilan County, from 2007 to 2013. We rely on three comprehensive datasets.

First, we use individual-level inpatient claims from Taiwan's National Health Insurance (NHI). These administrative records contain detailed information on patient demographics, diagnoses (including the Charlson Comorbidity Index), and the specific stent model implanted. The presence of unique patient and hospital identifiers allows us to construct a complete treatment history for every patient at each hospital in our sample.

Second, because NHI claims do not capture the out-of-pocket payments that are central to the patient's choice problem, we supplement the claims data with a manually collected panel of hospital-specific DES prices. These data are compiled from public reports on the NHI website. By imputing these prices for every cardiac stent across all hospitals and quarters, we construct a hospital-model-time panel of the exact prices faced by patients, which we successfully match to approximately 90% of the DES procedures in our claims data.[8]

[7] Only two hospitals in our Taipei sample belong to the same hospital system (Taipei Chang Gung Memorial Hospital and Keelung Chang Gung Memorial Hospital).

[8] We imputed cardiac stent prices using NHI-reported data through several steps. First, we obtained DES prices by type and brand directly from the NHI, covering 44 models across 102 hospitals between December 2006 and November 2013, with an average of 126 reports per hospital. To address missing observations, we calculated quarterly average prices for each hospital-model combination. Gaps between reports were filled under the assumption that hospitals maintained the same price between consecutive submissions—an assumption strongly supported by the data, as 91% of quarter-to-quarter observations for chosen DES show zero price change. In cases where hospitals delayed reporting, we assumed their initial reported price applied up to four quarters prior to the first submission. This constructed price series matched 89.99% of DES usage records in the NHI claims data. In the final estimation sample, price coverage is substantially higher: of the 29,646 DES procedures selected by patients, only 59 (0.20%) lack a matched price, concentrated in four hospitals and spread across the full sample period without systematic clustering by brand or time. The overall DES choice set, including unchosen alternatives, has a 98.4% price match rate. The small number of unmatched observations are excluded from estimation. Further details are provided in the appendix of Jin et al. (2025).

Third, to properly account for socioeconomic heterogeneity, we augment the claims data with a measure of patient income. While NHI premiums are based on salary and can serve as an income proxy, a large portion of our sample consists of retirees. We therefore incorporate data from the Ministry of Finance that report the average annual taxable income per household at the level of a patient's township of residence. Our demand model uses it as a robust proxy for a patient's economic status to explain variations in both price sensitivity and technology preference.

A final crucial element of our data construction is the systematic tracking of innovation. We categorize each DES model by manufacturer and by generation, defining generation by timing to market within each manufacturer rather than across manufacturers. This brand-specific definition reflects the fact that manufacturers improve DES along different clinical and design dimensions and also matches how patients are likely to perceive "newness." For example, Medtronic generations trace successive zotarolimus-eluting platforms within Medtronic's own product line, while Abbott and Boston Scientific follow separate product sequences; we therefore do not equate a third-generation Medtronic product with a third-generation product from another manufacturer. We consolidate near-simultaneous releases with similar specifications within the same brand, such as Abbott generations 1 and 2 and Boston Scientific generations 4 and 5, and group infrequently used brands into an "Other" category. This classification is a market-level proxy for perceived technological newness and need not imply that every generation step represents the same objective clinical improvement. These product-level data define the demand-side newest-generation indicator and the two hospital actions in the supply model: *adoption*, when a hospital first offers a brand, and *upgrade*, when it replaces an existing generation with a newer one of the same brand. For all analyses, we assume the basic BMS option is available in every hospital's portfolio in all periods. A complete list of major DES models and the generation assignments is reported in the Online Appendix.

As detailed in Jin et al. (2025), only major and minor teaching hospitals are allowed to perform PTCA procedures and cardiac stent placements in Taiwan, but the volume of stent treatment (and subsequent financial implications) vary significantly among hospitals. We limit our analysis to hospitals that performed at least 75 PTCA cases and 50 stent cases annually in our sample period. This is to ensure that our analysis was not driven by hospitals that had few stent surgeries and thus may not engage in strategic pricing and portfolio management as other hospitals. In spite of this criterion, we still capture more than 98% of stent cases performed by hospitals in Taipei area, with 61,645 patient claim records across 20 hospitals in our final data sample (2007-2013).

2.5 Descriptive Evidence and Stylized Facts

The data reveal several key stylized facts about the Taipei DES market. First, the market is characterized by strong intergenerational substitution. Figure 1 shows that following the introduction of a new DES generation by a manufacturer, the market share of the prior generation typically declines sharply. This rapid substitution suggests a powerful latent demand for technological novelty from patients and physicians, which motivates the inclusion of a newest-generation attribute in our demand model.



Figure 1: DES Quantity of Different Generations by Brand

Second, pricing patterns reflect both frontier competition and adjustment frictions. As shown in Figure 2, the average patient-paid price for a given brand tends to track the price of its newest available generation, indicating that the locus of competition is at the technological frontier. At the same time, substantial price dispersion remains across brands and hospitals, which points to frictions in hospitals' portfolio and pricing adjustments.

Third, the timing of technology adoption and upgrading demonstrates some heterogeneity across hospitals. Figure 3 plots the distribution of delays, measured in quarters from a new

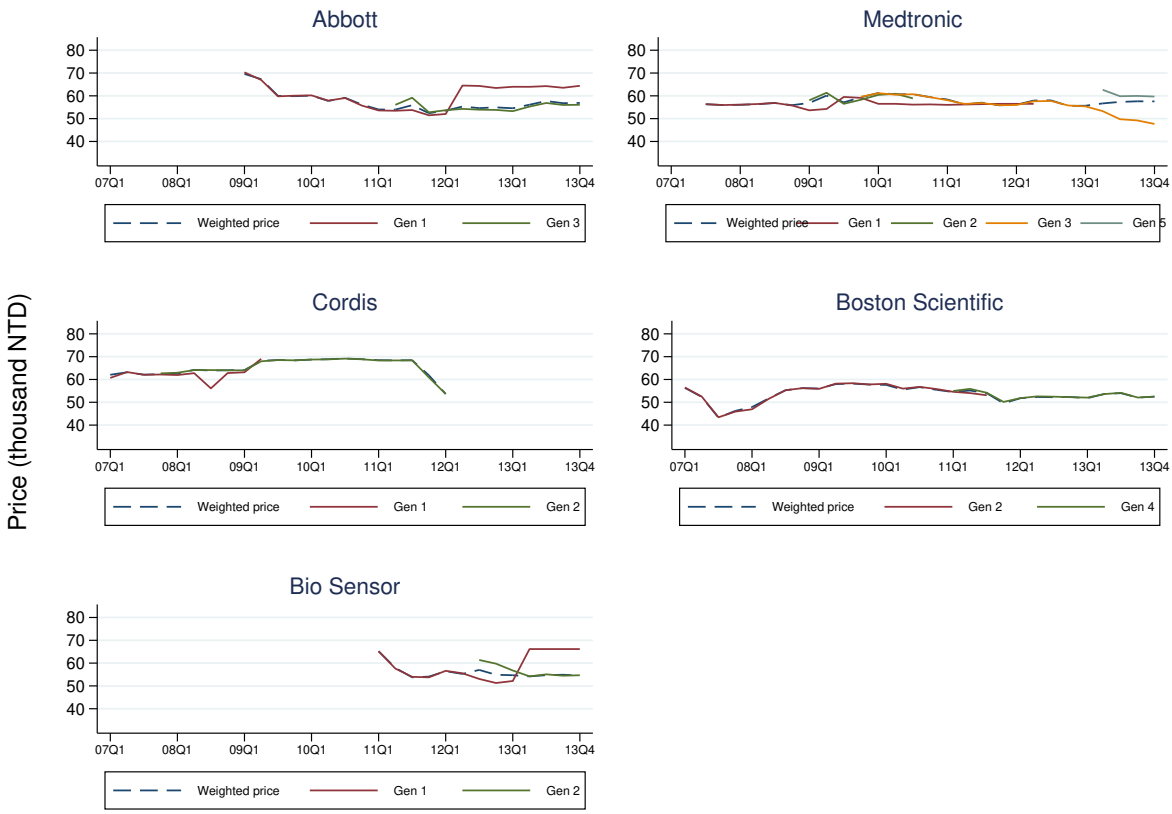


Figure 2: Average Price Across Generations

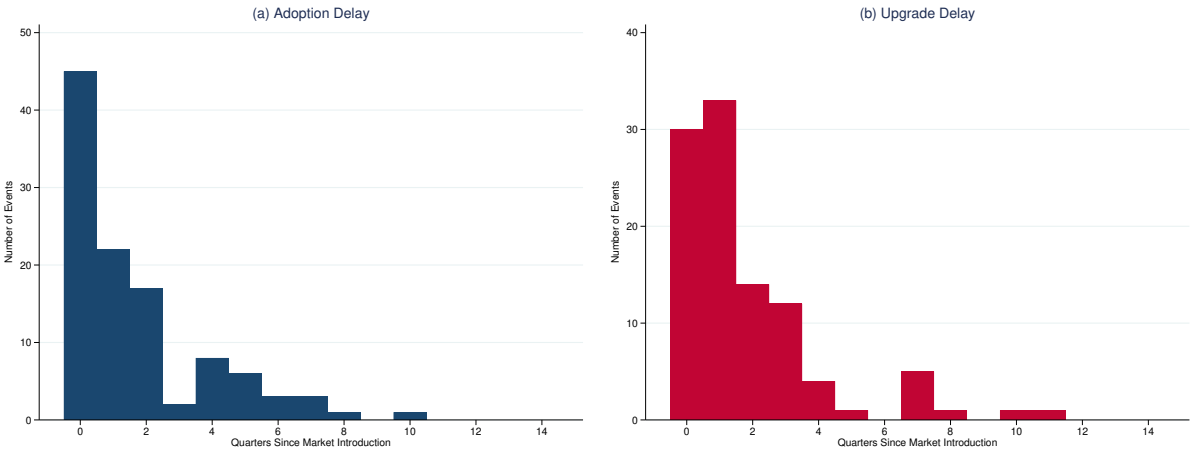


Figure 3: Distribution of Adoption and Upgrade Delays (in quarters)

product's market introduction to its appearance in a hospital's portfolio, conditional on the hospitals that eventually include the product in their portfolio. While the observed dispersion in delay suggests some market friction, most adoption and upgrading events occur relatively

quickly, typically within one to three quarters of a new product's market availability, which limits the salience of long-term strategic waiting. This heterogeneity is a central feature of the diffusion process and calls for hospital-specific fixed costs of portfolio adjustment, a key component of our supply-side model.

Finally, summary statistics confirm the powerful role of patient price sensitivity. As shown in Table 1, despite the clinical advantages of DES, the free BMS option consistently retains more than 50% of the market in our sample period (2007-2013). The out-of-pocket prices for DES are substantial, averaging 52,000–65,000 NTD, a figure comparable to the average monthly wage in Taiwan during this period and almost three times the average reimbursement rate from the NHI. This underscores the salience of the patient's cost-sharing burden and motivates our focus on out-of-pocket price as a primary determinant of demand.

Table 1: Summary Statistics on Patients' Stent Usage and Choices

	Mean	Std. Dev.	Min	Max
Number of patients per hospital-quarter	108.15	92.29	1	377
Patients using DES	52.56	62.94	1	279
Patients using BMS	56.74	39.90	1	251
Number of hospital-brand choices	90.18	19.19	41	114
Number of brand choices within hospital	4.55	1.39	1	7
Patient-paid prices of DES (1,000 NTD)				
Abbott	56.90	7.66	30.66	78.75
Bio Sensor	55.40	7.44	44.36	68.22
Cordis	65.69	3.90	43.31	74.56
Medtronic	57.40	5.95	29.71	67.96
Boston Scientific	52.45	9.52	26.09	72.00
Other	52.22	6.08	32.92	70.05
Patient chooses DES (1 = DES, 0 = BMS)	48.09%	49.96%	0%	100%

Taken together, these patterns—rapid but incomplete intergenerational substitution, pricing that reflects both frontier competition and adjustment frictions, heterogeneous adoption timing, and persistent price sensitivity—illustrate a market shaped by a complex interplay of forces. A structural model is therefore necessary to disentangle these mechanisms and to quantify their impact on hospital strategy, patient welfare, and the diffusion of innovation.

3 Model

To formally analyze the diffusion of drug-eluting stents (DES), we develop and estimate a structural model of the Taipei metropolitan market. The model is designed to capture the equilibrium interactions among three types of agents: (i) upstream manufacturers, who periodically introduce new DES generations; (ii) downstream hospitals, which strategically manage their DES portfolios and set prices; and (iii) patients, who choose a hospital and a stent for treatment.

Given that all five major stent manufacturers in our sample period are foreign to Taiwan (four are based in the US and one is headquartered in Singapore), their R&D decisions are made in a global context, and the Taiwanese market is relatively small, we treat the arrival of new DES generations as an exogenous process. Our analysis therefore focuses on the downstream market equilibrium. We do not model physicians as separate strategic agents, because they are salaried hospital employees in Taiwan, which aligns their primary incentives with those of the hospitals. The model therefore centers on the two key strategic actors in the downstream market, patients and hospitals, whose interactions are formulated as a sequential game.

3.1 Patients' Demand

On the demand side, we specify a discrete-choice model of a patient's joint selection of a hospital and a stent. Let i index patients, h hospitals, m stent brands, and t quarters. The choice set for a patient in quarter t consists of all hospitals in Taipei ($h \in H$) and the specific stent portfolios, M_{ht}, that they offer. The observable characteristics of each available stent $m \in M_{ht}$ include its patient-paid price, p_{hmt}, its generation, g_{hmt}, and an indicator for whether it is a DES ($isDES_m = 1$) or a BMS ($isDES_m = 0$).[9] To capture the value of innovation highlighted in our stylized facts, we define a variable New_{hmt}, which equals one if stent m is the newest generation of its brand available anywhere in the market in quarter t.[10] The cost of access is measured by dis_{ih}, the straight-line distance between the centroid of the patient's township of residence and the hospital's location.

We model this choice under an assumption of full information. This is justified by two key institutional features: first, the NHI publicly discloses the prices of all DES models at

[9] We denote the BMS option as m=0 without distinguishing between different BMS brands. However, BMS offered by different hospitals is treated as distinct treatment options.

[10] The definition of New is brand-relative for two reasons. First, we lack sufficient information and domain expertise to establish comparable generations across brands. Second, this brand-specific framing aligns with how patients perceive "newness" in the market. Importantly, newness reflects perceived novelty rather than an objective measure of technological improvement or meaningful functional differences.

every hospital, and second, Taiwanese media frequently report on these price differences, reinforcing patient awareness. While we formally model the patient as the decision-maker, we interpret the observed choice as the outcome of a joint decision process between the patient and their physician. Consistent with NHI policy, we assume every hospital offers a BMS ($m = 0$) in every period at zero out-of-pocket cost to the patient ($p_{h0t} = 0$), and we normalize its "newness" attribute to zero ($New_{h0t} = 0$).

Patient i's indirect utility from choosing stent brand m at hospital h in quarter t is specified as:

$$\begin{aligned} u_{ihmt} = \beta_{1i} isDES_m + \beta_{2i} New_{hmt} + \beta_{3i} dis_{ih} + \beta_{4i}(isDES_m \times dis_{ih}) \\ + \beta_{5i} p_{hmt} + \xi_{hmt} + \varepsilon_{ihmt} \equiv \delta_{ihmt} + \varepsilon_{ihmt}, \end{aligned} \tag{1}$$

where δ_{ihmt} represents the mean utility component. The preference parameters β_{ki} are patient-specific, allowing for heterogeneity in tastes for DES status (β_{1i}), newest-generation technology (β_{2i}), travel distance (β_{3i}), the interaction between DES status and distance (β_{4i}), and price (β_{5i}). The term ξ_{hmt} captures unobserved quality attributes of a hospital-brand combination, part of which can be decomposed into hospital and manufacturer fixed effects (ξ_h and ξ_m). The idiosyncratic error term, ε_{ihmt}, is assumed to follow a Type I extreme value distribution. The mean utility from the BMS option ($m = 0$) at hospital h is normalized to the hospital fixed effect, $\delta_{ih0t} = \xi_h$.

To capture rich, observable sources of preference heterogeneity, we allow the random coefficients to vary with patient characteristics:

$$\beta_{ki} = \beta_{k0} + \beta_{k1} charlson_i + \beta_{k2} income_i + \beta_{k3} male_i + \beta_{k4} age_i + v_{ki}, \quad \text{for } k = 1, \ldots, 5.$$

These characteristics include the patient's health status (Charlson Comorbidity Index), the average taxable income of their township of residence, their gender, and their age. The vector $\mathbf{v}_i = (v_{1i}, \ldots, v_{5i})$ represents unobserved individual-specific taste shocks, which are assumed to follow a distribution $G(\cdot)$. Given this mixed logit specification, the probability that patient i chooses stent brand m at hospital h is:

$$s_{ihmt}(\mathcal{P}_t, \mathcal{M}_t) = \frac{\exp(\delta_{ihmt})}{\sum_{h' \in H} \sum_{m' \in M_{h't}} \exp(\delta_{ih'm't})}, \tag{2}$$

where the denominator sums over all available hospital-brand combinations in the market. The market share for each product is then obtained by integrating these individual choice

probabilities over the joint distribution of observed demographics and unobserved preference shocks:

$$s_{hmt}(\mathcal{P}_t, \mathcal{M}_t) = \iint s_{ihmt}\, dF(Income_i, Charlson_i, male_i, age_i)\, dG(\mathbf{v}_i). \tag{3}$$

Price Endogeneity and Instrumental Variables. A primary concern in this specification is that posted prices, p_{hmt}, may be endogenous if they are correlated with hospital-brand attributes captured in the error term ξ_{hmt} that are observed by market participants but not by the econometrician. Examples include the quality of a hospital's catheterization lab, the reputation of its interventional cardiology team, specific delivery-system characteristics of a stent (e.g., crossing profile, trackability), or bundled post-operative care services. These attributes may remain relevant even after we control for hospital and manufacturer fixed effects (ξ_h and ξ_m). Hospitals with more favorable unobserved attributes may charge higher prices, generating a positive correlation between price and ξ_{hmt}.

We address this using a control-function approach. We specify a set of instrumental variables that are correlated with hospital costs but are plausibly orthogonal to unobserved demand shocks. These instruments include: (i) quarterly average exchange rates (NTD/USD and NTD/SGD), which affect the import costs of stents; (ii) historical adoption patterns, such as the number of other hospitals that adopted the same brand prior to the current period, which may proxy for learning or scale effects in procurement; and (iii) measures of historical peer adoption of the latest generation, which capture competitive cost pressures. The control function includes the residual from a first-stage regression of price on these instruments and other exogenous covariates. This residual is then included as an additional regressor in the utility function to purge the endogeneity bias. We carry this residual forward into all subsequent demand computations, supply-side estimation, and counterfactual analyses; for any new combinations of hospital, brand, and time that arise in counterfactuals, we set the residual to zero. This procedure ensures that price endogeneity is controlled for consistently throughout the analysis.

3.2 Supply Model

The supply side of the market is characterized by the strategic behavior of hospitals, which compete by adjusting their DES portfolios and setting prices on a quarterly basis. We model this interaction as a sequential game, which allows us to distinguish between hospitals' long-run investment decisions and their short-run pricing conduct. The game unfolds in three stages:

- **Stage 1 (Portfolio Adjustment)**: Hospitals simultaneously choose their technology portfolios. Each hospital can decide to: (i) adopt a new DES brand not previously offered, (ii) upgrade an existing brand to its latest generation, or (iii) maintain its current portfolio. This is the primary investment decision.
- **Stage 2 (Pricing Competition)**: After observing the portfolio choices of all rivals, hospitals simultaneously set patient-paid prices for all DES models they offer.
- **Stage 3 (Patient Choice)**: Patients observe the available portfolios and prices across all hospitals and make their treatment choices, as described in the demand model.

The model is solved via backward induction, beginning with the pricing stage.

Stage 2: Hospitals' Pricing Decisions

In the second stage, given a fixed configuration of hospital portfolios across the market, $\mathcal{M}_t = \{M_{ht}\}$, each hospital sets prices for the DES models in its portfolio to maximize its current-period variable profit. We assume that hospitals engage in Bertrand-Nash pricing competition. The profit function for hospital h is:

$$\pi_h(M_{ht}, \mathcal{M}_t) = \max_{\{p_{hmt}\}_{m \in M_{ht}}} \left\{ \sum_{m \in M_{ht}, m \neq 0} (p_{hmt} + r_t - c_{hmt})\, s_{hmt}(\mathcal{P}_t, \mathcal{M}_t) + 0.2 r_t s_{h0t}(\mathcal{P}_t, \mathcal{M}_t) \right\}, \tag{4}$$

where s_{hmt} is the market share of product m at hospital h, r_t is the fixed NHI reimbursement for a stent procedure, and c_{hmt} is the hospital's marginal cost for DES model m. During our sample period, the NHI lowered r_t twice, which we incorporate as raw data. The term $(p_{hmt} + r_t)$ represents the total revenue per DES procedure. For the BMS option ($m = 0$), which is free to patients, we assume hospitals earn a 20% margin on the reimbursement payment. This assumption reflects the NHI's own pricing mechanism: the NHI determines stent reimbursement rates through periodic market transaction price surveys, setting the payment equal to the volume-weighted average procurement price across hospitals plus a "reasonable zone" markup of approximately 20% to cover handling, storage, and administrative costs (Hsu and Lu, 2015). When the NHI reduced the base stent reimbursement in 2009 and 2012, it did so because the market survey revealed that wholesale BMS prices had fallen; accordingly, both the reimbursement rate and the underlying wholesale prices declined in tandem,

preserving the approximate 20% margin by design.[11] The first-order conditions from this pricing problem define a system of equations for all hospitals. By inverting this system, we can recover the implied marginal cost for each DES model offered in the market:

$$\mathbf{c}_t = \mathbf{r}_t + \mathbf{p}_t + (\Delta_{-0t})^{-1} (0.2\mathbf{r}_t\Delta_{0t} + \mathbf{s}_t), \tag{5}$$

where Δ_{-0t} is the matrix of own- and cross-price derivatives of DES shares with respect to DES prices, and Δ_{0t} is the vector of derivatives of BMS shares with respect to DES prices. These recovered marginal costs are crucial inputs for the subsequent analysis, as they allow us to compute the expected profits that drive hospitals' portfolio decisions in the first stage.

Stage 1: Hospitals' Portfolio Adjustment Decisions

In the first stage of the game, each hospital chooses its technology portfolio to maximize its total expected payoff. Based on our empirical observation that hospitals rarely make more than one portfolio change in a given quarter, we simplify the choice set: a hospital can make at most one adjustment per period, either adopting one new brand, upgrading one existing brand, or making no change. The value to hospital h of choosing a new portfolio M'_{ht}, given its current portfolio M_{ht} and the choices of its rivals $\mathcal{M}'_t$, is:

$$V_{ht}(M'_{ht}, M_{ht}, \mathcal{M}'_t) = \pi_h(M'_{ht}, \mathcal{M}'_t) - f^M_{ht}(M_{ht}, M'_{ht}), \tag{6}$$

where $\pi_h(\cdot)$ is the expected variable profit from the second-stage pricing equilibrium and $f^M_{ht}(\cdot)$ is the one-time fixed cost of the portfolio change (either an adoption or an upgrade). The hospital also receives a private payoff shock v_{ht}, assumed to be i.i.d. Type I extreme value, which enters additively but is omitted from V_{ht} for notational clarity.

We parameterize the fixed cost separately for adoptions (adding a brand not previously offered) and upgrades (replacing an existing brand with its latest generation). For an adoption event,

$$\begin{aligned} f^{\text{adopt}}_{ht} = {} & \gamma_0 + \gamma_1 \, \mathit{PortfolioSize}_{ht-1} + \gamma_2 \, \mathit{NumNewest}_{ht-1} \\ & + \gamma_3 \, \mathit{FracSameBrand}_{-h,t} + \gamma_4 \, \mathit{FracLatestSameBrand}_{-h,t} + \alpha_b + \alpha_y + \alpha_h, \end{aligned} \tag{7}$$

[11] This mechanism is directly analogous to the "drug price gap" (yao-chia-cha) adjustment process in Taiwan's pharmaceutical pricing system, where the NHI periodically surveys actual transaction prices between hospitals and manufacturers and adjusts the reimbursement rate to maintain a target margin (Hsu and Lu, 2015). Because BMS are free to patients, there is no patient-level price variation from which to recover BMS wholesale costs structurally; the 20% margin assumption is the most natural alternative given the institutional design.

where $PortfolioSize_{ht-1}$ is the number of DES brands in hospital h's portfolio at the end of the previous period, $NumNewest_{ht-1}$ counts how many of those brands are at the latest generation, $FracSameBrand_{-h,t}$ is the fraction of other hospitals already carrying the focal brand, $FracLatestSameBrand_{-h,t}$ is the fraction carrying its latest generation, and α_b, α_y, α_h denote brand, year, and hospital fixed effects. Intuitively, a larger existing portfolio lowers the incremental cost of adding another brand (economies of scope), while already carrying many frontier products raises the complexity of further expansion (diseconomies at the technological frontier). The peer-adoption variables allow for social learning or competitive pressure: hospitals may find it easier (or harder) to adopt a brand that rival hospitals have already introduced.

For an upgrade event, the fixed cost takes a similar form with additional time-since-introduction controls:

$$\begin{aligned} f_{ht}^{\text{upgrade}} = \delta_0 &+ \delta_1 \, PortfolioSize_{ht-1} + \delta_2 \, NumNewest_{ht-1} \\ &+ \delta_3 \, FracSameBrand_{-h,t} + \delta_4 \, FracLatestSameBrand_{-h,t} \\ &+ \delta_5 \, Duration_{ht} + \delta_6 \, NewDuration_t + \delta_7 \, NewDuration_t^2 + \alpha_b + \alpha_y + \alpha_h, \end{aligned} \tag{8}$$

where $Duration_{ht}$ measures how long hospital h has carried the current generation, $NewDuration_t$ captures the time since the newest generation first became available in the market (with a quadratic term to allow for non-linear effects), and the remaining variables are defined as above. The duration terms capture the idea that the urgency and cost of upgrading may depend on how "stale" the hospital's current technology is and how established the new generation has become in the broader market.

We note that the absolute magnitudes of f_{ht}^{adopt} and f_{ht}^{upgrade} should be interpreted relative to the scale of the unobserved profit shock v_{ht}. Because the T1EV distribution normalizes the variance of v_{ht}, the estimated cost parameters are identified only up to this scale. In particular, the dollar values implied by the estimates should not be read as literal NTD amounts, but rather as costs measured in units of the unobserved shock. This point is relevant for interpreting the magnitudes reported in Table 3.

We do not impose a non-negativity constraint on the implied fixed costs. In preliminary specifications that restricted $f_{ht} \geq 0$, the model fit deteriorated substantially, as the constraint prevented the cost function from capturing combinations of hospital, brand, and time where the net institutional adjustment cost is negative. A negative implied cost arises naturally when a portfolio change yields ancillary benefits, such as procurement leverage from consolidating volume with a single supplier, favorable consignment terms offered by manufacturers seeking market entry, or administrative savings from retiring a low-volume

legacy product, that exceed the direct transition costs enumerated above. The unconstrained specification therefore provides a more flexible and empirically accurate characterization of the cost distribution, while the estimated median and interquartile range remain positive and economically interpretable.

Although portfolio decisions are inherently dynamic, we model this stage as a static game in which hospitals maximize current-period profits net of adjustment costs. This simplification is motivated by three considerations. First, as our descriptive evidence shows, most adoption and upgrading events occur relatively quickly, typically within one to three quarters of a new product's market availability, which limits the salience of long-term strategic waiting. Second, the set of competing hospitals in our sample is stable, mitigating concerns about strategic portfolio choice aimed at deterring entry. Third, a fully dynamic model with this many products and firms would be computationally infeasible, as the state space of all possible portfolio combinations is immense.[12] Our static approach is consistent with recent work analyzing complex product-portfolio problems, such as Wollmann (2018) and Demirer and Olssen (2024).

The probability that hospital h chooses portfolio M'_{ht} follows a standard logit form, which we use as the basis for a maximum likelihood estimator of the parameters of the adoption and upgrade cost functions.

$$Pr_t(M'_{ht}|M_{ht}, \mathcal{M}_t) = \frac{\exp\left(V_{ht}(M'_{ht}, M_{ht}, \mathcal{M}_t)\right)}{\sum_{M' \in \mathcal{A}_{ht} \cup \{M_{ht}\}} \exp\left(V_{ht}(M', M_{ht}, \mathcal{M}_t)\right)}, \tag{9}$$

where $\mathcal{A}_{ht}$ is the set of all feasible single portfolio adjustments (adoptions or upgrades) for hospital h at time t.

The logit choice probability above can be interpreted as the equilibrium choice probability of a simultaneous portfolio game with private hospital-specific payoff shocks. The T1EV assumption smooths the discrete portfolio decision and yields strictly positive probabilities for each feasible adjustment, a standard approach in empirical models of entry and product portfolio choice (Head and Mayer, 2019; Wollmann, 2018; Demirer and Olssen, 2024). Additional equilibrium details are provided in the Online Appendix.

4 Estimation Results

We begin with demand-side estimates, which reveal rich patterns of patient heterogeneity and provide crucial insights into how preferences for innovation and price sensitivity shape

[12]For example, if each of the 20 hospitals can choose whether or not to carry each of the 5 major DES brands, there are $(2^5)^{20}$ possible portfolio states, even before accounting for generations.

market outcomes. We then turn to the supply side, showing how hospitals' strategic decisions regarding technology adoption and pricing respond to demand and competition.

4.1 Demand Estimation

We estimate the demand model using detailed, individual-level data on patient treatment choices. For each patient in a given quarter, the choice set encompasses all available DES and BMS options across every hospital in the Taipei area. Specified as a mixed logit, the model allows for rich, unobserved heterogeneity in patient preferences alongside heterogeneity based on observable patient characteristics. We estimate the model via simulated maximum likelihood. Acknowledging that out-of-pocket prices can be endogenous due to unobserved product-hospital quality attributes, we employ a control-function approach using the instrumental variables detailed in Section 3. An important advantage of Taiwan's single-payer NHI system for demand estimation is that all patients are enrolled in the same mandatory insurance plan with identical coverage rules. Unlike the U.S. setting, where patients may strategically select insurance plans based on anticipated medical needs, introducing a well-known source of selection bias, patients in our sample face a uniform coverage structure that is orthogonal to their stent choices. This eliminates insurance-plan selection as a confound and strengthens the identification of price sensitivity from observed out-of-pocket variation.

The identification of the random coefficients in our mixed logit specification relies on two distinct sources of variation. The four observable interaction terms, Price $\times$ Income, DES $\times$ Charlson, DES $\times$ Age, and Newest Gen $\times$ Income, are identified from within-market demographic variation in choices. Within any given hospital-quarter market, patients with different demographic profiles face the *same* set of product attributes (prices, generations, brand availability) but make systematically different choices. For example, the Price $\times$ Income interaction is identified because higher- and lower-income patients in the same market make different stent choices despite facing identical prices. With 61,645 patient encounters across 20 hospitals and 28 quarters, there is substantial within-market demographic variation to pin down these interaction terms. The unobservable random coefficient on price is identified from the pattern of substitution across products, following the standard logic of Berry et al. (1995) and Berry and Haile (2014). In a logit model without random coefficients, the substitution pattern is fully determined by market shares (the IIA property). Deviations from IIA in the observed substitution patterns, for instance, if a price increase for one DES brand disproportionately shifts demand to other DES brands rather than to BMS, identify the variance of the unobserved taste heterogeneity. The two NHI reimbursement cuts (2009

and 2012), which shifted DES prices but not BMS prices, provide particularly informative variation for tracing out substitution patterns that discipline this random coefficient.

Table 2 presents the demand estimates, moving from a standard conditional logit to our preferred mixed logit specification with a full set of interactions. The superior fit and refined insights from the latter specification confirm the need to model both observed and unobserved heterogeneity.

The estimation results for our preferred specification reveal two primary determinants of patient choice. First, patients are sensitive to out-of-pocket costs, as indicated by the negative and statistically significant mean price coefficient. This price sensitivity represents a significant friction in the adoption of new technologies. Second, patients exhibit a strong preference for clinical novelty, evidenced by the large, positive, and significant coefficient on "newest-generation". This preference for new technology creates a demand-side incentive for hospitals to update their stent portfolios.

The model further reveals substantial heterogeneity in patient choices. The estimated interaction terms show that price sensitivity is significantly attenuated for higher-income patients, while the preference for DES over BMS increases with patient income and age. More notably, two sets of interaction coefficients involving the Charlson Comorbidity Index, an index of disease severity, warrant careful interpretation: both DES $\times$ Charlson (-0.109) and Price $\times$ Charlson (-0.004) are negative and significant.

These coefficients are best understood jointly as reflecting *clinical indication heterogeneity* rather than pure preference heterogeneity. Patients with higher Charlson scores are more likely to have contraindications for DES, primarily because DES requires prolonged dual antiplatelet therapy (DAPT, typically 6 to 12 months) to prevent stent thrombosis, and patients with multiple comorbidities, such as those at elevated bleeding risk, with hepatic or renal impairment, or facing imminent non-cardiac surgery, are often clinically directed toward BMS, which requires only 1 month of DAPT (Valgimigli et al., 2018). Since BMS is fully covered by NHI and carries zero out-of-pocket cost, the clinical channeling of high-Charlson patients toward BMS mechanically generates a negative DES $\times$ Charlson coefficient (these patients are less likely to choose DES for clinical reasons) and simultaneously amplifies the negative Price $\times$ Charlson coefficient (these patients disproportionately choose the zero-price option). We therefore caution against interpreting the Price $\times$ Charlson coefficient as evidence that sicker patients are intrinsically more price-sensitive; rather, it partially reflects the same underlying clinical selection that drives the DES $\times$ Charlson result.[13]

[13]While we cannot fully decompose the Price $\times$ Charlson coefficient into a "clinical channeling" component and a "true price sensitivity" component without exogenous variation in clinical guidelines, the institutional fact that BMS carries zero patient cost under Taiwan's top-up system means that any clinical preference for BMS is observationally equivalent to heightened price sensitivity in the estimated model.

Table 2: Demand Estimation Results

Variable	(1) Clogit 1	(2) Clogit 2	(3) Mixlogit 1	(4) Mixlogit 2
Price Effects				
Price	-0.009***(0.001)	-0.026***(0.004)	-0.036***(0.002)	-0.025***(0.005)
Price × Charlson		-0.004***(0.001)		-0.004***(0.001)
Price × Income		0.005***(0.002)		0.004*(0.002)
Price × Male		-0.004***(0.001)		-0.004**(0.002)
Price × Age		-0.001***(0.000)		-0.001***(0.000)
Newest Generation Effects				
Newest Generation	0.318***	0.659***(0.142)	0.470***(0.021)	0.636***(0.143)
Newest Gen × Charlson		-0.027(0.026)		-0.029(0.026)
Newest Gen × Income		-0.188***(0.061)		-0.175***(0.063)
Newest Gen × Male		0.038(0.050)		0.040(0.050)
Newest Gen × Age		-0.000(0.002)		0.001(0.002)
Distance Effects				
Distance	-0.072***(0.000)	-0.061***(0.005)	-0.105***(0.001)	-0.088***(0.005)
Distance × Charlson		-0.006***(0.001)		-0.007***(0.001)
Distance × Income		-0.026***(0.003)		-0.012***(0.003)
Distance × Male		-0.002(0.001)		-0.002(0.002)
Distance × Age		-0.000(0.000)		0.000(0.000)
DES Indicator Effects				
DES Indicator	-1.163***(0.056)	-0.232(0.262)	0.620***(0.097)	-0.297(0.295)
DES × Charlson		-0.106**(0.043)		-0.109**(0.049)
DES × Income		0.462***(0.108)		0.538***(0.125)
DES × Male		0.097(0.086)		0.090(0.097)
DES × Age		0.057**(0.003)		0.057*(0.003)
DES Interaction with Distance				
DES × Distance	0.002***(0.000)	-0.003**(0.001)	0.001***(0.000)	-0.004***(0.001)
DES × Dist × Charlson		0.000*(0.000)		0.000*(0.000)
DES × Dist × Income		0.005***(0.001)		0.006***(0.001)
DES × Dist × Male		0.001***(0.000)		0.001***(0.000)
DES × Dist × Age		0.000(0.000)		0.000(0.000)
Control Function				
Control Function Residual	0.010(0.020)	0.037***(0.002)	0.037***(0.002)	0.037***(0.002)
Standard Deviations of Random Coefficients				
SD Price			-0.000(0.002)	-0.000(0.002)
SD DES Indicator			0.100(0.148)	0.124(0.171)
SD Newest Generation			-0.141(0.142)	-0.160(0.147)
SD Distance			0.073***(0.001)	0.072***(0.001)
Hospital FE	Yes	Yes	Yes	Yes
Manufacturer FE	Yes	Yes	Yes	Yes
N	5579861	5579861	5579861	5579861
Log-likelihood	-2.45e+05	-2.28e+05	-2.29e+05	-2.28e+05

Notes: Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Separately, the interaction of newest-generation with income (Newest Gen $\times$ Income $= -0.175$) is negative and significant, indicating that higher-income patients place less additional value on the "newest" label conditional on choosing DES. A plausible explanation is informational: higher-income patients (and the physicians at the hospitals they frequent) may be better informed about the modest incremental clinical improvements between successive DES generations, making them less responsive to the "newest" marketing signal. This is distinct from the Charlson-driven patterns above and represents genuine preference heterogeneity along the income dimension.

Finally, additional estimates support the validity of our model specification. The coefficient on the control-function residual is positive and highly significant, confirming the presence of price endogeneity and the necessity of our instrumental variable strategy. The coefficient on travel distance is negative and significant, consistent with geographic competition. The disutility from travel is larger for sicker and richer patients, likely reflecting higher opportunity costs. These ancillary results are consistent with economic theory and support the robustness of our main findings.

Implications for Innovation Diffusion. The demand estimates reveal two key insights for innovation diffusion. First, our estimated individual price elasticity is mainly within the -3 to -1 range, while the mode is around -2. Such pronounced price responsiveness means that high out-of-pocket costs can limit adoption even when patients value innovation. Second, we estimate that patients' willingness to pay (WTP) for the newest-generation attribute is approximately 5.69 thousand NTD, nearly 10% of the average DES price. The strong preference for newest-generation DES creates incentives for hospitals to upgrade, but the distribution of these preferences suggests that the benefits of innovation may not be equally accessible to all patients. The Online Appendix reports the elasticity calculations and willingness-to-pay calculations.

The gap between the marginal WTP for the newest-generation DES (approximately 5,690 NTD) and the typical out-of-pocket DES price (50,000 to 80,000 NTD) is striking, but it does not imply that patients undervalue DES. Rather, the observed sorting between BMS and DES reflects two channels. Clinically, patients with greater comorbidity burden are more likely to be directed toward BMS because DES requires prolonged dual antiplatelet therapy. Financially, the out-of-pocket cost of approximately 60,000 NTD, roughly one month of average household income in our sample, is a substantial barrier for lower-income patients. The patient coupon counterfactual in Section 5.2 directly addresses this equity dimension.

Finally, we note that the demand estimates are robust to alternative assumptions about the supply side. The identification of patient preferences requires only that patients observe posted prices and choose to maximize utility. It does not depend on the mechanism that

generates those prices, whether strategic markup optimization under Bertrand-Nash competition, inventory clearing under fixed manufacturer allocations, or any other pricing rule. Thus, even if one holds a different view of the supply-side model, the demand-side results, including price elasticities, income-based heterogeneity in WTP, and substitution patterns between BMS and DES, remain valid. We present institutional evidence supporting active hospital portfolio choice in Section 2.3.

4.2 Supply Estimation

Supply-side estimation proceeds in two steps. First, we recover marginal costs at the hospital, brand, and quarter level by inverting the Bertrand-Nash first-order conditions using demand estimates. Second, we estimate fixed costs of portfolio adjustment, including adoption of new brands and upgrades of existing lines, using a discrete-choice framework fit to observed portfolio changes. These primitives discipline the counterfactuals on pricing and diffusion.

This two-step approach is standard in empirical IO studies of differentiated products. The inversion step uses the estimated demand elasticities to map observed prices and shares into implied marginal costs under the assumed pricing game, thereby separating preference-driven demand variation from cost-driven supply behavior. The second step models portfolio modification as adjustment over time along the extensive margin. While we do not fully solve a dynamic program, the discrete-choice specification captures state dependence via observed portfolio composition and fixed effects and provides a tractable representation of adjustment costs that is sufficient for the counterfactual exercises below.

Recovering Marginal Costs. Using the preferred demand estimates, we recover marginal costs by inverting the first-order conditions from hospitals' pricing decisions. Under Bertrand-Nash competition, these conditions imply

$$\mathbf{c}_t = \mathbf{r}_t + \mathbf{p}_t + (\Delta_{-0t})^{-1}\Big(0.2\,\mathbf{r}_t\,\Delta_{0t} + \mathbf{s}_t\Big),$$

where $\mathbf{p}_t$ stacks DES prices at the hospital and brand level, $\mathbf{r}_t$ is the reimbursement vector, and $\mathbf{s}_t$ denotes DES market shares. The matrix Δ_{-0t} contains own- and cross-price derivatives of DES shares with respect to DES prices, while Δ_{0t} collects the derivatives of BMS shares with respect to DES prices, both evaluated at observed prices and demand parameters. We compute these derivatives from demand estimates using simulation draws and the control-function residuals retained from demand.

Operationally, we construct Δ_{-0t} and Δ_{0t} at the hospital and brand level each quarter, invert Δ_{-0t} using a numerically stable routine with Tikhonov regularization for near-singular systems, and obtain implied marginal costs $\mathbf{c}_t$. We drop cells with predicted shares below

10^{-5} to avoid numerical artifacts.[14] For combinations of hospital, brand, and quarter not observed in the data but required for counterfactuals, we impute marginal costs via a linear model with hospital, brand, and quarter fixed effects fit on the recovered $\mathbf{c}_t$.

Recovered costs indicate modest cross-manufacturer dispersion around 50,000 NTD per unit (Cordis highest at 58,000; "Other" lowest at 45,000). Costs decline from roughly 55,000 NTD in 2008 to 45,000 NTD in 2012, consistent with the 2009 reimbursement cut and intensified entry. Major teaching hospitals face lower marginal costs (48,000 NTD) than minor teaching hospitals (50,000 NTD), consistent with scale and bargaining power in procurement.

Adoption and Upgrade Cost Estimation. The second step of our supply-side estimation quantifies the fixed costs that hospitals incur when adjusting their DES portfolios. We model the hospital's decision to adopt a new brand or upgrade an existing one as a discrete choice, where the cost of each action is a function of hospital characteristics, current portfolio composition, and broader market conditions. Because it is rare to observe any hospital engaging in multiple adoption or upgrade actions in a quarter, we limit at most one action when a hospital adjusts its DES portfolio in a quarter. The action options available to a hospital-quarter depends on the brand-generation available in the market, and the hospital's existing portfolio. The model includes an unobserved cost component assumed to follow a Type I extreme value distribution.

Formally, the log-likelihood for the observed sequence of portfolio choices across all hospitals and quarters is

$$\mathcal{L}(\boldsymbol{\gamma}, \boldsymbol{\delta}, \boldsymbol{\alpha}) = \sum_{h=1}^{H} \sum_{t=1}^{T} \ln Pr_t\big(M_{ht}^{\text{obs}} \;\big|\; M_{h,t-1}, \mathcal{M}_{-h,t}^{\text{obs}}\big)\,, \tag{10}$$

where M_{ht}^{obs} is hospital h's observed portfolio in quarter t, $M_{h,t-1}$ is its portfolio in the preceding quarter, and $\mathcal{M}_{-h,t}^{\text{obs}}$ denotes the observed portfolios of all rival hospitals. The choice probability $Pr_t(\cdot)$ is given by Eq. (9), with the deterministic payoff V_{ht} computed using the recovered variable profits $\pi_h(\cdot)$ from Stage 2 and the parameterized fixed costs f_{ht}^{adopt} or f_{ht}^{upgrade} from Eqs. (7) and (8). Note that we condition on rivals' *observed* portfolio choices rather than their predicted choice probabilities, treating the estimation as a conditional (partial) likelihood. This is consistent with our Bayesian Nash interpretation: each hospital optimizes against the realized distribution of rival portfolios, and the likelihood aggregates these conditional choice probabilities across the panel.

We maximize this likelihood with respect to the full parameter vector $(\boldsymbol{\gamma}, \boldsymbol{\delta}, \boldsymbol{\alpha})$, which includes the adoption and upgrade cost coefficients and the brand, year, and hospital fixed

[14] Results are insensitive to thresholds between 10^{-6} and 10^{-4}.

effects. This approach allows for a flexible specification that includes a rich set of fixed effects to control for unobserved, time-invariant heterogeneity at the brand, year, and hospital levels. As a robustness check on the importance of these controls, we evaluated a sequence of specifications: a model with only the 11 core covariates yields a log-likelihood of -228.643; successively adding brand, year, and finally hospital fixed effects improves the fit to -215, -195, and -170, respectively. The material gains at each step confirm that unobserved heterogeneity is a key feature of the data and that its inclusion is critical for obtaining reliable cost estimates. While alternative estimation approaches for portfolio-adjustment models exist, such as moment-inequality estimators, they are computationally infeasible in our setting. With 73 parameters in our full specification, evaluating the necessary inequality conditions would require exploring a parameter space of intractable dimensionality (e.g., 10^{73} combinations for a coarse 10-point grid search). Our MLE approach is therefore the most viable strategy that accommodates the institutional richness of the market.

Table 3: Adoption and Upgrade Cost Estimates

	Specification 1		Specification 2	
Variable	Est.	Std. Err.	Est.	Std. Err.
Adoption Cost				
Portfolio size	-9.213	1.3269	-9.3706	1.3306
No. of hospitals using same brand/20			7.3858	5.8832
No. of hospitals using the latest same brand/20			-7.6596	4.1622
No. of latest generation within own portfolio	4.5766	1.1429	4.4923	1.0971
Upgrade Cost				
Portfolio size	-2.5913	0.6946	-2.6266	0.7244
No. of hospitals using same brand/20			4.4148	7.8712
No. of hospitals using the latest same brand/20			3.1591	2.4551
No. of latest generation within own portfolio	2.5763	0.4771	2.5751	0.4918
Duration of current generation			-0.0069	0.0732
Duration of newest gen since introduction			-0.8143	0.4963
Duration of newest gen since introduction, sq			0.0498	0.0302
Brand FE	Yes		Yes	
Year FE	Yes		Yes	
Hospital FE	Yes		Yes	
Likelihood	-174.843		-170.373	

Table 3 reports the parameter estimates from our preferred specification with a full set of fixed effects. The results indicate the presence of both economies and diseconomies of scope. The negative and significant coefficient on portfolio size in both the adoption and upgrade

cost functions (-9.371 and -2.627, respectively) implies substantial economies of scope: hospitals with larger existing portfolios face lower incremental costs for both adding new product lines and upgrading existing ones. In contrast, the positive and significant coefficient on the number of newest-generation stents already in the portfolio (4.492 for adoption; 2.575 for upgrade) points to diseconomies of scope as a hospital's portfolio becomes more technologically advanced. To accommodate potential social learning between hospitals, we allow a hospital's fixed costs of adoption or upgrade to depend on the fraction of other hospitals using the same brand or the latest generation of the same brand up to the study quarter. None of these social learning variables turns out significant, confirming our understanding that social learning and other strategic dynamics between competing hospitals do not play an important role in hospitals' portfolio adjustment.

To translate these parameters into economically meaningful magnitudes, we compute the implied distribution of fixed costs. The median fixed cost for adopting a new brand is estimated to be 5.85 million NTD, with an interquartile range of 0.90 to 8.96 million NTD. The cost to upgrade an existing product line is considerably lower, with a median of 0.83 million NTD and an interquartile range of -1.38 to 2.69 million NTD. The negative lower quartile indicates that roughly a quarter of feasible upgrade opportunities carry negative implied costs, cases where the ancillary benefits of the transition (e.g., improved consignment terms, volume consolidation, or retirement of a low-demand legacy product) exceed the direct adjustment costs. This significant cost differential between adoption and upgrade provides a clear economic explanation for the empirical pattern wherein hospitals are more inclined to incrementally upgrade their existing technology portfolios than to adopt products from entirely new manufacturers.

These magnitudes are plausible relative to the DES profits that finance portfolio changes: the median upgrade cost is roughly half of a hospital-brand quarter's median DES profit, whereas the median adoption cost corresponds to several quarters of such profit. The fixed cost should also be interpreted as the total hospital-level adjustment cost, including device committee review, contract renegotiation, inventory transition, NHI billing updates, and physician or staff familiarization, rather than only a physician learning cost. The Online Appendix further shows that hospitals that actually adopted or upgraded faced much lower implied costs than hospitals with feasible opportunities that chose not to adjust.

Figure 4 displays the full distribution of implied fixed costs separately for adoption and upgrade events. The distributions confirm two key features of the cost structure: adoption costs are systematically higher than upgrade costs, and both distributions exhibit substantial dispersion across hospital-brand-quarter opportunities.

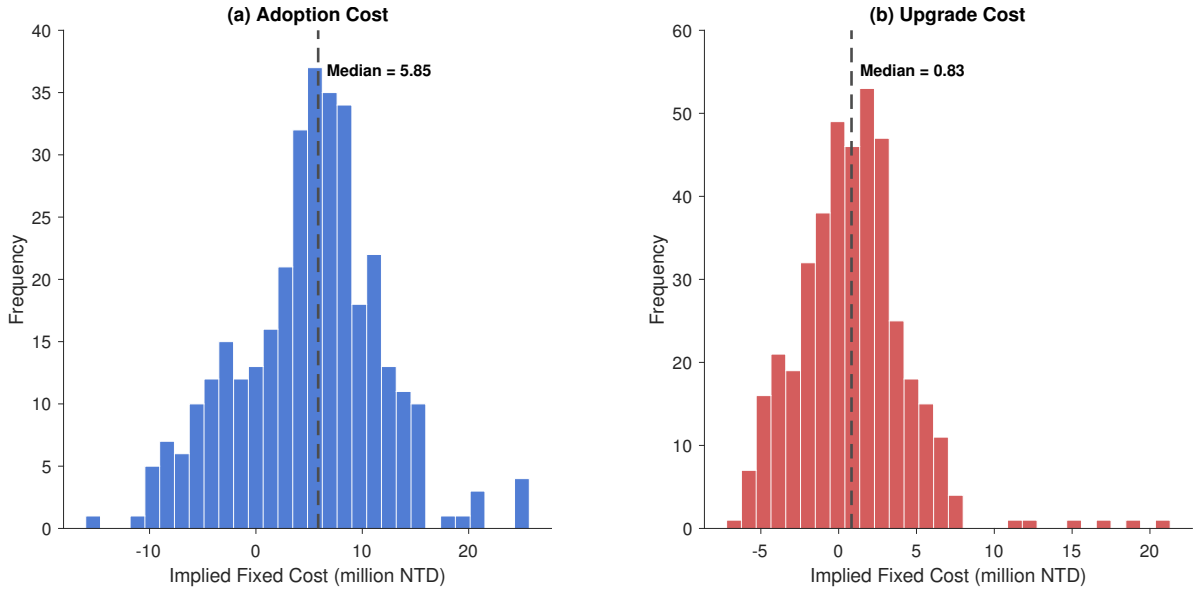


Figure 4: Distribution of Implied Fixed Costs

Note: The figure shows the distribution of implied fixed costs for adoption (left panel) and upgrade (right panel) events, computed using the parameter estimates from Specification 2 of Table 3. Each observation corresponds to a feasible portfolio adjustment at the hospital, brand, and quarter level. Vertical dashed lines indicate the median. Negative values indicate combinations of hospital, brand, and quarter where ancillary benefits of the portfolio change exceed direct transition costs.

The recovered marginal costs in Figure 5 provide a final view of the underlying cost structure of the market. The distribution by manufacturer centers around approximately 50,000 NTD, while the temporal trend shows a decline from approximately 55,000 NTD in 2008 to 45,000 NTD in 2012. The cost differential between major and minor teaching hospitals is consistent with scale or bargaining advantages, providing an economic rationale for the observation that larger institutions are often the earliest adopters of new technologies.

These estimates identify the central forces used in the counterfactuals: patients value newest-generation technology but remain price sensitive, while hospitals face substantially higher costs for adding a new brand than for upgrading an existing one. The Online Appendix reports the model-fit exercises and the comparison of actual adopters with non-adopters.

5 Counterfactual Simulations

The continuous introduction of new DES models by upstream manufacturers presents a fundamental economic challenge: can market mechanisms provide sufficient incentives for downstream adoption while maintaining affordable patient access? To answer this question,

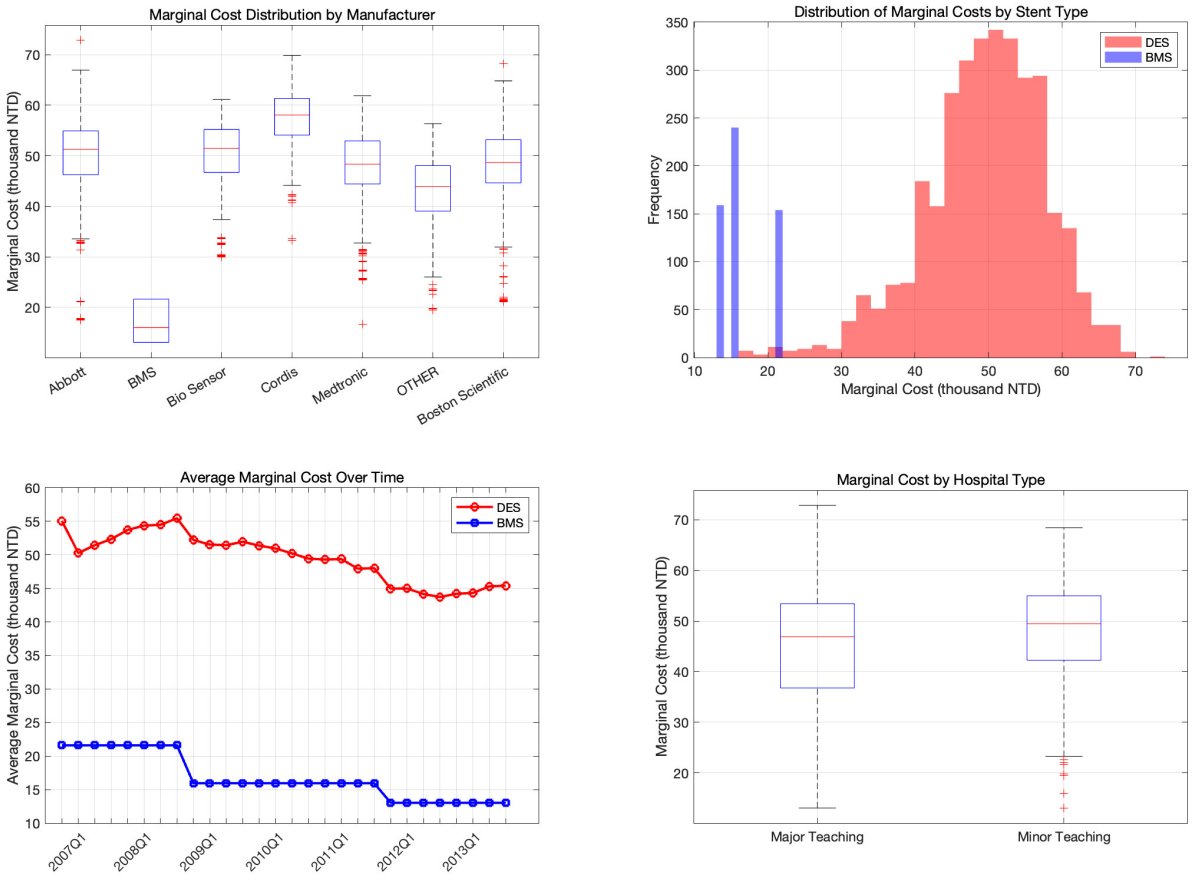


Figure 5: Summary of Recovered Marginal Costs

our counterfactual simulations leverage the structural estimates to decompose the diffusion process and quantify welfare implications. This analysis reveals how the interaction between market structure and policy design determines both the pace of technological progress and its distributional consequences.

We proceed in two complementary dimensions. First, we examine how varying degrees of market concentration interact with heterogeneous patient preferences for innovation to shape equilibrium outcomes, revealing a fundamental tension between competitive pricing and investment incentives. Second, we evaluate alternative reimbursement mechanisms, ranging from simple reimbursement rate change to supply-side selective contracting and demand-side patient subsidies. When we quantify their differential impacts, we also highlight how the efficacy of these policies vary by market structure.

5.1 Market Competition and Patient Demand

Before turning to specific policy designs, we use the structural model to decompose the forces shaping the market equilibrium we observe. The exercises in this subsection are not intended to replicate realistic policy scenarios; rather, they isolate and quantify, one dimension at a time, how supply-side factors (market structure) and demand-side factors (patient preferences for innovation) contribute to hospitals' joint decisions on pricing and portfolio composition. This decomposition establishes a baseline understanding of the mechanisms at work, which then disciplines the interpretation of the policy counterfactuals that follow. We note that these market structure variations also have indirect empirical relevance: Taiwan's hospital market exhibits substantial cross-regional variation in concentration, ranging from one or two PCI-capable hospitals in rural eastern counties to over a dozen competitors in southern metropolitan areas. Our $N \in \{2, 10, 20\}$ scenarios thus span the range of competitive environments that exist across Taiwan's healthcare landscape.

The canonical industrial organization question regarding the competition-innovation nexus (Aghion et al., 2005) takes on a distinct character in markets where innovation diffusion operates through intermediaries. In our setting, hospitals function as gatekeepers whose portfolio choices determine both the availability and pricing of new technologies. This intermediation creates a fundamental tension: while competition enhances patient access through lower prices and expanded aggregate choice, it simultaneously erodes the quasi-rents that incentivize individual hospitals to adopt costly innovations. Furthermore, the transmission of patient preferences into market outcomes is not uniform; as we show, market structure mediates how demand shocks are distributed, determining whether changes in valuation are passed through to consumers or captured by intermediaries.

Market Structure and Equilibrium Outcomes. To quantify the relationship between market concentration and technology diffusion, we simulate equilibria under three competitive environments: a duopoly ($N = 2$), an intermediate case ($N = 10$), and the baseline market observed in our data ($N = 20$). For scenarios with reduced hospital counts, we randomly select N hospitals and average outcomes across ten draws; computational details are reported in the Online Appendix.

Table 4 shows that competition improves access primarily through prices. Moving from duopoly to the observed $N = 20$ market reduces the average patient-paid DES price from 70.13 to 57.10 thousand NTD and nearly triples consumer surplus, while per-hospital profit falls from 21.30 to 2.72 million NTD. Total social surplus therefore rises monotonically with competition, indicating that the consumer gains from lower prices dominate the loss of producer rents.

Table 4: Counterfactual 1: Market Competition Effects on DES Market Outcomes

Measure	N=2	N=10	N=20
Panel A: Quarterly Market Outcomes			
Consumer surplus (million NTD)	60.81	138.00	175.84
Average patient-paid DES price (thousand NTD)	70.13	59.83	57.10
Hospital profit (per hospital, million NTD)	21.30	4.59	2.72
DES price elasticity	−2.40	−2.19	−2.13
Government subsidy (million NTD)	43.97	43.97	43.97
Subsidy allocated to DES	16.38	19.69	20.79
Social surplus (CS + profit - subsidy) (million NTD)	59.45	139.90	186.29
DES utilization rate (%)	38.08	45.92	48.26
Panel B: Technology Adoption and Patient Access			
Portfolio modification probability per hospital:			
Overall (%)	43.69	35.61	36.34
Adoption (%)	27.01	25.86	25.75
Upgrade (%)	33.92	23.36	24.23
Hospital-level DES combinations	4.32	4.20	4.25
With newest generation	3.58	3.39	3.41
Proportion of newest-generation options (%)	82.87	80.68	80.42
Patients receiving newest-generation DES (%)	34.48	41.71	44.43

Competition also changes the composition of innovation access. The portfolio modification probability falls from 43.7% in duopoly to roughly 36% in more competitive markets, consistent with weaker innovation rents. Yet patient access to newest-generation DES rises from 34.5% to 44.4%, because lower prices and a larger aggregate choice set more than offset the decline in the average hospital's portfolio modification. This pattern summarizes the paper's central mechanism: competition weakens individual hospitals' adoption incentives but improves aggregate access to frontier technology. Additional mechanism plots in the Online Appendix decompose these price, profit, utilization, and portfolio-adjustment channels.

The Role of Patient Preferences for Innovation. To isolate demand-side innovation pull, we scale the estimated preference parameter for newest-generation technology (the coefficient β_{2i} on the New_{hmt} attribute defined in Section 3.1) by a factor ζ. Setting $\zeta = 2$ doubles patients' marginal utility from the newest generation of each brand relative to its previous generation, while $\zeta = 0.5$ halves it relative to the baseline.

Table 5 shows that stronger patient preferences for novelty accelerate diffusion in all market structures. Doubling ζ raises newest-generation DES utilization from 44.4% to 55.6% in the observed $N = 20$ market, while halving it lowers utilization to 39.1%. Market structure

Table 5: Market Outcomes Under Alternative Patient Newness Preferences

	N=2 Hospitals			N=10 Hospitals			N=20 Hospitals		
Newness Preference Factor (ζ)	0.5	1 (Baseline)	2	0.5	1 (Baseline)	2	0.5	1 (Baseline)	2
Panel A: Patient and Product Outcomes									
Consumer Surplus (million NTD)	57.06	60.81	69.63	132.69	138.00	151.14	169.99	175.84	189.99
DES Utilization Rate (%)	34.37%	38.08%	45.86%	41.25%	45.92%	55.86%	43.29%	48.26%	58.54%
Newest-Generation DES Utilization Rate (%)	30.27%	34.48%	43.66%	36.69%	41.71%	52.89%	39.10%	44.43%	55.62%
Avg. DES Models Available per Hospital	4.30	4.32	4.33	4.20	4.20	4.20	4.25	4.25	4.24
Avg. Newest-Gen. Models per Hospital	3.54	3.58	3.77	3.37	3.39	3.47	3.41	3.41	3.44
Panel B: Hospital Outcomes									
Hospital Profit (per hospital, million NTD)	19.20	21.30	26.27	4.27	4.59	5.29	2.57	2.72	3.04
Social Surplus (million NTD)	51.49	59.45	78.19	131.39	139.90	160.07	177.50	186.29	206.80
Avg. Patient-Paid Price (thousand NTD)	68.71	70.13	73.39	59.51	59.83	60.61	57.03	57.10	57.31
Hospital Portfolio Modification Rate (%)	42.27%	43.69%	48.02%	35.40%	35.61%	37.50%	36.34%	36.34%	36.87%
Panel C: NHI Subsidy Outcomes									
Total Subsidy (million NTD)	43.97	43.97	43.97	43.97	43.97	43.97	43.97	43.97	43.97
Subsidy Allocated to DES (million NTD)	14.75	16.38	19.78	17.65	19.69	24.06	18.61	20.79	25.33

Notes: "Hospital Portfolio Modification Rate" represents the frequency of adoption or upgrade events per hospital among hospitals where such modifications were feasible. $\zeta = 1$ corresponds to the estimated preference; $\zeta = 2$ doubles the marginal utility from the newest generation relative to the previous generation within the same brand; $\zeta = 0.5$ halves this marginal utility. All outcomes are averaged across simulations. Units are New Taiwan Dollars (NTD).

determines how these demand shocks are distributed: concentrated hospitals raise prices more when patients value newness more, whereas competitive markets keep prices nearly flat and pass more of the value to patients.

5.2 Alternative Reimbursement Designs for Enhanced Innovation Diffusion

We evaluate three alternative reimbursement designs, each operating through a distinct economic channel. The first adjusts DES-specific reimbursement rates, decoupling them from the BMS benchmark to directly influence provider margins and pricing decisions. The second employs selective contracting, where the insurer negotiates wholesale discounts with specific manufacturers in exchange for exclusive reimbursement eligibility. The third introduces targeted patient subsidies that reduce out-of-pocket costs for lower-income patients without disrupting market pricing mechanisms. These interventions offer different combinations of benefits and costs across three key dimensions: innovation diffusion, fiscal burden, and distributional equity.

Alternative Design #1: DES Reimbursement Adjustment. The simplest NHI reform is adjusting the reimbursement rate specifically for DES while holding BMS reimbursement constant. We parameterize this through a multiplicative factor r^d applied to DES procedures, so that the NHI pays $r^d \cdot r_t$ per DES while maintaining r_t for BMS. Setting $r^d < 1$ cuts DES reimbursement below the status quo; $r^d > 1$ raises it.

Table 6 presents market outcomes under two alternative regimes, doubling the DES reimbursement ($r^d = 2$) and halving it ($r^d = 0.5$). A higher per-DES reimbursement makes each sale more profitable, so the profit-maximizing hospital lowers the patient price to expand volume; a lower reimbursement does the reverse. How fully this pass-through reaches patients depends on market structure. In competitive markets ($N = 20$), doubling the reimbursement lowers patient prices by 35% (from 57.1 to 37.0 thousand NTD) and raises utilization by 17 percentage points, whereas duopolists retain a larger share of the transfer, cutting prices by only 20% and lifting utilization by 12 points. Halving the reimbursement is the mirror image, raising the patient-paid price by about 17% and contracting utilization by 8 points at $N = 20$. The raw upside and downside look uneven—a 35% price drop versus a 17% rise—but this simply reflects the unequal shock sizes, since doubling is twice the reimbursement change of halving; per unit of reimbursement the price response is essentially symmetric (about 20 thousand NTD per r_t at $N = 20$, 14 thousand at $N = 2$). The genuine asymmetry is across market structure: hospitals transmit reimbursement shocks to patients far more completely where competition is intense.

Table 6: Market Outcomes Under Alternative DES Reimbursement Factors (r^d)

	N=2 Hospitals			N=10 Hospitals			N=20 Hospitals		
DES Reimbursement Factor	0.5	1.0	2.0	0.5	1.0	2.0	0.5	1.0	2.0
Panel A: Patient and Product Outcomes									
Consumer Surplus (million NTD)	55.44	60.81	73.78	129.71	138.00	159.75	166.69	175.84	200.08
DES Utilization Rate (%)	32.53	38.08	49.66	38.22	45.92	62.06	40.04	48.26	65.25
Newest-Gen DES Rate (%)	29.49	34.48	45.31	34.70	41.71	56.79	36.76	44.43	60.52
Avg DES Models per Hospital	4.27	4.32	4.35	4.18	4.20	4.21	4.24	4.25	4.26
Avg Newest-Gen Models per Hospital	3.56	3.58	3.69	3.39	3.39	3.41	3.43	3.41	3.43
Panel B: Hospital Profit and Pricing									
Hospital Profit (per hospital, million NTD)	17.96	21.30	29.56	4.06	4.59	5.76	2.49	2.72	3.23
Avg Patient-Paid Price (thousand NTD)	77.30	70.13	56.22	69.33	59.83	40.73	67.13	57.10	36.98
Hospital Portfolio Modification Rate (%)	42.75	43.69	45.28	35.54	35.61	35.90	36.56	36.34	36.87
Panel C: Government Subsidy									
Total Subsidy (million NTD)	37.04	43.97	65.62	35.86	43.97	71.02	35.44	43.97	72.52
Subsidy Allocated to DES (million NTD)	6.93	16.38	43.30	8.11	19.69	54.10	8.54	20.79	57.10
Social Surplus (million NTD)	54.31	59.45	67.28	134.44	139.90	146.31	181.05	186.29	192.13

Notes: The table compares outcomes for $r^d = 0.5$ (halved DES reimbursement), $r^d = 1$ (baseline), and $r^d = 2$ (doubled DES reimbursement). Consumer Surplus, Hospital Profit (per hospital), Social Surplus, Total Subsidy, and Subsidy Allocated to DES are reported in millions of NTD; Average Patient-Paid Price is in thousands of NTD; usage rates and portfolio modification probabilities are percentages.

The gains from a richer reimbursement run through the demand margin, not the supply margin. As prices fall, newest-generation utilization climbs steeply (from 44.4% to 60.5% at $N = 20$ when $r^d = 2$), yet hospitals' offerings barely change: the average number of newest-generation models per hospital moves only from 3.41 to 3.43, and portfolio modification rates shift by at most about one percentage point, not even uniformly in the same direction across market structures. Adoption fixed costs thus generate portfolio inertia, leaving the extensive margin of technology supply far less responsive to reimbursement than patient demand is. The welfare accounting reflects a sharp fiscal trade-off: doubling the reimbursement raises both consumer surplus and social surplus (the latter from 186.3 to 192.1 million NTD at $N = 20$, net of the transfer), but inflates NHI expenditure by roughly 65% (from 43.97 to 72.52 million NTD), while halving it saves about 20% fiscally at the cost of constrained access. DES reimbursement adjustment is therefore a blunt and fiscally expensive lever: it reaches patients chiefly through prices in competitive markets and leaves the extensive margin essentially untouched.

Alternative Design #2: Selective Contracting. Selective contracting uses the government insurer's bargaining power to obtain wholesale discounts from selected manufacturers in exchange for exclusive reimbursement eligibility, a mechanism analogous to Taiwan's prior negotiations over hepatitis C treatments.[15] We model this policy by applying a discount factor $r^c < 1$ to hospital input costs for the two market leaders, Abbott and Medtronic, while making non-contracted DES ineligible for reimbursement. Rather than solving for an optimal contract, the exercise asks whether a simple selective contract can simultaneously lower patient costs, accelerate the diffusion of frontier technology, sustain manufacturer participation, and contain public spending.

Table 7 compares two selective-contracting designs that share the same exclusivity rule (only Abbott, Medtronic, and BMS remain reimbursable, while non-targeted DES are excluded) but differ in the negotiated wholesale discount: the "Excl. only" column applies no discount ($r^c = 1$), while the "Excl. + discount" column applies a 50% wholesale discount ($r^c = 0.5$). Throughout this subsection, the magnitudes we quote are *marginal* effects of adding the wholesale discount on top of exclusivity—the difference between the "Excl. + discount" and "Excl. only" columns—rather than total effects measured against the no-policy status quo of Table 6, in which all DES remain reimbursed. Layering the discount onto exclusivity lowers targeted DES prices and expands utilization in every market structure: in the observed $N = 20$ market, for example, targeted DES prices fall by 41% and DES utilization

[15] Taiwan negotiated prices for direct-acting antivirals before reimbursement began in 2017, substantially lowering treatment costs and expanding access (Chien et al., 2021; Pwu et al., 2025).

rises by about 15.5 percentage points. Targeted-brand portfolio modification, meanwhile, climbs from 25.5% to 41.2% under duopoly but only from 18.2% to 19.1% under $N = 20$.

Table 7: Market Outcomes under Selective Contracting Policy Designs

	N=2 Hospitals		N=10 Hospitals		N=20 Hospitals	
	Excl. only	Excl. + discount	Excl. only	Excl. + discount	Excl. only	Excl. + discount
Reimbursement factor (r^d)	1	1	1	1	1	1
Cost factor (r^c)	1	0.5	1	0.5	1	0.5
Consumer-related:						
Consumer Surplus (M NTD)	59.20	73.15	131.48	149.36	168.52	187.65
% of DES usage	36.37%	49.35%	39.95%	54.97%	41.77%	57.29%
% using newest generation	32.43%	45.61%	35.82%	49.21%	38.14%	52.09%
Hospital-related:						
Profit (per hospital, M NTD)	14.70	22.28	4.13	5.17	2.54	2.99
Avg. Patient-Paid Price (k NTD)	70.46	51.12	65.30	43.15	63.00	40.32
Targeted DES	64.41	45.34	58.38	35.99	56.11	33.13
Non-targeted DES	84.99	90.53	80.50	81.88	78.31	78.89
NHI-related:						
Subsidy (M NTD)	39.41	41.21	38.74	40.30	38.48	40.09
Targeted DES	10.93	18.54	11.73	19.96	12.34	20.84
BMS	28.48	22.67	27.00	20.34	26.15	19.26
Market Structure & Dynamics:						
Avg DES models per hospital	4.26	4.31	4.20	4.22	4.24	4.27
of which newest-gen	3.53	3.64	3.39	3.41	3.43	3.44
Hospital Portfolio Modification Rate (%)	40.98	40.18	35.55	36.34	36.78	37.17
Targeted DES (option-level)	25.53	41.18	17.42	18.93	18.24	19.07
Non-targeted DES (option-level)	14.52	11.41	13.26	13.03	13.46	13.29

Notes: Both columns impose the same exclusivity rule—only the targeted brands (Abbott and Medtronic) and BMS are reimbursable—and differ only in the wholesale discount: none in "Excl. only" ($r^c = 1$), 50% in "Excl. + discount" ($r^c = 0.5$). The lower NHI subsidy than in the no-policy status quo of Table 6 (≈39 vs. 43.97 million NTD) reflects only the exclusion of non-targeted DES from reimbursement; the per-procedure rate is unchanged.

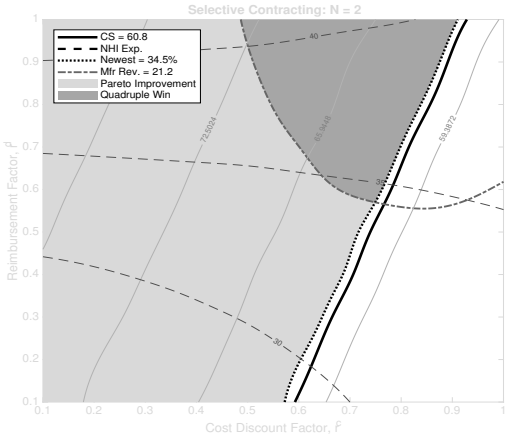
The cross-market-structure contrast in the last rows reflects a single mechanism. In the duopoly case, equilibrium markups are large and pass-through of the wholesale discount is incomplete: Bertrand competition between two firms still leaves prices well above marginal cost, so hospitals retain a sizable share of the cost reduction as added gross profit. The 50% discount lowers patient-paid targeted-DES prices by about 19 thousand NTD per procedure, but this partial pass-through is still enough to generate large consumer-surplus gains, because consumer surplus is a steep function of patient prices when baseline markups are high. The retained rent, in turn, funds a supply-side response: targeted-brand portfolio modification jumps from 25.5% to 41.2%, as hospitals become more willing to incur the fixed cost of modifying their portfolios. Newest-generation diffusion rises in parallel, to a 45.6% share at the representative point $(r^c, r^d) = (0.5, 1)$, well above the 34% CF1 baseline; this gain combines demand-side substitution toward the targeted, newest-generation-heavy brands with the distinctively duopolistic supply response.

The competitive $N = 20$ market shows the opposite picture. Because baseline competition has already compressed hospital margins close to marginal cost, the discount is passed through almost in full to patients—about 92% of the per-unit cost reduction, versus about 76% under duopoly. The resulting consumer-surplus gain is substantial (19.1 million NTD at the representative $r^c = 0.5$ policy), yet little rent stays with hospitals: targeted-brand portfolio modification barely moves, from 18.2% to 19.1%, in stark contrast to the jump to 41.2% under duopoly. Newest-generation diffusion still rises, from a 38.1% to a 52.1% share, but the increase is now purely demand-side—consumers shifting toward targeted brands whose portfolios already skew newest-generation—rather than reflecting any expansion of hospitals' offerings.

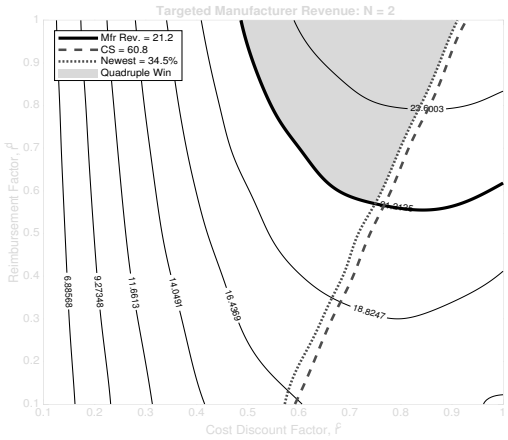
The table, however, fixes the discount at 50% and omits one outcome central to the policy's feasibility—the revenue of the contracted manufacturers. Figure 6 fills both gaps, tracing the full policy space spanned by the wholesale-discount factor r^c and the reimbursement multiplier r^d and adding the manufacturer dimension. Across this space, selective contracting moves the four outcomes that define its success in partly opposing directions, each measured against the no-contracting CF1 baseline of Table 4. Consumer surplus rises as the discount deepens (r^c falls) and as reimbursement grows more generous (r^d rises), since both lower patient out-of-pocket prices. Newest-generation diffusion, the policy's central objective, moves the same way, as cheaper targeted brands draw volume toward frontier devices. Targeted-manufacturer revenue rises with the exclusivity-driven volume gain but erodes once the discount cuts too deeply into per-unit margins, so it peaks at a modest discount. NHI expenditure scales with r^d, so the payer favors leaner reimbursement; but exclusivity alone already holds spending below the no-contracting baseline of 43.97 million NTD everywhere in the plane, so the payer's budget never binds.

Because these outcomes move in partly opposing directions, improving all of them at once is possible, if at all, only over a limited part of the policy space. A policy delivers a Pareto improvement when consumer surplus and newest-generation diffusion both weakly exceed their baselines while NHI expenditure stays weakly below; because the payer's constraint is slack everywhere, this reduces in practice to clearing the consumer-surplus and diffusion thresholds. It is a "quadruple win" when it additionally holds targeted-manufacturer revenue weakly above baseline, so that all four success dimensions—consumer surplus, newest-generation diffusion, targeted-manufacturer revenue, and NHI savings—improve simultaneously, with hospitals also gaining from the negotiated discount.

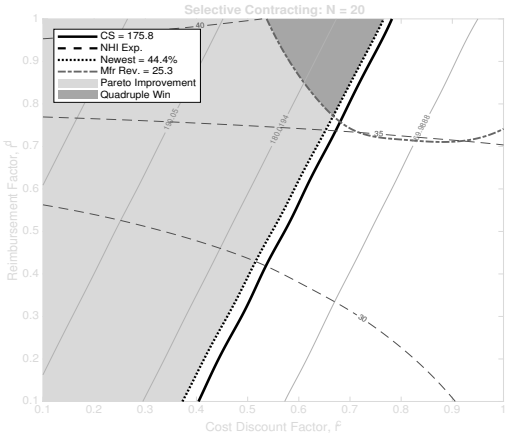
The figure plots these regions for each market structure, with the exclusivity-only benchmark at the upper-right corner and deeper discounts or leaner reimbursement moving left and down. The quadruple-win set is broad under duopoly and much narrower under $N = 20$,



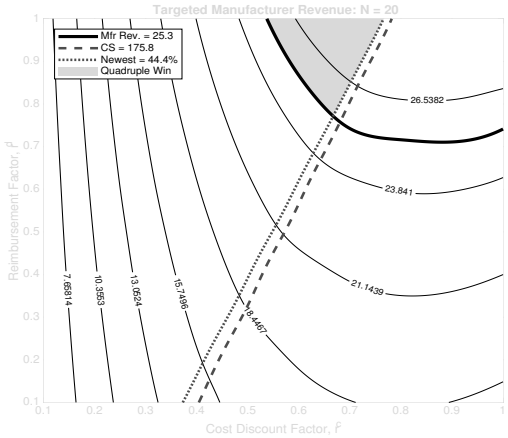
(a) Baseline iso-curves and Pareto region, $N = 2$



(b) Manufacturer revenue, $N = 2$



(c) Baseline iso-curves and Pareto region, $N = 20$



(d) Manufacturer revenue, $N = 20$

Figure 6: Policy Outcomes under Selective Contracting: Varying (r^c, r^d)

Notes: Over the (r^c, r^d) plane, iso-curves are drawn at the CF1 competition-only baseline levels for consumer surplus (solid), NHI expenditure (dashed), newest-generation DES share (dotted), and targeted-manufacturer revenue (dash-dot); r^c is the wholesale cost-discount factor and r^d the NHI reimbursement multiplier, with $r^c = r^d = 1$ at the no-discount, baseline-reimbursement corner. Left panels shade the Pareto-improvement region light gray and the "quadruple-win" sub-region—which additionally requires targeted-manufacturer revenue weakly above baseline—dark gray; right panels shade the same quadruple-win region in gray. Thin dashed contours mark NHI expenditure at 30, 35, and 40 million NTD.

with the intermediate $N = 10$ case (not plotted) in between. This follows directly from the rent mechanism above. Under duopoly, the rent hospitals retain keeps consumer surplus, diffusion, and—through the exclusivity-driven volume gain—targeted-manufacturer revenue all above baseline across a wide band, so the four conditions hold at once (Panel (a)). Under $N = 20$, near-complete pass-through leaves little rent, and because patient prices were already low the volume gain is too small to offset the lower wholesale price: targeted-

manufacturer revenue slips below baseline (Panel (d)) and becomes the binding constraint, shrinking the region to a narrow strip (Panel (c)).

The broader implication is that a policy working through the extraction and redistribution of channel rents has the most leverage where those rents are largest. Selective contracting thus operates as a complement to, rather than a substitute for, downstream competition: it delivers its largest welfare gains and strongest diffusion response precisely where competition is weakest.

Alternative Design #3: Patient Coupon Program. Unlike the reimbursement and selective-contracting policies above, which operate through hospital margins, a patient coupon is a demand-side subsidy that lowers DES out-of-pocket costs directly for targeted lower-income patients.

The motivation is the steep socioeconomic gradient in DES access under the status quo ($N = 20$), where DES uptake is 41.7% for the bottom income decile and 59.2% for the top 5% (see Online Appendix Figure A.1); this gradient is present before any coupon policy is introduced. We examine two designs: one targeted to the bottom decile (0–10%) and one extending tiered coverage to the bottom half (0–50%) of the income distribution. Under both, patients in the bottom decile receive a 36,000 NTD coupon, about 63% of the baseline patient-paid DES price (57.10 thousand NTD at $N = 20$, from Table 6); the broader design then steps the coupon down by 6,000 NTD per decile—30, 24, 18, and 12 thousand NTD for the 10–20% through 40–50% groups—and offers nothing above the median.

Table 8 reports the market impact. Coupons sharply raise DES uptake among targeted patients: in the observed $N = 20$ market, bottom-decile uptake rises from 41.7% under the status quo to about 70% under either design. Broader coverage lifts uptake among the 10–50% groups as well, but it raises fiscal costs far more steeply.

The strategic response is concentrated in prices rather than portfolios. When the coupon is limited to the bottom decile, or the market is competitive enough, hospitals absorb the subsidized demand shock with little price response, because raising the list price would forfeit demand from the larger unsubsidized population. In the duopoly case, however, broad 0–50% coverage makes a list-price increase worthwhile: the average patient-paid price climbs (from 71.3 to 73.4 thousand NTD at $N = 2$), and because the list price applies to everyone, the increase spills over onto unsubsidized patients, whose own DES uptake falls (from 41.1% to 39.5% in the top group) even as eligible patients gain (Table 8). The fiscal trade-off is equally direct: because the coupon is a fixed transfer, broader eligibility and higher uptake—not the price change—drive total coupon cost from roughly 5 to 14–17 million NTD per quarter. Portfolio modification, by contrast, barely moves under either design, since the fixed cost of adding or upgrading a brand remains far larger than the

Table 8: Market Outcomes of the Tiered Patient Coupon Program: 10% vs. 50% Coverage

	$N = 2$ Hospitals		$N = 10$ Hospitals		$N = 20$ Hospitals	
			Coupon coverage			
Statistic	0-10%	0-50%	0-10%	0-50%	0-10%	0-50%
DES Usage by Income Group (%):						
0-10% (Lowest)	61.65	60.19	67.93	67.62	70.48	70.41
10-20%	31.34	55.32	38.45	64.34	41.53	67.69
20-30%	33.49	52.37	42.12	62.79	44.49	65.43
30-40%	34.01	47.43	41.90	57.37	44.36	60.23
40-50%	36.11	44.49	46.11	56.45	48.16	58.79
50-100% (Highest)	41.13	39.53	50.39	50.18	52.82	52.83
Market Outcomes:						
Avg. Patient-Paid Price (000 NTD)	71.32	73.39	60.05	60.37	57.18	57.24
Hospital Portfolio Modification Rate (%)	43.69	44.90	35.74	35.74	36.56	36.56
Government Budget (Million NTD):						
Hospital Reimbursement	43.97	43.97	43.97	43.97	43.97	43.97
Coupon Cost	4.81	13.97	5.29	16.37	5.49	17.11
Total Expenditure	48.78	57.94	49.27	60.34	49.47	61.08

Notes: The table compares outcomes between the 0-10% Coupon Program (coupons for lowest 10%) and the 0-50% Coupon Program (coupons for lowest 50%). Coupon amounts are tiered: 36k NTD (0-10%), 30k (10-20%), 24k (20-30%), 18k (30-40%), 12k (40-50%). Income groups 50-100% receive no coupon. Hospital portfolio modification rates use the same hospital-level definition as in Table 5: the share of hospitals with a feasible adoption or upgrade opportunity that actually modify their portfolio.

localized demand gain a coupon creates. A coupon is therefore best read as a distributional instrument: narrowly targeted or in competitive markets it improves access with little price distortion, but it neither stimulates broad technology diffusion nor, under broad coverage in concentrated markets, avoids imposing a price externality on the unsubsidized.

6 Conclusion

This paper examines how hospital competition and reimbursement policy shape the diffusion of medical innovations in Taiwan's drug-eluting stent market. Our structural analysis reveals a key tradeoff: while competition effectively disciplines prices, it simultaneously erodes the surplus required to incentivize hospitals to adopt frontier technologies. Counterfactual simulations suggest that, as the number of hospitals increases, intensified price competition dominates reduced adoption incentives, leading to a net gain in consumer surplus and total welfare.

Conditional on market structure, our counterfactual simulations demonstrate that selective contracting and targeted patient coupons offer promising pathways to further improve the balance of access, costs, and equity. By leveraging exclusivity to align manufacturer and hospital incentives, selective contracting can achieve a "quadruple-win" that raises consumer surplus, lowers government expenditure, accelerates newest-generation diffusion, and sustains contracted manufacturers' revenue, with hospitals also gaining from the negotiated discount. However, the feasibility of such outcomes is contingent on market structure: the parameter space for a "quadruple-win" shrinks significantly in competitive markets, where the scarcity of exclusivity rents limits the policy's leverage.

In comparison, targeted patient coupons can effectively mitigate socioeconomic disparities in access without distorting market prices, if the program has limited coverage or the market is sufficiently competitive. Yet in these situations, they fail to stimulate broader technology adoption (the extensive margin) due to the high fixed costs of portfolio modification. Similarly, reimbursement rate adjustments across all DES models exhibit heterogeneous pass-through, with competitive markets responding more elastically than concentrated ones. Unlike selective contracting, both targeted patient coupons and DES reimbursement adjustments entail extra government spending on DES utilization. Together, these results highlight the importance of accounting for hospitals' strategic responses and market structure when designing policies to promote medical innovation diffusion.

As with any structural counterfactual analysis, these results warrant appropriate qualification. The market structure and patient preference exercises in Section 5.1 are designed as mechanism decompositions that isolate how supply-side competition and demand-side preferences for innovation separately shape hospital portfolio and pricing decisions, rather than as predictions for specific market configurations. For configurations far from the observed data, the quantitative magnitudes should be interpreted with caution. The policy counterfactuals in Section 5.2—reimbursement-rate adjustment, selective contracting, and targeted patient coupons—vary reimbursement design parameters while holding market structure fixed at the observed 20 hospitals, involve more modest departures from observed conditions, and are more directly disciplined by the data.

Our findings have broader implications for healthcare markets where intermediaries control technology access. Traditional policies focusing solely on upstream R&D incentives prove insufficient when hospitals act as strategic gatekeepers. Effective policy must recognize this intermediation role and align hospital incentives with social objectives. Future research should explore several extensions: incorporating dynamic considerations of how current adoption affects future innovation by manufacturers, examining heterogeneous hospital objectives beyond profit maximization, and testing the generalizability across different

medical technologies and healthcare systems. As healthcare costs continue rising globally, understanding how intermediaries influence the diffusion-affordability trade-off becomes increasingly critical for policy design.

References

Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt. 2005. "Competition and Innovation: an Inverted-U Relationship." *The Quarterly Journal of Economics* 120 (2): 701–728.

Agostoni, Pierfrancesco, Giuseppe G.L. Biondi-Zoccai, Gabriele L. Gasparini, Maurizio Anselmi, Giorgio Morando, Marco Turri, Antonio Abbate, Eugene P. McFadden, Corrado Vassanelli, Piero Zardini, Antonio Colombo, and Patrick W. Serruys. 2005. "Is bare-metal stenting superior to balloon angioplasty for small vessel coronary artery disease? Evidence from a meta-analysis of randomized trials." *European Heart Journal* 26 (9): 881–889.

Arrow, Kenneth J. 1962. "Economic Welfare and the Allocation of Resources for Invention." In *The Rate and Direction of Inventive Activity*, 609–626, Princeton University Press.

Bangalore, Sripal, Bora Toklu, Neil Patel, Frederick Feit, and Gregg W. Stone. 2018. "Newer-Generation Ultrathin Strut Drug-Eluting Stents Versus Older Second-Generation Thicker Strut Drug-Eluting Stents for Coronary Artery Disease." *Circulation* 138 (20): 2216–2226.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4): 841–890.

Berry, Steven T., and Philip A. Haile. 2014. "Identification in Differentiated Products Markets Using Market Level Data." *Econometrica* 82 (5): 1749–1797.

Chandra, Amitabh, David Cutler, and Zirui Song. 2011. "Chapter Six - Who Ordered That? The Economics of Treatment Choices in Medical Care." In *Handbook of Health Economics*, edited by Mark V. Pauly, Thomas G. Mcguire, and Pedro P. Barros, Volume 2. of Handbook of Health Economics 397–432, Elsevier.

Chandra, Amitabh, Evan Flack, and Ziad Obermeyer. 2024. "The Health Costs of Cost Sharing." *The Quarterly Journal of Economics* 139 (4): 2037–2082.

Chandra, Amitabh, and Jonathan Skinner. 2012. "Technology Growth and Expenditure Growth in Health Care." *Journal of Economic Literature* 50 (3): 645–80.

Chernew, Michael E., and Joseph P. Newhouse. 2011. "Chapter One - Health Care Spending Growth." In *Handbook of Health Economics*, edited by Mark V. Pauly, Thomas G. Mcguire, and Pedro P. Barros, Volume 2. of Handbook of Health Economics 1–43, Elsevier.

Chien, Rong-Nan, Sheng-Nan Lu, Raoh-Fang Pwu, Grace Hui-Min Wu, Wen-Wen Yang, and Chia-Ling Liu. 2021. "Taiwan accelerates its efforts to eliminate hepatitis C." *Global Health & Medicine* 3 (5): 293–300.

Conley, Timothy G, and Christopher R Udry. 2010. "Learning About a New Technology: Pineapple in Ghana." *American Economic Review* 100 (1): 35–69.

Credevo. 2024. "Guide to Medical Device Approval Process in Taiwan: Regulations, Requirements & Steps." Accessed March 2026.

Dafny, Leemore. 2005. "How Do Hospitals Respond to Price Changes?" *American Economic Review* 95 (5): 1525–1547.

Decarolis, Francesco. 2015. "Medicare Part D: Are insurers gaming the low income subsidy design?" *American Economic Review* 105 (4): 1547–1580.

Demirer, Mert, and Alexander L. Olssen. 2024. "Drug Rebates and Formulary Design: Evidence from Statins in Medicare Part D." *Working Paper*.

Duggan, Mark, and Fiona Scott Morton. 2010. "The Effect of Medicare Part D on Pharmaceutical Prices and Utilization." *American Economic Review* 100 (1): 590–607.

Dunn, Abe, Lasanthi Fernando, and Eli Liebman. 2023. "A Direct Measure of Medical Innovation on Health Care Spending: A Condition-Specific Approach." *Health Management, Policy & Innovation*.

Dunn, Abe, Lasanthi Fernando, and Eli Liebman. 2024. "How Much Are Medical Innovations Worth? A Detailed Analysis Using Cost-Effectiveness Studies." *working paper*.

Grennan, Matthew. 2013. "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices." *American Economic Review* 103 (1): 145–177.

Grennan, Matthew. 2014. "Bargaining Ability and Competitive Advantage: Empirical Evidence from Medical Devices." *Management Science* 60 (12): 3011–3025.

Grennan, Matthew, and Ashley Swanson. 2020. "Transparency and Negotiated Prices: The Value of Information in Hospital-Supplier Bargaining." *Journal of Political Economy* 128 (4): 1234–1268.

Head, Keith, and Thierry Mayer. 2019. "Brands in Motion: How Frictions Shape Multinational Production." *American Economic Review* 109 (9): 3073–3124.

Ho, Kate, and Robin S. Lee. 2019. "Equilibrium Provider Networks: Bargaining and Exclusion in Health Care Markets." *American Economic Review* 109 (2): 473–522.

Hsu, Jason C., and Christine Y. Lu. 2015. "The evolution of Taiwan's National Health Insurance drug reimbursement scheme." *DARU Journal of Pharmaceutical Sciences* 23 (1): 15.

Jin, Ginger Zhe, Hsienming Lien, and Xuezhen Tao. 2025. "Top-up Design and Health Care Expenditure: Evidence from Cardiac Stents." *China Economic Review* 93.

Kaiser, Christoph, Søren Galatius, Paul Erne, Franz Eberli, Hannes Alber, Hans Rickli, Giovanni Pedrazzini, Burkhard Hornig, Osmund Bertel, Piero Bonetti, Stefano De Servi, Hans-Peter Brunner-La Rocca, Ingrid Ricard, and Matthias Pfisterer. 2010. "Drug-Eluting versus Bare-Metal Stents in Large Coronary Arteries." *New England Journal of Medicine* 363 (24): 2310–2319.

Kandzari, David E, Laura Mauri, Jacques J Koolen, Joseph M Massaro, Gheorghe Doros, Hector M Garcia-Garcia, Johan Bennett, Ariel Roguin, Elie G Gharib, Donald E Cutlip, and Ron Waksman. 2017. "Ultrathin, bioresorbable polymer sirolimus-eluting stents versus thin, durable polymer everolimus-eluting stents in patients undergoing coronary revascularisation (BIOFLOW V): a randomised trial." *The Lancet* 390 (10105): 1843–1852.

Kyle, Margaret K. 2007. "Pharmaceutical Price Controls and Entry Strategies." *Review of Economics and Statistics* 89 (1): 88–99.

Liu, Charles, Hiroshi Gotanda, Dhruv Khullar, Thomas Rice, and Yusuke Tsugawa. 2021. "The Affordable Care Act's Insurance Marketplace Subsidies Were Associated With Reduced Financial Burden For US Adults." *Health Affairs* 40 (3): 496–504.

McClellan, Mark. 2011. "Reforming Payments to Healthcare Providers: The Key to Slowing Healthcare Cost Growth While Improving Quality?" *Journal of Economic Perspectives* 25 (2): 69–92.

Navarese, Eliano Pio, Mariusz Kowalewski, David Kandzari, Alexandra Lansky, Bartosz Górny, Łukasz Kołtowski, Ron Waksman, Sergio Berti, Giuseppe Musumeci, Ugo Limbruno, Rene J. van der Schaaf, Malte Kelm, Jacek Kubica, and Harry Suryapranata. 2014. "First-generation versus second-generation drug-eluting stents in current clinical practice: updated evidence from a comprehensive meta-analysis of randomised clinical trials comprising 31 379 patients." *Open Heart* 1 (1): e000064.

Nexon, David, and Stephen Ubl. 2010. "Implications Of Health Reform For The Medical Technology Industry." *Health Affairs* 29 (7): 1325–1329.

Olateju, Olajumoke A., Chan Shen, and James Douglas Thornton. 2024. "The Affordable Care Act and income-based disparities in health care coverage and spending among nonelderly adults with cancer." *Health Affairs Scholar* 2 (5): qxae050.

Pakes, A., J. Porter, Kate Ho, and Joy Ishii. 2015. "Moment Inequalities and Their Application." *Econometrica* 83 (1): 315–334.

Pwu, Raoh-Fang, Wen-Wen Yang, Grace Hui-Min Wu, Sheng-Nan Lu, and Rong-Nan Chien. 2025. "Investing in health: Policy lessons from reimbursing direct-acting antivirals for hepatitis C in Taiwan." *Journal of the Formosan Medical Association* 124 S141–S148, Toward Hepatitis C Elimination in Taiwan: 2025 Report.

Schumpeter, Joseph Alois. 1942. *Capitalism, Socialism and Democracy.*. New York: NY: Harper & Brothers.

Skinner, Jonathan, and Douglas Staiger. 2015. "Technology Diffusion and Productivity Growth in Health Care." *Review of Economics and Statistics* 97 (5): 951–964.

Sorensen, Alan T. 2003. "Insurer-Hospital Bargaining: Negotiated Discounts in Post-Deregulation Connecticut." *The Journal of Industrial Economics* 51 (4): 469–490.

Starc, Amanda, and Ashley Swanson. 2021. "Preferred Pharmacy Networks and Drug Costs." *American Economic Journal: Economic Policy* 13 (3): 406–446.

Tebaldi, Pietro. 2025. "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA." *The Review of Economic Studies* 92 (1): 586–620.

Tu, Jack V., James Bowen, Maria Chiu, Dennis T. Ko, Peter C. Austin, Yaohua He, Robert Hopkins, Jean-Eric Tarride, Gord Blackhouse, Charles Lazzam, Eric A. Cohen, and Ron Goeree. 2007. "Effectiveness and safety of drug-eluting stents in Ontario." *New England Journal of Medicine* 357 (14): 1393–1402.

Valgimigli, Marco, Héctor Bueno, Robert A. Byrne, Jean-Philippe Collet, Francesco Costa, Anders Jeppsson, Peter Jüni, Adnan Kastrati, Philippe Kolh, Laura Mauri, Gilles Montalescot, Franz-Josef Neumann, Mate Petricevic, Marco Roffi, Philippe Gabriel Steg, Stephan Windecker, and Jose Luis Zamorano. 2018. "2017 ESC Focused Update on Dual Antiplatelet Therapy in Coronary Artery Disease Developed in Collaboration with EACTS." *European Heart Journal* 39 (3): 213–260.

Wollmann, Thomas G. 2018. "Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles." *American Economic Review* 108 (6): 1364–1406.

Zhu, Zheng, Quan Wang, Qiang Sun, Joel Lexchin, and Li Yang. 2023. "Improving access to medicines and beyond: the national volume-based procurement policy in China." *BMJ Global Health* 8 (7): e011535.

Online Appendix

to

"Diffusing Innovations Under Market Competition: Evidence from Drug-Eluting Stents"

Appendix A DES Product Classification and Generation Assignment

The empirical analysis requires mapping each observed DES model to a manufacturer-specific product generation. We define generations within manufacturer rather than across manufacturers because DES makers improve products along different clinical and engineering dimensions, and patients are more likely to perceive "newness" relative to the product line offered by the same brand. This definition is also the one used to construct the newest-generation attribute in the demand model.

We consolidate products released nearly simultaneously when their technical specifications are sufficiently similar for the purposes of our market-level analysis. For example, Abbott generations 1 and 2 and Boston Scientific generations 4 and 5 are treated as the same effective generation. Models that are globally available but not observed in our claims and price data are listed for transparency but not used as distinct alternatives. Less common brands are aggregated into the "Other" category in estimation. Table A.1 reports the resulting classification and shows how the product data map into the adoption and upgrade actions used in the supply model.

Table A.1: Major DES Models and Generation Assignment

Brand Name (Abbr.)	Gen	Initial Use	Stent Material	Coating Material	Length Choices	Diameter Choices	Note
Abbott (AB)	1	Oct 2008	L-605 CoCr Alloy	Everolimus and polymers	6	6	
	2	Nov 2008	L-605 CoCr Alloy	Everolimus and polymers	6	4	Consolidated with Gen 1
	3	Feb 2011	L-605 CoCr Alloy	Everolimus and polymers	8	5	
Medtronic (M4)	1	Mar 2007	Co alloy and PC coating	ABT-578	8	6	
	2	Jun 2008	Co alloy and PC coating	Zotarolimus	8	6	
	3	Jul 2009	Co alloy w/ BioLinx coating	Zotarolimus	8	6	
	4	Apr 2011	Co alloy w/ BioLinx coating	Zotarolimus	9	6	Rarely used, dropped
	5	Jan 2013	Co alloy w/ BioLinx coating	Zotarolimus	11	6	
Cordis (CD)	1	Nov 2006	316L stainless steel w/ Titanium Oxynitride	Sirolimus and polymers	6	8	
	2	Jul 2007	316L stainless steel	Sirolimus and polymers	6	5	
Boston Scientific (SB)	1	Nov 2006	316L stainless steel w/ polymers	Paclitaxel	7	8	Consolidated with Gen 2
	2	Oct 2006	316L stainless steel w/ Translute Polymer Carrier	Paclitaxel	7	5	
	3		316L stainless steel w/ Translute Polymer Carrier	Paclitaxel	8	8	Not observed in data
	4	Jan 2011	PtCr alloy w/ Translute Polymer Carrier	Paclitaxel	8	8	
	5	Nov 2010	PtCr alloy w/ PBMA-poly coating	PVDF-HFPpoly and Everolimus	8	6	Consolidated with Gen 4
Bio Sensor (BS)	1	Oct 2010	316L stainless steel	Biolimus w/ PLA	7	7	
	2	Feb 2012	316L stainless steel	Biolimus w/ PLA	8	6	

Notes: This table details the major DES models from five prominent manufacturers. Generations are consolidated if released nearly simultaneously with similar specifications. Some globally available models are omitted as they do not appear in our data. Less common brands are aggregated into an "Other" category for analysis and are not listed here.

Appendix B Equilibrium Computation

This appendix details the computational procedure for counterfactual equilibria. The algorithm solves a two-stage game each period: hospitals first adjust portfolios, then compete in prices. This sequential structure captures the distinction between long-run technology investments and short-run pricing decisions.

Let M_{ht} denote hospital h's portfolio in period t, $\mathcal{M}_t$ the market-wide portfolio profile, and $\mathcal{P}(\mathcal{M}_t)$ the Nash equilibrium prices given portfolios. Hospital payoffs are $\pi_{ht}(\cdot)$ net of variable costs, with fixed adjustment costs $f^M(M_{ht}, M_{h,t-1})$ for portfolio changes. The algorithm proceeds via iterated best responses until convergence, as detailed in Algorithm 1.

Algorithm 1 (Counterfactual Equilibrium Simulation)**.**

1. ***Initialization for period*** t***.*** *At the beginning of period* t*, each hospital observes the simulated portfolios and prices from period* $t-1$*. For the initial iteration (*$s = 0$*), set each hospital* h*'s portfolio equal to its previous-period portfolio:*

$$M_{ht}^{(0)} = M_{h,t-1}.$$

2. ***Iterative best responses in portfolios.*** *For iteration* $s \geq 1$*, consider hospitals sequentially. For a given hospital* h*, take other hospitals' portfolios from iteration* $s-1$*, denoted by* $\mathcal{M}_t^{(s-1)}$*, and evaluate each feasible portfolio* M *for* h *by:*

   (a) *Solving the pricing subproblem given the candidate profile* $\big(M, \mathcal{M}_t^{(s-1)}\big)$ *to obtain the vector of equilibrium prices*

$$\mathcal{P}\big(M, \mathcal{M}_t^{(s-1)}\big).$$

   (b) *Computing hospital* h*'s operating payoff at those prices,*

$$\pi_{ht}\Big(M, \mathcal{P}\big(M, \mathcal{M}_t^{(s-1)}\big); \mathcal{M}_t^{(s-1)}\Big).$$

   (c) *Subtracting the fixed adjustment cost relative to period* $t-1$ *to form net payoff,*

$$\pi_{ht}\Big(M, \mathcal{P}\big(M, \mathcal{M}_t^{(s-1)}\big); \mathcal{M}_t^{(s-1)}\Big) \; - \; f^M\big(M, M_{h,t-1}\big).$$

   Update hospital h*'s portfolio to the maximizer,*

$$M_{ht}^{(s)} \in \arg\max_M \Big\{\pi_{ht}\Big(M, \mathcal{P}\big(M, \mathcal{M}_t^{(s-1)}\big); \mathcal{M}_t^{(s-1)}\Big) - f^M\big(M, M_{h,t-1}\big)\Big\}.$$

After all hospitals have updated in iteration s, denote the resulting portfolio profile by $\mathcal{M}_t^{(s)}$. If $\mathcal{M}_t^{(s)} = \mathcal{M}_t^{(s-1)}$, declare convergence for period t and proceed. Otherwise, continue to iteration $s+1$.

3. ***Finalize the period-t outcome.*** *The converged profile $\mathcal{M}_t$ and the associated equilibrium prices $\mathcal{P}(\mathcal{M}_t)$ constitute the simulated market outcome for period t.*

4. ***Advance to period $t+1$.*** *Use $\mathcal{M}_t$ as the starting point for period $t+1$ and repeat the steps above until the final period.*

When the best-response procedure cycles rather than converging, we randomly select from the cycling portfolios as the equilibrium outcome. To ensure robustness, we average results across five independent simulation runs with varying initial conditions.

Appendix C Baseline DES Adoption by Patient Income

This appendix documents the baseline socioeconomic gradient in DES access that motivates the patient-coupon counterfactual in Section 5.2. Figure A.1 plots the share of patients choosing DES and the share choosing newest-generation DES against patient income (with the underlying patient-count distribution overlaid on the right axis) under the status-quo $N = 20$ market. DES utilization rises monotonically with income from about 41.7% in the bottom decile to about 59.2% in the top 5%—a roughly 17.5 percentage-point gap that is already present before any coupon policy is introduced. The newest-generation share follows the same gradient, suggesting that both extensive-margin (any DES) and intensive-margin (frontier DES) adoption are constrained by the out-of-pocket cost burden for lower-income patients.

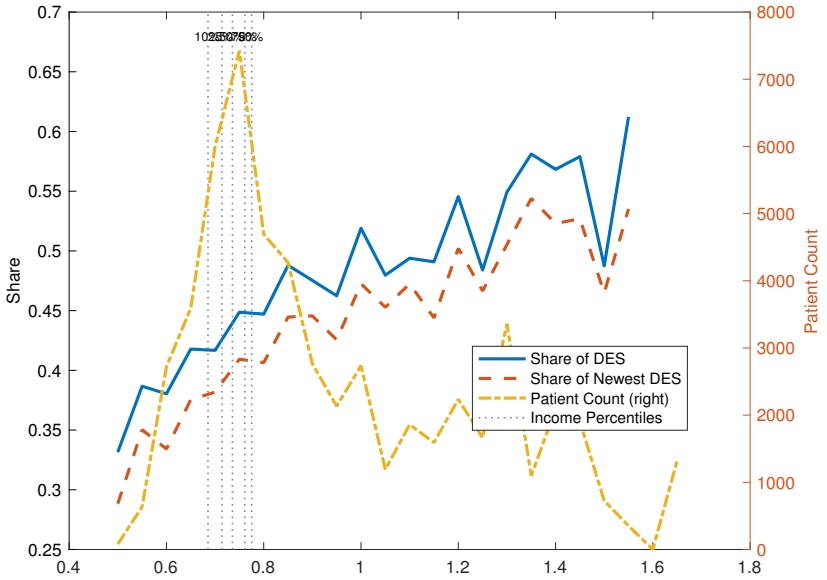


Figure A.1: Distribution of Income and DES Uptake under the Status Quo ($N = 20$)

Notes: The solid line plots the share of patients choosing any DES; the dashed line plots the share choosing newest-generation DES; the dash-dot line shows the underlying patient count (right axis). Dotted vertical lines mark the 10th, 25th, 50th, 75th, and 90th income percentiles. Income on the horizontal axis is a relative index; higher values correspond to higher-income patients. Shares are computed from CF1 baseline simulations of the observed $N = 20$ market.

Appendix D Patient Valuation Measures

This appendix quantifies patients' valuation of technological innovation through willingness to pay (WTP) and willingness to travel (WTT) measures. Using our demand estimates, we translate the utility gain from accessing newest-generation technology into monetary and distance equivalents that inform the welfare analysis in the main text.

Following standard discrete choice theory, WTP for the newest-generation attribute equals the ratio of its marginal utility to the marginal utility of income (negative of the price coefficient):

$$\text{WTP}_{\text{newest},i} = -\frac{\beta_{2i}}{\beta_{5i}}$$

where β_{2i} captures patient i's preference for newest-generation technology and β_{5i} is the price sensitivity parameter. Both coefficients vary with patient characteristics through the interaction terms estimated in the preferred demand specification in the main text.

For the average patient in our sample, we compute WTP by evaluating these coefficients at the sample means (Charlson Index: 0.915, Income category: 1.007, Male: 0.766, Age: 65.304):

$$\begin{aligned}
\text{WTP}_{\text{avg}} &= -\frac{A_{\text{avg}}}{B_{\text{avg}}}, \\
A_{\text{avg}} &= 0.636 + (-0.029 \times 0.915) + (-0.175 \times 1.007) \\
&\quad + (0.040 \times 0.766) + (0.001 \times 65.304), \\
B_{\text{avg}} &= -0.025 + (-0.004 \times 0.915) + (0.004 \times 1.007) \\
&\quad + (-0.004 \times 0.766) + (-0.001 \times 65.304).
\end{aligned} \tag{11}$$

This calculation yields a WTP of approximately 5.69 thousand NTD, indicating that the typical patient values access to the newest generation at nearly 10% of the average DES price.

Analogously, WTT measures the additional distance a patient would travel to access newest-generation technology:

$$\text{WTT}_{\text{newest},i} = -\frac{\beta_{2i}}{\beta_{3i}}$$

where β_{3i} represents the disutility from travel distance. At sample means:

$$\begin{aligned}
\text{WTT}_{\text{avg}} &= -\frac{C_{\text{avg}}}{D_{\text{avg}}}, \\
C_{\text{avg}} &= 0.636 + (-0.029 \times 0.915) + (-0.175 \times 1.007) \\
&\quad + (0.040 \times 0.766) + (0.001 \times 65.304), \\
D_{\text{avg}} &= -0.088 + (-0.007 \times 0.915) + (-0.012 \times 1.007) \\
&\quad + (-0.002 \times 0.766) + (0.000 \times 65.304).
\end{aligned} \tag{12}$$

This yields approximately 4.9 kilometers, a substantial distance given the dense hospital network in the Taipei metropolitan area.

These valuations reveal strong patient preferences for innovation that create powerful adoption incentives on the supply side. The magnitude of these preferences helps explain why hospitals invest in costly portfolio upgrades despite the fixed costs documented in our supply estimates, and why selective contracting policies that restrict access to certain innovations generate substantial welfare losses for excluded patients.

Appendix E Demand Elasticity Calculations

This appendix details the methodology for computing demand elasticities and documents the substantial heterogeneity that shapes our welfare conclusions. We derive individual-specific own-price elasticities for each DES alternative using the preferred mixed logit parameter estimates reported in the main text.

Given our linear price specification in the utility function, the own-price elasticity for individual i considering product j at time t takes the standard logit form:

$$\eta_{ijt} = \beta_i^p \cdot p_{jt} \cdot (1 - \hat{P}_{ijt}) \tag{13}$$

where p_{jt} denotes the out-of-pocket payment in thousand NTD, $\hat{P}_{ijt}$ represents the predicted choice probability from our demand model, and β_i^p captures the individual-specific price sensitivity incorporating observed heterogeneity:

$$\beta_i^p = \beta_0^p + \beta_1^p \cdot \text{Charlson}_i + \beta_2^p \cdot \text{Income}_i + \beta_3^p \cdot \text{Male}_i + \beta_4^p \cdot \text{Age}_i \tag{14}$$

The empirical distribution of calculated elasticities demonstrates uniformly elastic demand throughout the market. Individual elasticity values cluster primarily within the -3 to -1 range, with the mode occurring around -2. Such pronounced price responsiveness stems from the significant co-payments required under Taiwan's supplementary reimbursement framework, where patients bear the incremental cost above the base reimbursement rate.

Table A.2: Heterogeneity in Own-Price Demand Elasticity for DES

Stratum	Mean Elasticity	Count	Stratum	Mean Elasticity	Count
Panel A: By Brand			*Panel C: By Income Decile*		
Abbott	−1.998	7,206	1 (Lowest)	−2.106	2,164
Bio Sensor	−1.964	2,002	2	−2.077	2,540
Cordis	−2.306	3,004	3	−2.061	2,599
Medtronic	−2.021	9,083	4	−2.038	3,098
Boston Scientific	−1.861	5,662	5	−2.032	3,042
Other	−1.873	2,404	6	−2.000	3,017
Panel B: By Charlson Comorbidity Index			7	−2.024	3,044
0	−1.840	11,766	8	−1.947	3,000
1	−2.038	13,000	9	−1.897	3,287
2	−2.243	3,700	10 (Highest)	−1.876	3,570
3	−2.422	681	*Panel D: By Gender*		
4	−2.577	162	Female	−1.910	6,758
5	−2.736	52	Male	−2.024	22,603
			Panel E: By Age Quartile		
			Q1 (Youngest)	−1.893	7,026
			Q2	−1.966	7,891
			Q3	−2.035	7,760
			Q4 (Oldest)	−2.103	6,684

Notes: Each panel reports the mean of individual-level own-price elasticities for DES, $\eta_{ijt} = \beta_i^p \cdot p_{jt} \cdot (1 - \hat{P}_{ijt})$, evaluated using the preferred mixed-logit estimates (Column 4 of Table 2) and the realized choice probabilities. Counts denote the number of patient-product observations in each cell. Panel A varies the focal product brand at fixed patient characteristics; Panels B–E partition the patient sample. The brand-level standard deviations underlying Panel A range from 0.30 (Medtronic) to 0.40 (Boston Scientific), so the brand rankings are stable within sampling variation.

Substantial heterogeneity emerges along both product and patient dimensions, summarized in Table A.2. Panel A reveals systematic variation across manufacturers. Cordis exhibits the highest mean elasticity (−2.31), about 15 percent more elastic than Abbott, Medtronic, and Bio Sensor (all clustered near −2.00), while Boston Scientific and the residual "Other" category sit at the inelastic end of the distribution near −1.87. This ordering tracks each brand's positioning relative to the technological frontier: Cordis offered older-generation devices for a larger fraction of our sample period, so the patients who selected Cordis tended to be those weighing a moderate clinical premium against a substantial out-of-pocket cost. Brands with stronger newest-generation portfolios attract patients whose clinical preference for the frontier device dominates the price calculus, lowering their measured elasticity. These brand-level differences underpin the asymmetric welfare effects of selective contracting documented in Section 5.2: a wholesale discount applied to a high-elasticity brand mechanically

generates a larger volume response, and hence more pass-through to consumers, than the same discount applied to a low-elasticity brand.

Patient-level heterogeneity is even more striking. Panels B–E establish three robust gradients. First, price sensitivity rises monotonically with comorbidity burden, from -1.84 at a Charlson score of zero to -2.74 at a score of five, an almost 50 percent increase. This pattern is consistent with the clinical-channeling interpretation discussed in the demand section: high-Charlson patients face stronger medical contraindications for DES (longer dual antiplatelet therapy, bleeding risk, comorbid renal or hepatic impairment) and are therefore closer to the margin of substituting toward the fully reimbursed BMS option, so even modest out-of-pocket increases shift their choices appreciably. Second, elasticity declines steadily with income, from -2.11 in the bottom decile to -1.88 in the top decile; the lowest-income patients exhibit price sensitivities roughly 12 percent greater than their highest-income counterparts, a gap that compounds along the entire income distribution rather than being driven by tail outliers. Third, age has a quantitatively similar effect: patients in the oldest quartile are about 11 percent more elastic than those in the youngest quartile, consistent with lower lifetime earnings and a shorter expected horizon over which to amortize the DES clinical benefit. Gender effects are smaller in magnitude, with male patients only modestly more elastic than female patients (-2.02 versus -1.91).

These heterogeneity patterns provide the economic foundation for the counterfactual results in Section 5. The patient-coupon program (Alternative Design #3) leverages the income gradient directly: a means-tested transfer shifts a high-elasticity subgroup along its steep portion of the demand curve, generating relatively large welfare gains per dollar of fiscal outlay. Selective contracting (Alternative Design #2) operates through the brand-elasticity ordering: directing exclusive volume to a moderate-elasticity brand translates the negotiated wholesale discount into a sizable utilization increase without requiring large absolute changes in the policy parameters. Uniform DES reimbursement expansion (Alternative Design #1), by contrast, treats elastic and inelastic submarkets identically and therefore yields the smallest welfare gain per fiscal dollar of the three policies, a comparative-statics pattern that recurs across all three counterfactual exercises.

Appendix F Supply Model Validation

This appendix presents two complementary validation exercises for the supply-side model. The first compares aggregate counts of adoption, upgrade, and no-change events between the data and the model's predicted probabilities, both at the brand level (Table A.3) and over time (Figure A.2). The second exploits the cross-section of feasible portfolio-adjustment

opportunities to check whether the implied fixed costs are systematically lower for hospitals that took an action than for hospitals that had the same option but chose the status quo (Table A.4). Neither exercise is a calibration target of the likelihood, so agreement with the data provides genuine evidence on the model's economic content rather than mechanical overfitting.

We begin with aggregate event counts. The imputed values sum the model-predicted probabilities of each event across all hospitals and quarters.

Table A.3: Imputed vs. Actual Adoption and Upgrade Occurrences

	Abbott	Bio Sensor	Cordis	Medtronic	Other	Boston Sci.
Actual adoptions	19	15	3	17	15	0
Imputed adoptions	19	15	5	17	15	0
Actual upgrades	14	13	11	45	0	13
Imputed upgrades	14	13	10	36	0	9
Actual no-change	153	77	65	145	51	82
Imputed no-change	153	79	64	157	54	88

Notes: Imputed statistics sum predicted probabilities across all hospitals for each DES brand. "No-change" refers to instances with no portfolio modification when at least one adjustment option was feasible.

The brand-by-brand counts in Table A.3 show a close match between predicted and observed adoption and upgrade events. The model recovers actual adoptions exactly for Abbott (19), Bio Sensor (15), Medtronic (17), and the residual "Other" category (15), and to within two events for Cordis (5 imputed versus 3 actual). On the upgrade margin the fit is also tight for Abbott and Bio Sensor, while Medtronic and Boston Scientific show modest under-prediction (36 versus 45 upgrades for Medtronic; 9 versus 13 for Boston Scientific), consistent with these brands' richer multi-generation product trajectories during the sample. The no-change row, which counts feasible opportunities that ended in status quo, is within a few percent of the observed total for every brand. Because the likelihood targets individual choice probabilities rather than these aggregated counts, their alignment provides reassurance that the estimated fixed-cost parameters are not merely rationalizing the average adoption rate but also matching the brand-level allocation of those events.

Figure A.2 extends the comparison to the temporal dimension by plotting model-imputed and observed counts of adoption (left panel) and upgrade (right panel) events in each quarter of the sample. The two series track each other closely throughout the 2007–2013 window, including the cluster of adoption activity that follows the first NHI reimbursement cut in 2009 and the wave of upgrades in 2011–2012 around the introduction of third-generation DES products. The model neither smooths these episodes into a flat baseline nor amplifies

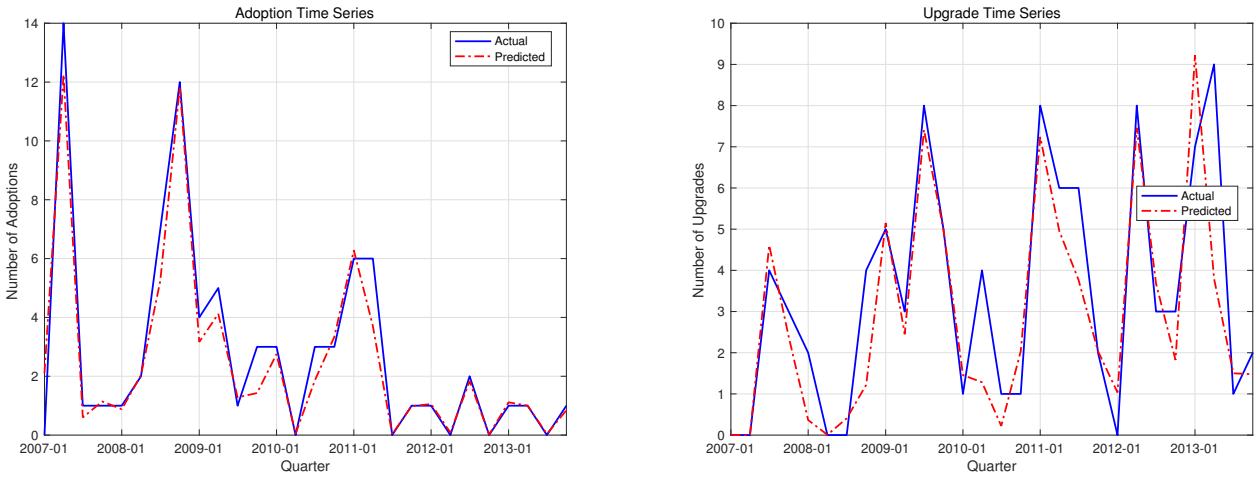


Figure A.2: Imputed vs. Actual Frequency of Adoption and Upgrade Over Time

Notes: "Imputed" is the sum of model-predicted probabilities of hospitals' adoptions and upgrades in each quarter; "Actual" is the observed count.

them into spurious volatility, indicating that the coefficients on peer-adoption, portfolio-size, and time-since-introduction variables successfully capture the dynamic forces driving portfolio adjustment. Residual deviations are concentrated in the lowest-activity quarters, where small absolute discrepancies translate into large percentage gaps; these quarters carry little weight in the counterfactual exercises, which operate at higher activity levels.

We next compare implied costs for hospitals that actually took an action with hospitals that could have taken the same action but chose not to, restricting attention to feasible opportunities only.

Table A.4: Implied Fixed Costs: Actual vs. Counterfactual Decisions

	Hospitals that acted			Hospitals that did not act		
	N	Median	IQR	N	Median	IQR
Panel A: Adoption (new brand)						
All brands	79	−3.79	[−6.29, −1.54]	261	7.23	[4.47, 10.27]
Panel B: Upgrade (same brand)						
All brands	100	−1.34	[−3.28, 0.17]	308	1.69	[−0.42, 3.38]

Notes: Each observation is a feasible portfolio adjustment opportunity at the hospital, brand, and quarter level. "Hospitals that acted" are those that actually adopted (Panel A) or upgraded (Panel B); "hospitals that did not act" had the option available but chose the status quo. Implied costs are in millions of NTD. Negative values indicate that the net payoff gain exceeded the adjustment cost.

Table A.4 provides a more direct check on the structural interpretation of the implied fixed-cost estimates. For adoption (Panel A), the median implied cost among the 79 hospital-

brand-quarter cells where a brand was actually added is −3.79 million NTD, meaning that the net flow gain from the new brand exceeded the estimated adoption cost by about 3.8 million on average. By contrast, the 261 cells where adoption was feasible but did not occur have a median implied cost of +7.23 million, an eleven-million NTD wedge between actors and non-actors. The upgrade margin (Panel B) shows the same ordering with a smaller gap (−1.34 versus +1.69 million), reflecting that upgrading is less costly than adopting a new brand and that the marginal decision is therefore tighter. The interquartile ranges, reported in brackets, do not overlap in the adoption panel and overlap only slightly in the upgrade panel, so the wedge is not driven by a handful of extreme observations. Because nothing in the likelihood directly targets the median implied cost within each subgroup, the alignment with observed behavior supports the interpretation that our estimated fixed-cost surface tracks a genuine economic adjustment margin rather than absorbing residual noise.