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

October 6, 2025

Abstract

This paper investigates how hospital competition and insurance reimbursement policy shape the diffusion of medical innovations. Using patient-level data from Taiwan's drug-eluting stent (DES) market in the Taipei metropolitan area, we estimate a structural model of hospital behavior, incorporating patient demand and hospitals' endogenous portfolio and pricing decisions. Our analysis reveals a fundamental trade-off: intense competition lowers prices but weakens hospitals' incentives to adopt new technologies, with total welfare peaking at an intermediate level of concentration. We then evaluate policy solutions. Selective contracting—where the government insurer negotiates exclusive wholesale discounts—can achieve a "quadruple-win" for consumers, hospitals, participating manufacturers, and the payer. In contrast, increasing the insurer's DES-specific reimbursement encourages hospitals to expand DES technology adoption and lower patient payment per DES, but creates a substantial fiscal burden for the insurer. Alternatively, a patient coupon program targeting low-income patients improves equity but has limited market-wide impact, as hospitals barely modify their price or portfolio decisions in response. These findings highlight that effective technology diffusion policies must account for downstream hospitals' strategic responses and their market competition.

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, 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, and University of Pennsylvania 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, Fernando and Liebman (2023), Dunn, Fernando and Liebman (2024)). This dual impact creates a fundamental policy trade-off: 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 trade-off, the effectiveness of these policies is not automatic (Nexon and Ubl (2010), McClellan (2011), Chandra, Flack and Obermeyer (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, Cutler and Song (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 empirically study this intermediary dynamics, we analyze the market for drug-eluting stents (DES) in the Taipei metropolitan area of Taiwan. The advent of DES transformed the treatment of coronary artery disease, offering substantial clinical advantages over the older generation of bare-metal stents (BMS) at a significant cost premium. This market 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 trade-off 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 system-wide, 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—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—whether to adopt a new brand, upgrade an existing one, or maintain the status quo—weighing expected profits against the fixed costs of adjustment. In the second stage, conditional on the market-wide portfolio configuration, hospitals engage in Bertrand-Nash pricing competition.

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-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 trade-off: while more intense competition disciplines prices and benefits consumers, it also erodes hospital profit margins, thereby weakening the incentive to invest in frontier technologies. As a consequence, total welfare follows an inverted-U pattern and is maximized at an intermediate level of concentration. 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 assess the efficacy of different reimbursement designs. A supply-side policy of "selective contracting," wherein the NHI leverages its bargaining power to negotiate wholesale price discounts with specific manufacturers in exchange for exclusive reimbursement, is found to generate a "quadruple-win" by increasing surplus for consumers, hospitals, and participating manufacturers while containing public expenditure. In contrast, increasing the NHI's DES-specific reimbursement encourages hospitals to expand DES technology adoption and lower patient payment per DES, but creates a substantial fiscal burden for the NHI. Alternatively, a demand-side "patient coupon" targeting the bottom 10% low-income individuals improves equity but with a limited market-wide impact because hospitals barely modify their price or portfolio management decisions in response. These results underscore that effective policy design must account for the equilibrium responses of the downstream intermediaries who ultimately determine how medical innovation translates into public value.

1.1 Related Literature

Our research contributes to five main strands of economic literature. First, our work speaks directly to the classic industrial organization literature on the relationship between market competition and innovation. While the foundational debate—from Schumpeter (1942) and Arrow (1962) to modern theoretical and empirical analyses such as Aghion et al. (2005)—has centered on the incentives for R&D investment at the producer level, we shift the analytical focus to the *diffusion of innovations by intermediaries*. In our setting, hospitals act as downstream agents who face strategic trade-offs analogous to those of upstream innovators, including the cannibalization of existing product lines and business stealing from rivals (Berry and Waldfogel, 1999; Mankiw and Whinston, 1986). By modeling and quantifying these forces at the intermediary level, we show how portfolio choice—specifically, the discrete decision to adopt a new brand versus upgrade an existing one—becomes the key strategic margin governing the speed and scope of diffusion.

Second, we build upon the literature on the diffusion of medical technologies. Prior work has identified numerous determinants of technology adoption, including intellectual property regimes (Kyle, 2007), social learning among physicians (Conley and Udry, 2010), and the potential for new technologies to exacerbate health disparities (Skinner and Staiger, 2015). Our primary contribution here is to model the *hospital*, rather than the individual physician, as the key strategic decision-maker. By structurally modeling the hospital's portfolio choice, we endogenize the availability of new technologies to patients at the market level. This framework allows us to distinguish between the distinct economic considerations of adopting an entirely new brand versus upgrading an existing one, revealing how patient preferences for "newness" are transmitted through hospital incentives to shape market-wide diffusion paths, a dynamic also explored in studies of consumer learning (Ching, Erdem and Keane, 2013; Collard-Wexler, Grennan and Steck, 2024).

Third, our paper contributes to the research on provider responses to reimbursement policies. A rich body of work has documented a wide range of behavioral responses to payment design, including patient upcoding (Dafny, 2005), adjustments to treatment intensity (Jin, Lien and Tao, 2025), and quality investment decisions driven by payer mix (Garthwaite, Ody and Starc, 2022). We highlight a complementary and economically significant channel: the strategic management of technology portfolios and the associated pricing of new products. Our findings are consistent

with related studies showing that payment policy can either encourage or discourage technology adoption (Dunn, Fernando and Liebman, 2023, 2024) and affect hospitals' capacity for capital investment (Yurukoglu, Liebman and Ridley, 2017). Our unified equilibrium framework advances this literature by allowing for the joint evaluation of diffusion rates, patient welfare, hospital profits, and insurer spending, answering the call for such integrated analysis by McClellan et al. (2017).

Fourth, we inform the literature on selective contracting in healthcare. Selective contracting is a widely studied cost-containment tool, with applications in prescription drug formularies (Duggan and Morton, 2010; Olssen and Demirer, 2024), provider networks (Sorensen, 2003; Pakes et al., 2015; Ho and Lee, 2019), and pharmacy benefit management (Starc and Swanson, 2021). The bulk of this research centers on the trade-off between the lower prices achieved through bargaining and the potential welfare losses from restricted patient access. We offer a new perspective by analyzing selective contracting through the lens of innovation diffusion. We demonstrate how this policy tool can be used not just to contain costs, but to actively steer technology adoption and strategically redistribute surplus among patients, providers, and manufacturers, creating the potential for "quadruple-win" outcomes.

Finally, our work contributes to the specific economics of the cardiac stent market. Prior research in this area has provided valuable insights into the upstream segment of the supply chain, including hospital–manufacturer price negotiations (Grennan, 2013, 2014; Grennan and Swanson, 2020) and the role of marketing and lobbying in procurement decisions (Bergman, Grennan and Swanson, 2021, 2022). We complement this body of work by focusing on the downstream market where patients and hospitals interact. In doing so, we illuminate how hospitals' strategic portfolio and pricing decisions ultimately mediate the process by which upstream innovations reach patients and determine the prices they face.

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

This section details the empirical context of our study: the market for cardiac stents in Taiwan. We begin by outlining the key institutional features of this market, 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. Together, this institutional and data foundation provides the necessary context for the structural model developed in Section 3.

2.1 Institutional Setting: Stent Implantation in Taiwan

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 generation 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.

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. DES are substantially more expensive than BMS, and the drug coating can delay arterial healing, which in early-generation models was associated with a higher risk of stent thrombosis (blood clot formation), particularly for older patients or those with severe comorbidities. Consequently, DES require a longer course of dual antiplatelet therapy (DAPT) to mitigate this risk. The clinical risk profile of DES has improved over successive generations, a key feature of the technological "newness" that we model on the demand side.

The healthcare environment in which these clinical decisions are made is defined by Taiwan's National Health Insurance (NHI) program. During our study period, the NHI employed a "top-up" reimbursement model for stents. Under this system, hospitals received a fixed base payment for any stent procedure, set at a level sufficient to cover the 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 difference between the hospital's posted retail price and 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 in 2009 and a further 18% in 2012). 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, in contrast to the U.S. healthcare system, most physicians in Taiwan—including cardiologists performing stent implantation—are employed by hospitals. Because detailed information on hospital compensation structures and physician–patient interactions is unavailable, it is not feasible to model individual physician selection or the agency relationship between physicians and patients. Accordingly, throughout the paper we abstract from the role of individual physicians in hospital pricing and portfolio management. We further 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 typically an elective procedure; the cost of the procedure itself (except for the out-of-pocket payment required 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.

2.2 Data Sources and Sample Construction

Our empirical analysis relies on a combination of three comprehensive datasets, which together provide the granular information necessary to identify the parameters of our structural model. The 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.

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.[1]

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 retired individuals. 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. This measure serves as a robust proxy for a patient's economic status, which our demand model uses to explain variation 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 its manufacturer and, most importantly, by its generation. As shown in Table 1, the major manufacturers in our sample introduced multiple DES generations over the study period. We consolidate near-simultaneous releases with similar technical specifications and group infrequently used brands into an "Other" category. These product-level data allow us to

[1]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. 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. Further details are provided in the appendix of Jin, Lien and Tao (2025).

define the two key hospital actions in our supply-side model: an *adoption* (when a hospital first offers a brand) and an *upgrade* (when a hospital replaces an existing generation with a newer one). For all analyses, we assume the basic BMS option is available in every hospital's portfolio in all periods.

As detailed in Jin, Lien and Tao (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.

2.3 Descriptive Evidence and Stylized Facts

The data reveal several key stylized facts about the Taipei DES market. These empirical patterns motivate the specific features of our structural model, which must be able to rationalize these observed outcomes.

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. However, the speed and completeness of these transitions vary across brands, and newer generations do not always drive out older ones entirely, suggesting that the preference for newness interacts with brand loyalty, hospital pricing, and other market frictions.

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, we observe substantial price dispersion across brands for clinically similar products. We also observe pricing anomalies, where older generations are sometimes priced higher than newer

Table 1: List of Major Drug-Eluting Stent Models

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 (e.g., Abbott Gen 1 and 2). Some globally available models (e.g., Boston Scientific Gen 3) 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.

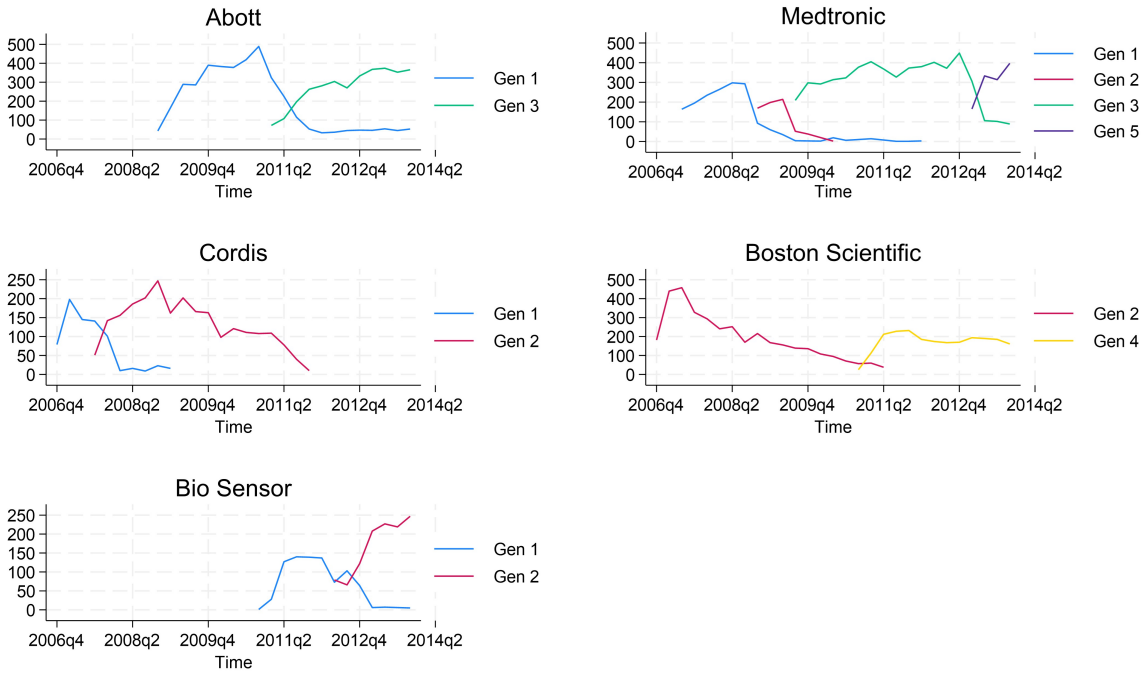


Figure 1: DES Quantity of Different Generations by Brand

ones in the market. This evidence points to significant brand differentiation and, more importantly, frictions in hospitals' portfolio and pricing adjustments, motivating a supply-side model where such adjustments are costly.

Third, the timing of technology adoption and upgrading is highly heterogeneous across hospitals. Figure 3 plots the distribution of delays, measured in quarters from a new product's market introduction to its appearance in a hospital's portfolio. The significant and widely dispersed delays, which result in a right-skewed distribution, are inconsistent with a frictionless market. This heterogeneity is a central feature of the diffusion process and provides strong evidence for the existence of substantial, hospital-specific fixed costs of adjustment, a key component of our supply-side model.

Finally, summary statistics confirm the powerful role of patient price sensitivity. As shown in Table 2, despite the clinical advantages of DES, the free BMS option consistently retains a large share of the market. 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. 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.

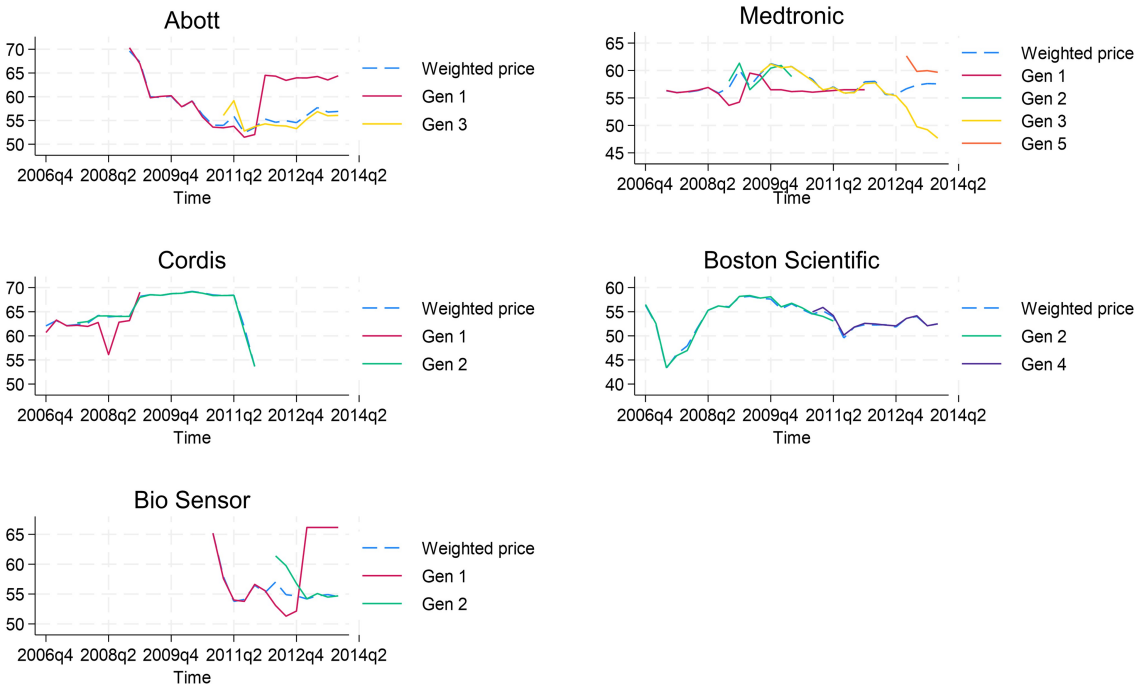
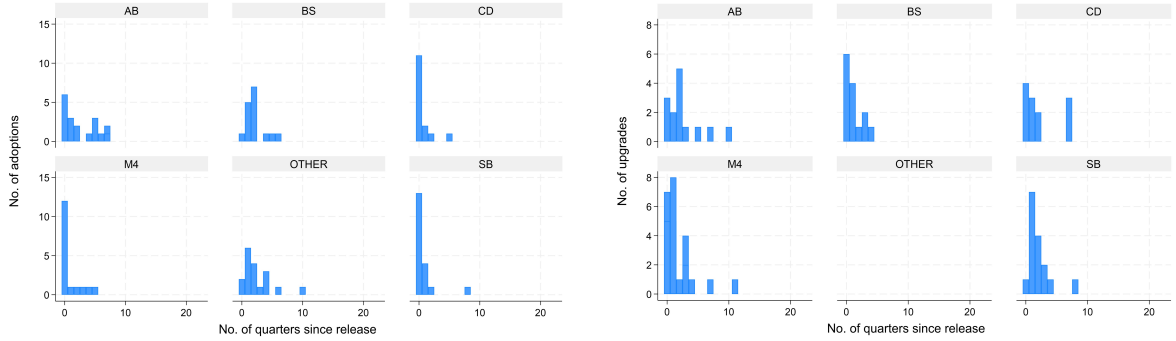


Figure 2: Average Price Across Generations



(a) Adoption Delay

(b) Upgrade Delay

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

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.

Table 2: 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
Percentage of patients using DES per hospital-quarter	40.17%	20.20%	0%	100%

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 manufacturers' 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 ($isDES_m = 0$).[2] 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. 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 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:

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

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:

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

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

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 unobserved product quality captured in the error term, ξ_{hmt}, even after we control for hospital and manufacturer fixed effects (ξ_h and ξ_m). 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 hospital-brand-time combinations 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, consistent

with NHI regulations for medical devices and drugs.[3]. 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:

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

where $\pi_h(\cdot)$ is the expected variable profit from the second-stage pricing equilibrium, $c^M_{ht}(\cdot)$ is the one-time, fixed cost of the portfolio change (either an adoption or an upgrade), and v_{ht} is an i.i.d. Type I extreme value shock private to the hospital.

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

[3]Under this policy, reimbursement is based on the average input cost across hospitals, with an additional 20% allocated as hospital profit.

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.[4] Our static approach is consistent with recent work analyzing complex product-portfolio problems, such as Wollmann (2018) and Olssen and Demirer (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.

$$Pr_t(M'_{ht}|M_{ht}, \mathcal{M}_t) = \frac{\exp\left(V_{ht}(M'_{ht}, M_{ht}, \mathcal{M}_t)\right)}{\sum_{M' \in \mathcal{A}_{ht} \cup \{M_{ht}\}} \exp\left(V_{ht}(M', M_{ht}, \mathcal{M}_t)\right)}, \tag{7}$$

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

4 Estimation Results

This section presents the estimation results for both sides of our structural model. We begin with the demand-side estimates, which reveal rich patterns of patient heterogeneity and provide crucial insights into how preferences for innovation and price sensitivity shape market outcomes. We then turn to the supply-side results, showing how hospitals' strategic decisions regarding technology adoption and pricing respond to these demand patterns and competitive pressures. Together, these estimates provide the foundation for understanding the complex dynamics of medical technology diffusion in the DES market.

4.1 Demand Estimation

The parameters of the demand model are estimated using detailed, individual-level data on patient treatment choices. For each patient in a given quarter, the choice set is comprehensive, encompassing all available Drug-Eluting Stent (DES) and Bare-Metal Stent (BMS) options across every hospital in the Taipei metropolitan area. The model is specified as a mixed logit, which is

[4]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.

particularly well-suited for this context as it allows for rich, unobserved heterogeneity in patient preferences alongside heterogeneity based on observable patient characteristics. We estimate the model parameters via simulated maximum likelihood. Acknowledging that out-of-pocket prices are endogenous—potentially correlated with unobserved product-hospital quality attributes—we employ a control-function approach using the instrumental variables detailed in Section 3 to ensure consistent estimates.

Table 3 presents the estimation results, moving from a standard conditional logit (Column 1) to our preferred mixed logit specification with a full set of interactions (Column 4). The superior fit and refined insights from the latter specification confirm the necessity of modeling both observed and unobserved heterogeneity, which forms the basis for our discussion.

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 the "newest-generation" status. This preference for new technology creates a demand-side incentive for hospitals to update their stent portfolios.

The model further reveals substantial heterogeneity in these preferences across the patient population. The estimated interaction terms show that price sensitivity is significantly attenuated for higher-income patients but is more pronounced for patients with greater comorbidity (a higher Charlson Index). This suggests that the cost of new technology is a more significant consideration for lower-income and sicker patients. The preference for DES technology over BMS also varies, increasing with patient income and age. This observed heterogeneity in preferences is a key feature of the market, indicating that different patient segments evaluate the trade-off between price and technology differently.

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 higher-income 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, the substantial price sensitivity (with own-price elasticities generally ranging from -3 to -1, see Appendix C for details) means that high out-of-pocket costs can significantly limit adoption of new technologies, even when patients value innovation. Second, the strong preference for newest-generation stents creates powerful incentives for hospitals to invest in technology upgrades, but the distribution of these preferences across patient types suggests that the benefits of innovation may not be equally accessible to all patients. To quantify the value patients place on innovation, we calculate their willingness to pay (WTP) and willingness to travel (WTT) for newest-generation stents. For an average patient, these are substantial, at approximately 5,690 NTD and 5.03 kilometers, respectively (see Appendix B for calculation details). These values underscore the powerful market incentives for hospitals to invest in the latest technologies.

4.2 Supply Estimation

Supply-side estimation proceeds in two steps. First, we recover hospital–brand–quarter marginal costs by inverting the Bertrand–Nash first-order conditions using the Column (4) demand estimates. Second, we estimate fixed costs of portfolio adjustment—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 a dynamic adjustment 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.

Table 3: 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)

Variable	(1) Clogit 1	(2) Clogit 2	(3) Mixlogit 1	(4) Mixlogit 2
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

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

Recovering Marginal Costs

Using the demand estimates from Column (4) of Table 3, 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 hospital–brand DES prices, $\mathbf{r}_t$ 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 the mixed-logit estimates using simulation draws and the control-function residuals retained from demand.

Operationally, we construct Δ_{-0t} and Δ_{0t} at the hospital–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; results are insensitive to thresholds between 10^{-6} and 10^{-4}. For hospital–brand–quarter combinations 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. The model includes an unobserved

cost component assumed to follow a Type I extreme value distribution.

We estimate the structural parameters of the cost function via maximum likelihood, using the observed sequence of hospital portfolio choices. 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 4 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, suggesting increasing marginal costs as a hospital's portfolio becomes more technologically advanced.

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 5.30 million NTD, with an interquartile range of 1.11 to 8.84 million NTD. The cost to upgrade an existing product line is considerably lower, with a median of 1.50 million NTD and an interquartile range of 0.59 to 2.89 million NTD. This significant cost differential provides a clear economic explanation for the empirical pattern wherein hospitals are more inclined to incrementally

Table 4: 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	

upgrade their existing technology portfolios than to adopt products from entirely new manufacturers. The magnitude of these costs underscores that portfolio adjustments represent significant strategic investments for hospitals.

Model Validation and Fit

To validate the supply-side model, we assess its ability to replicate key patterns of technology adoption and upgrading observed in the data. A well-fitting model should capture not only the cross-sectional distribution of portfolio adjustments across different product lines but also the aggregate dynamics of diffusion over time. We conduct two validation exercises corresponding to these criteria.

First, we examine the model's cross-sectional fit by comparing the actual number of adoption and upgrade events for each manufacturer with the model-imputed counterparts. The imputed values are calculated by summing the model-predicted probabilities of each event across all hos-

pitals and quarters. Table 5 presents this comparison. The model demonstrates a high degree of accuracy, closely matching the actual event counts for most major brands (e.g., Abbott, Bio Sensor, Medtronic). While minor deviations exist—such as a slight overprediction of adoptions for Cordis and a slight underprediction of upgrades for Boston Scientific—the overall correspondence confirms that the model successfully captures the salient features of cross-brand heterogeneity in portfolio adjustments.

Second, we assess the model's ability to reproduce the temporal patterns of diffusion. Figure 4 plots the actual and imputed time series for the total number of adoptions and upgrades per quarter. The figure shows that the model tracks the historical evolution of these events remarkably well, capturing both the timing and the intensity of portfolio adjustments over the sample period. The close alignment between the imputed and actual series provides strong evidence that the model's underlying parameters are well-identified and that it adequately represents the dynamic incentives facing hospitals.

Taken together, the results from these validation exercises indicate that the estimated model successfully replicates both the cross-sectional and time-series dimensions of hospital portfolio decisions. This provides confidence in the model's structural integrity and supports its use for conducting meaningful counterfactual policy simulations.

Table 5: 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

Note: 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 recovered marginal costs, summarized graphically in Figure 5, provide a final, detailed view of the underlying cost structure of the market. The distribution of costs by manufacturer (panel a) reveals a market average of approximately 50,000 NTD, with modest but systematic variation

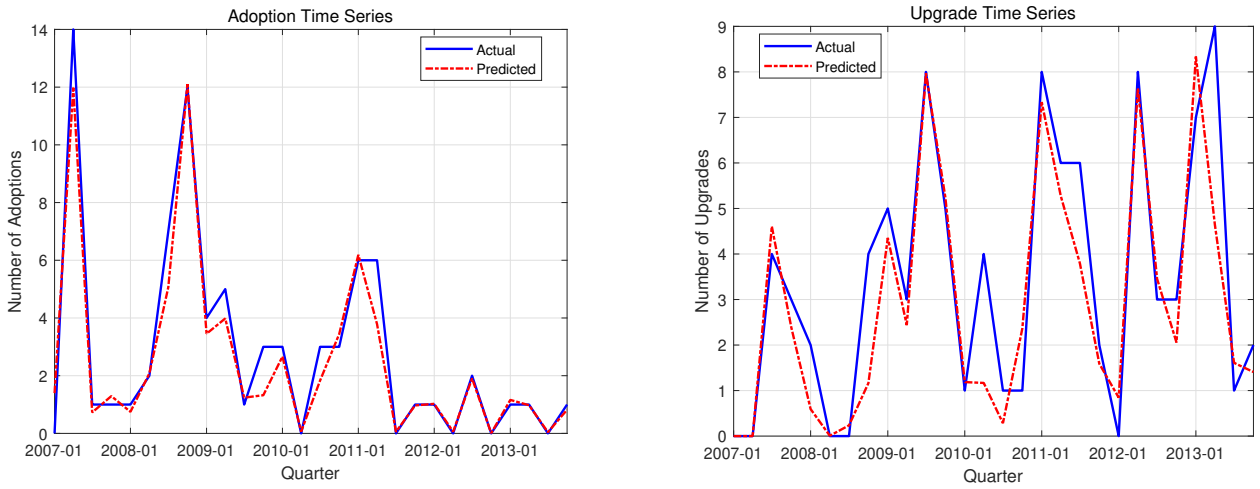


Figure 4: Imputed vs. Actual Frequency of Adoption and Upgrade Over Time

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

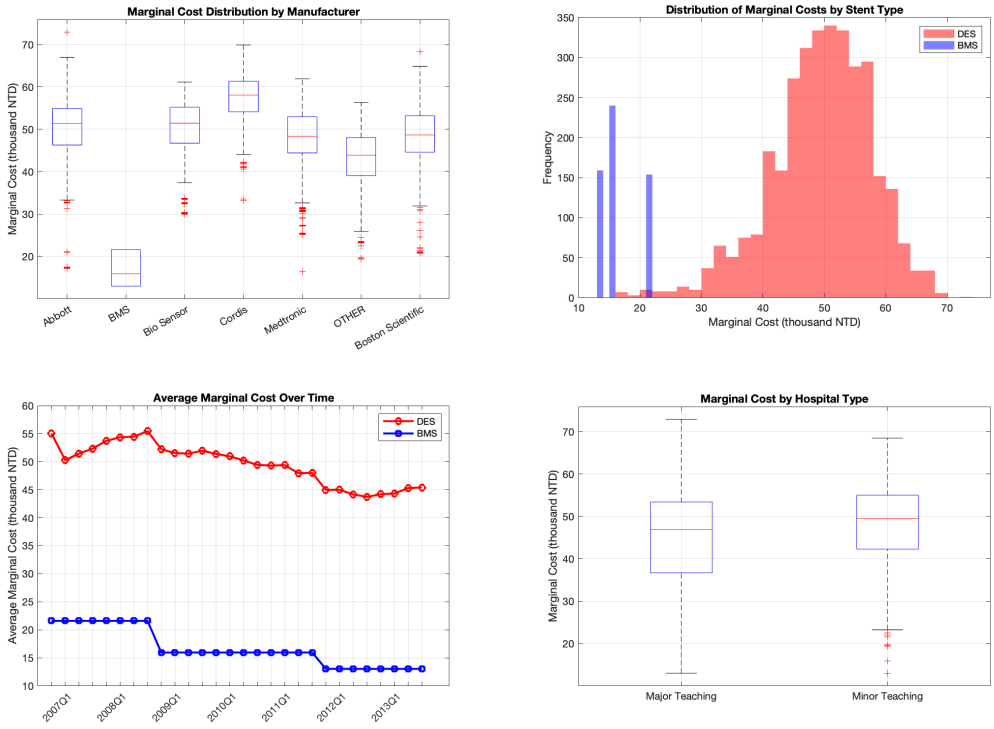


Figure 5: Summary of Recovered Marginal Costs

across firms. The temporal trend (panel c) shows a clear decline in average marginal costs for DES from approximately 55,000 NTD in 2008 to 45,000 NTD in 2012. This trend is consistent with the effects of increased competition from new entrants and a reimbursement policy change in 2009,

both of which likely increased hospitals' bargaining power with upstream suppliers. Furthermore, the cost differential between major and minor teaching hospitals (panel d) confirms the presence of scale or bargaining advantages, providing a direct economic rationale for the observation that larger institutions are often the earliest adopters of new technologies.

Collectively, the demand and supply-side estimates illuminate the central economic forces that jointly shape market outcomes in this industry. On the demand side, a strong patient preference for technological novelty creates a significant commercial incentive for hospitals to offer the latest products. However, this is counterbalanced by substantial price sensitivity, which acts as a considerable barrier to access, particularly for lower-income and clinically more vulnerable patient populations.

These findings set the stage for our counterfactual analysis. The tension between patient demand for innovation and the price sensitivity that limits access, combined with the strategic investment decisions of hospitals facing significant adoption costs, raises important questions about market performance and policy design. The next section uses our estimated structural model to evaluate how alternative market structures and policy interventions could better align private incentives with social welfare in the context of medical technology diffusion.

5 Counterfactual Simulations

The continuous introduction of new DES models by upstream manufacturers presents a fundamental economic challenge: determining whether market mechanisms provide sufficient incentives for downstream adoption while maintaining affordable patient access. 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.

Our approach proceeds 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 supply-side selective contracting to demand-side patient subsidies—quantifying their differential

impacts on technology diffusion, public expenditure, and the distribution of surplus among market participants.

5.1 Market Competition and Patient Demand

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. While upstream manufacturers undertake the primary R&D investments, the realized social value of these innovations depends critically on downstream adoption decisions. In our setting, hospitals function as gatekeepers whose portfolio choices determine both the availability and pricing of new technologies. This intermediation creates a two-sided incentive problem: hospitals must weigh the fixed costs of adoption against uncertain demand, while simultaneously considering competitive dynamics that may erode first-mover advantages through rapid imitation.

Market Structure and Equilibrium Outcomes

To quantify the relationship between market concentration and technology diffusion, we simulate equilibria under alternative competitive environments. We consider three scenarios that span the empirically relevant range: a highly concentrated duopoly ($N = 2$), an intermediate case with moderate competition ($N = 10$), and the baseline competitive market observed in our data ($N = 20$). For scenarios with reduced hospital counts, we randomly select the corresponding number of hospitals from our sample and compute equilibria using the iterative best-response algorithm detailed in Appendix A. Results represent averages across ten independent draws to ensure robustness.[5]

Table 6 presents the equilibrium implications of market structure. Panel A demonstrates the expected pro-competitive effects on pricing and consumer welfare: moving from duopoly to the full competitive market reduces average DES prices by 9.5% (from 65.72 to 59.46 thousand NTD) while more than tripling consumer surplus (from 1.73 to 5.57 million NTD). These magnitudes align with the substantial price elasticities documented in our demand estimates, confirming that

[5]Our treatment of unobserved demand shocks in counterfactual scenarios follows a conservative approach. For hospital-manufacturer-quarter combinations observed in the data, we retain the control-function residuals from our demand estimation. For new combinations arising endogenously in the simulation (e.g., when a hospital adopts a previously unoffered brand), we examined both zero residuals and imputation via a triple fixed-effects specification. The qualitative patterns and quantitative magnitudes prove remarkably stable across these alternatives, with key outcomes differing by less than 1%. We report results based on the zero-residual specification for transparency.

Table 6: Market Outcomes Under Alternative Market Structures

	Number of Hospitals		
	2	10	20
Panel A: Quarterly Market Outcomes			
Consumer surplus (million NTD)	1.73	4.78	5.57
Average patient-paid DES price (thousand NTD)	65.72	60.40	59.46
Hospital profit (million NTD)	26.18	28.47	25.67
Total welfare (CS+Profit, million NTD)	27.91	33.25	31.24
DES price elasticity	-2.26	-2.17	-2.15
Government subsidy (million NTD)	43.97	43.97	43.97
Subsidy allocated to DES	12.26	13.98	12.32
DES utilization rate	28.10%	32.09%	27.89%
Panel B: Technology Adoption and Patient Access			
Portfolio modification probability per hospital:			
Overall	13.36%	7.57%	5.14%
Adoption	4.85%	2.02%	1.67%
Upgrade	12.15%	8.32%	5.70%
Total hospital-model DES combinations	4.86	18.55	31.22
With newest generation	3.26	11.50	17.85
Proportion of newest-generation options	66.95%	61.99%	57.18%
Patients receiving newest-generation DES	22.40%	25.76%	21.09%

Notes: "Portfolio modification probability" represents the frequency of adoption or upgrade events among hospitals where such modifications were feasible. "Overall" indicates the percentage of cases where any portfolio change (adoption or upgrade) occurred.

market power translates directly into higher patient costs.

The producer-side outcomes reveal a more nuanced pattern. Hospital profits follow an inverted-U trajectory, increasing from 26.18 million NTD under duopoly to a maximum of 28.47 million NTD at intermediate concentration ($N = 10$), before declining to 25.67 million NTD in the fully competitive market. This non-monotonicity arises from offsetting forces: moderate competition expands market penetration (DES utilization increases from 28.1% to 32.1%), generating volume-based profit gains that initially dominate the margin compression from lower prices. However, as competition intensifies further, the price erosion effect dominates, reducing equilibrium profits. Notably, total welfare—the sum of consumer and producer surplus—is maximized at the intermediate market structure, suggesting that neither extreme concentration nor perfect competition optimizes social efficiency in this setting.

Panel B reveals the critical trade-off between competition and innovation diffusion that motivates our analysis. The probability of portfolio modification—encompassing both new brand adoption and existing brand upgrades—declines monotonically with market competition, falling from 13.36% in duopoly to 5.14% in the full market. This pattern reflects the erosion of innovation rents: in more competitive environments, the quasi-rents from offering differentiated technologies are rapidly competed away, weakening the business case for costly adoption. While the aggregate variety of offerings mechanically increases with hospital count (from 4.86 to 31.22 distinct hospital-model combinations), the *composition* shifts toward older technologies—the proportion featuring newest-generation models declines from 67.0% to 57.2%.

These portfolio dynamics translate directly into differential patient access to innovation. The proportion of patients receiving newest-generation DES follows an inverted-U pattern, rising from 22.4% under duopoly to peak at 25.8% with moderate competition ($N = 10$), before declining to 21.1% in the fully competitive market. This non-monotonicity crystallizes the central tension in our analysis: while competition enhances affordability through lower prices (the "access effect"), it simultaneously erodes hospitals' incentives to invest in frontier technologies (the "incentive effect"). The data reveal that these opposing forces achieve optimal balance at intermediate market structures, where sufficient competition disciplines pricing without entirely eliminating innovation rents.

Figure 6 decomposes these equilibrium patterns into their constituent mechanisms. Panel (a) captures the first-order welfare effects: consumer surplus rises monotonically with competition due to price reductions, while hospital profits follow the inverted-U trajectory discussed above. The steepest price declines occur in the transition from duopoly to moderate competition, with diminishing effects thereafter—evidence that portfolio differentiation sustains market power even as the number of competitors increases.

Panels (b) through (d) trace the innovation dynamics that underlie our main results. Panel (b) demonstrates how the competition-innovation trade-off shapes patient outcomes: both overall DES utilization and access to newest-generation technologies peak at intermediate concentration ($N \approx 10$). Panel (c) reveals the compositional shift driving this pattern—while total variety mechanically increases with hospital count, the growing gap between total and newest-generation offerings confirms that competitive markets systematically tilt toward older technologies. Panel

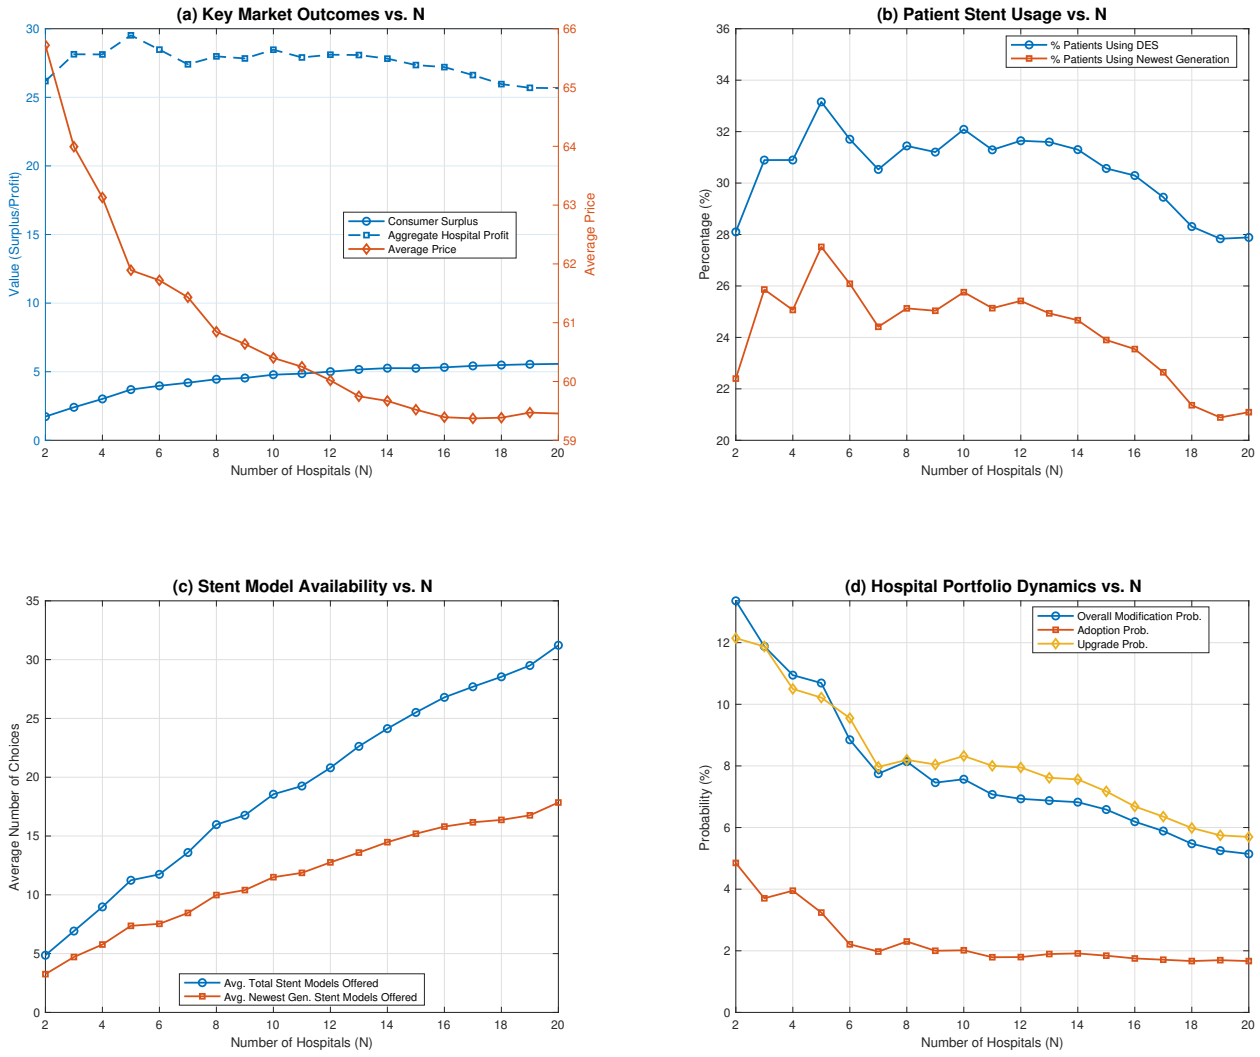


Figure 6: Market Summary for Various Number of Hospitals

(d) identifies the root cause: portfolio modification probabilities decline sharply as competition intensifies, with the steepest drops occurring precisely where business-stealing incentives lose their force. Together, these panels illustrate how market structure fundamentally shapes not just the price and quantity of medical technology, but its vintage and quality distribution across the patient population.

The Role of Patient Preferences for Innovation

The preceding analysis held demand parameters fixed while varying market structure. We now examine the complementary question: how do patient preferences for technological novelty shape equilibrium outcomes across different competitive environments? This question has immediate

policy relevance, as information campaigns or quality reporting initiatives could potentially shift the salience of innovation attributes in patient decision-making.

To isolate this mechanism, we conduct counterfactual simulations that scale the estimated preference parameter for newest-generation technology by a factor ζ. Setting $\zeta = 2$ doubles patients' marginal utility from accessing frontier innovations relative to the baseline ($\zeta = 1$), allowing us to trace how stronger demand-side pull for innovation propagates through the market equilibrium across our three competitive scenarios.

Table 7: Market Outcomes Under Alternative Patient Newness Preferences

	N=2 Hospitals		N=10 Hospitals		N=20 Hospitals	
Newness Preference Factor (ζ)	1 (Baseline)	2	1 (Baseline)	2	1 (Baseline)	2
Panel A: Patient and Product Outcomes						
Consumer Surplus (million NTD)	1.73	2.03	4.78	5.07	5.57	5.94
DES Utilization Rate (%)	28.10%	37.18%	32.09%	40.21%	27.89%	38.72%
Newest-Generation DES Utilization Rate (%)	22.40%	32.77%	25.76%	36.00%	21.09%	33.29%
Avg. DES Models Available per Hospital	4.86	5.02	18.55	18.30	31.22	33.56
Avg. Newest-Gen. Models per Hospital	3.26	3.65	11.50	12.30	17.85	20.44
Panel B: Hospital Outcomes						
Hospital Profit (million NTD)	26.18	30.93	28.47	33.19	25.67	32.13
Avg. Patient-Paid Price (thousand NTD)	65.72	67.06	60.40	60.92	59.46	59.35
Hospital Portfolio Modification Rate (%)	13.36%	15.27%	7.57%	8.19%	5.14%	6.58%
Panel C: NHI Subsidy Outcomes						
Total Subsidy (million NTD)	43.97	43.97	43.97	43.97	43.97	43.97
Subsidy Allocated to DES (million NTD)	12.26	16.24	13.98	17.61	12.32	17.05

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. All outcomes are averaged across simulations. Units are New Taiwan Dollars (NTD).

Table 7 reveals how heightened patient preferences for innovation fundamentally reshape market equilibria, with effects that vary systematically across competitive environments. The results demonstrate that while stronger demand for novelty universally accelerates diffusion, the distribution of resulting surplus—and crucially, the affordability of innovation—depends critically on market structure.

Panel A quantifies the demand-side impacts. Doubling patients' preference for technological novelty ($\zeta = 2$) substantially increases both overall DES utilization and access to frontier technologies across all market structures. The effects are particularly striking for newest-generation adoption: even in the highly competitive baseline market ($N = 20$), utilization rises from 21.1% to

33.3%, representing a 58% increase. These gains in innovation access translate directly into higher consumer surplus, though the magnitude varies with competition—consumers in more competitive markets capture a larger share of the value created by their own stronger preferences.

The supply-side responses documented in Panel B reveal how hospitals strategically adapt to stronger innovation demand. Portfolio modification rates increase uniformly (rising from 5.14% to 6.58% in the $N = 20$ case), confirming that demand-pull mechanisms can partially offset the innovation-dampening effects of competition identified earlier. However, the pricing responses diverge sharply by market structure. In concentrated markets ($N = 2$), hospitals exploit enhanced willingness-to-pay by raising prices substantially (from 65.72 to 67.06 thousand NTD). As competition intensifies, this price extraction diminishes—in the fully competitive market, prices remain essentially flat despite doubled innovation preferences, indicating complete pass-through of value to consumers.

Panel C illuminates an important fiscal externality of innovation preferences. While total government expenditure remains fixed by design, the allocation shifts dramatically toward DES as utilization increases. The share of subsidies flowing to DES rises by approximately 40% across all market structures when preferences double, highlighting how demand-side factors can reshape public spending patterns even under budget-neutral policies.

These findings synthesize into a crucial insight for policy design: patient preferences for innovation create powerful diffusion incentives, but without sufficient competition, these preferences may paradoxically reduce access by enabling price increases that offset the welfare gains from better technology. This fundamental tension—between harnessing demand-pull innovation incentives and maintaining affordability—motivates our examination of targeted reimbursement mechanisms that could better align private incentives with social objectives.

5.2 Alternative Reimbursement Designs for Enhanced Innovation Diffusion

Our preceding analyses reveal a fundamental tension in the market for medical innovation. Neither competitive forces nor patient demand alone can achieve both rapid technological diffusion and broad affordability. Competition reduces prices but weakens hospitals' incentives to invest in new technologies, while strong patient preferences for innovation accelerate adoption but enable price increases in concentrated markets. This inherent trade-off suggests an important role for targeted

policy interventions that better align private incentives with social welfare objectives.

We evaluate three policy instruments, 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 vulnerable populations 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.

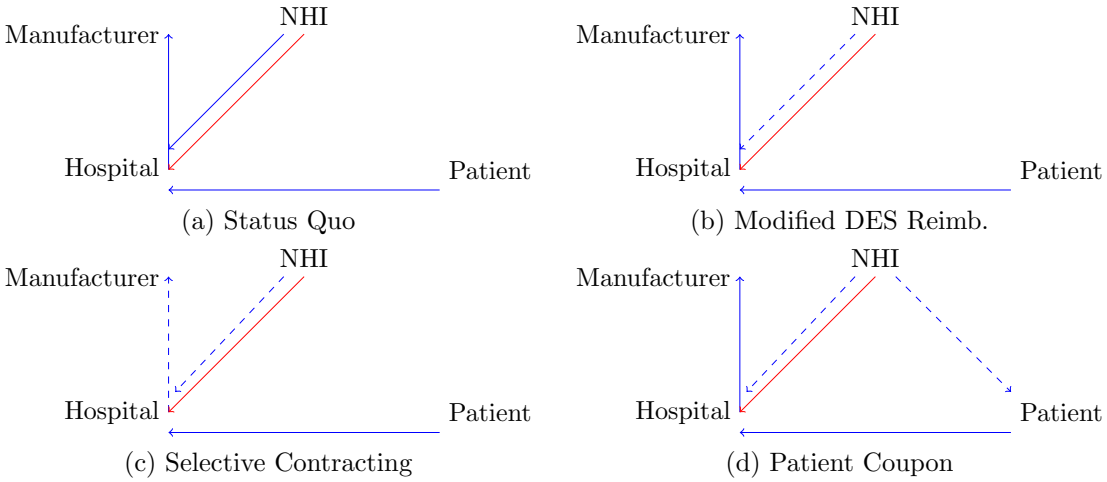


Figure 7: Comparison of Alternative NHI Reimbursement Designs

Figure 7 illustrates the payment flows under each policy design, with arrows representing financial transfers among market participants. Solid lines indicate baseline payments while dashed lines mark policy-induced changes. Red arrows denote BMS-related payments and blue arrows represent DES-specific transfers.

Panel (a) shows the status quo arrangement where the NHI provides a uniform reimbursement r_t calibrated to BMS costs. Hospitals charge patients an additional top-up fee for DES to cover the cost differential. While this design controls public spending, it creates substantial out-of-pocket burdens that may inefficiently limit access to innovation.

Panel (b) depicts the DES-specific reimbursement policy, which introduces a multiplicative factor r^d applied only to DES procedures, yielding a reimbursement of $r^d \cdot r_t$. When $r^d > 1$, the

enhanced payment reduces the coverage gap, enabling hospitals to lower patient prices.

Panel (c) illustrates selective contracting, where the NHI negotiates directly with chosen manufacturers to secure wholesale discounts (captured by $r^c < 1$) in exchange for exclusive reimbursement eligibility. This approach reduces hospitals' input costs for contracted brands, potentially lowering patient prices while steering demand toward selected products. The exclusivity requirement transforms the competitive landscape by creating a two-tier market of reimbursable and non-reimbursable technologies.

Panel (d) shows the patient coupon program, which provides direct subsidies to eligible patients (such as those in the lowest income decile) when they choose DES. Unlike supply-side interventions, this approach preserves market pricing while selectively reducing financial barriers for targeted populations. The program's effectiveness depends on how hospitals respond strategically to the demand shift from subsidized patients.

Alternative Design #1: Technology-Specific Reimbursement Rates

The simplest policy intervention adjusts the reimbursement rate for DES independently from BMS. We parameterize this through a multiplicative factor r^d applied to DES procedures, resulting in a reimbursement of $r^d \cdot r_t$ for DES while maintaining r_t for BMS. This approach allows policymakers to directly control the subsidy level for innovation, with $r^d < 1$ reducing support and $r^d > 1$ enhancing coverage.

Table 8 presents the market outcomes when doubling the DES reimbursement ($r^d = 2$). The enhanced subsidy generates substantial price reductions that vary systematically with market structure. In the duopoly setting, patient prices fall by 22%, while in more competitive markets ($N = 10$ and $N = 20$), the reduction exceeds 32%. This pattern reflects stronger pass-through in competitive environments, where hospitals have less ability to capture the subsidy as additional profit.

These price reductions dramatically expand technology adoption. DES utilization increases by 17 percentage points across all market structures, rising from approximately 28-32% to 45-49% of procedures. Access to newest-generation technologies shows similar gains, with utilization rates increasing by 16 percentage points. The uniformity of these quantity responses across different competitive environments suggests that affordability, rather than availability or clinical factors, represents the binding constraint on innovation diffusion.

While the policy succeeds in expanding access, it generates divergent welfare effects. Consumer surplus increases across all markets due to lower prices and expanded choice. Hospital profits rise even more dramatically—by 76% in duopoly and 43% in moderately competitive markets—as the enhanced reimbursement more than compensates for price reductions. These profit increases strengthen adoption incentives, with portfolio modification rates rising by approximately 50% across market structures.

The fiscal burden, however, is substantial. Total NHI expenditure increases by roughly 50%, while DES-specific spending more than triples. The government pays over 50,000 NTD per additional DES patient—a cost that may exceed the clinical value of upgrading from BMS to DES for marginal patients. This high fiscal cost motivates our examination of more targeted interventions that might achieve similar diffusion gains more efficiently.

Table 8: Market Outcomes Under Alternative DES Reimbursement Factors (r^d)

	N=2		N=10		N=20	
Reimbursement Factor	$r^d = 1$	$r^d = 2$	$r^d = 1$	$r^d = 2$	$r^d = 1$	$r^d = 2$
Consumer-related:						
Consumer Surplus (million NTD)	1.73	2.33	4.78	5.43	5.57	6.36
% of Patients Using DES	28.10%	45.51%	32.09%	49.27%	27.89%	49.18%
% of Patients Using Newest DES	22.40%	38.93%	25.76%	41.90%	21.09%	39.69%
Aggregate DES Choices	4.86	6.06	18.55	19.11	31.22	35.96
Aggregate Newest DES Choices	3.26	4.50	11.50	12.73	17.85	22.11
Hospital-related:						
Hospital Profit (million NTD)	26.18	46.07	28.47	40.66	25.67	38.65
Avg. Patient-Paid Price (thousand NTD)	65.72	51.23	60.40	40.95	59.46	39.09
Avg. Portfolio Modification Probability	13.36%	21.41%	7.57%	9.03%	5.14%	7.61%
NHI-related:						
Total NHI Subsidy (million NTD)	43.97	64.02	43.97	65.80	43.97	65.82
DES-Specific NHI Subsidy (million NTD)	12.26	40.10	13.98	43.65	12.32	43.70

Notes: The table compares outcomes for $r^d = 1$ (baseline) and $r^d = 2$ (doubled DES reimbursement). Consumer Surplus, Hospital Profit, Total NHI Subsidy, and DES-Specific NHI Subsidy are reported in millions of NTD; Average Patient-Paid Price is in thousands of NTD; usage rates and portfolio modification probabilities are percentages. Aggregate choices represent the total number of hospital-product combinations offered in the market.

Alternative Design #2: Selective Contracting

Selective contracting leverages the insurer's bargaining power to negotiate favorable terms with specific manufacturers. Unlike the broad reimbursement adjustment examined above, this approach targets specific products and manufacturers, potentially achieving more efficient outcomes through market segmentation. This mechanism, common in pharmaceutical benefit management, combines exclusive reimbursement with negotiated discounts. We examine whether it can simultaneously reduce patient costs, maintain adoption incentives, satisfy manufacturer participation constraints, and control public spending.

We model selective contracting through two complementary instruments. First, the NHI negotiates wholesale discounts with chosen manufacturers—in our simulations, the two market leaders Abbott and Medtronic—reducing hospital input costs to $r^c \cdot c_{hmt}$ where $r^c < 1$ represents the discount factor. Second, the NHI grants exclusive reimbursement eligibility to contracted brands, effectively creating a two-tier market where non-contracted DES become out-of-network products that patients must fully finance.

Table 9 presents outcomes under a representative selective contracting arrangement combining a 50% wholesale discount ($r^c = 0.5$) with baseline reimbursement rates ($r^d = 1$). The policy fundamentally restructures the market through two channels. First, it creates dramatic price differentiation: contracted DES prices fall by 40% (from 57,000 to 34,000 NTD in the $N = 20$ market), while non-contracted products become marginally more expensive as they lose competitiveness. Second, this price advantage drives substantial demand reallocation, with overall DES utilization rising from 24% to 38% in competitive markets.

The welfare effects reveal why selective contracting can achieve a "quadruple win." Consumers benefit from both lower prices and increased access to innovation—consumer surplus rises by 8-13% across market structures. Hospitals paradoxically profit despite offering lower prices, as the wholesale discounts exceed their pass-through to patients, expanding margins on higher volumes. The government maintains budget neutrality: while DES utilization increases substantially, the negotiated discounts offset the volume growth, keeping total expenditure stable. Finally, contracted manufacturers benefit from dramatic market share gains that more than compensate for per-unit price reductions.

The policy's distributional consequences merit attention. Non-contracted manufacturers face effective market exclusion, losing most of their reimbursed patient base. This concentration of market power raises potential dynamic concerns about long-term innovation incentives, though our static analysis shows that hospitals' portfolio modification rates for contracted products actually increase substantially (from 4% to 5% for targeted brands in competitive markets). The policy thus appears to channel existing innovation more efficiently rather than fundamentally altering innovation production.

Table 9: Market Outcomes under Selective Contracting Policy Designs

	N=2 Firms		N=10 Firms		N=20 Firms	
Reimbursement factor (r^d)	1 (Targeted)	1 (Targeted)	1 (Targeted)	1 (Targeted)	1 (Targeted)	1 (Targeted)
Cost discount factor (r^c)	1 (Targeted)	0.5 (Targeted)	1 (Targeted)	0.5 (Targeted)	1 (Targeted)	0.5 (Targeted)
Consumer-related:						
Consumer Surplus (million NTD)	1.67	2.40	4.60	5.20	5.47	5.91
% of DES usage	26.54%	47.15%	26.14%	43.55%	24.34%	37.79%
% of patients using newest generation	21.60%	42.16%	21.90%	39.96%	19.96%	33.38%
Hospital-related:						
Profit (million NTD)	23.66	43.71	24.59	36.80	23.82	32.43
Avg. Patient-Paid Price (thousand NTD)	78.08	69.72	76.90	71.04	76.31	72.25
Targeted DES	62.80	44.01	58.79	36.20	57.07	34.40
Non-targeted DES	84.68	90.22	80.63	81.71	79.84	80.18
NHI-related:						
Subsidy (million NTD)	40.32	41.91	39.54	40.88	39.28	40.17
Targeted DES	7.85	18.38	6.79	15.80	5.84	12.64
BMS	32.48	23.53	32.75	25.08	33.44	27.53
Market Structure & Dynamics:						
Total hospital-model DES choices	5.36	6.83	18.86	20.38	34.48	36.37
w/ newest generation	3.49	4.78	11.67	13.24	20.59	22.48
Portfolio Modification Probability per Hospital	15.19%	24.44%	7.67%	9.67%	6.48%	7.41%
Targeted DES	12.30%	32.12%	5.04%	7.21%	4.11%	4.77%
Non-targeted DES	4.79%	5.36%	2.86%	2.95%	2.58%	2.94%

Table 10: Preferred Direction of Policy Parameters

	NHI	Consumers	Hospitals	Manufacturers
Cost Factor, r^c	Lower	Lower	Lower	Mid-range
Reimbursement Factor, r^d	Lower	Higher	Higher	Higher

Building on these results, Table 10 summarizes directional preferences over the two policy parameters—r^c (the wholesale cost factor for targeted DES) and r^d (the reimbursement factor for targeted DES)—for each stakeholder. The NHI, seeking to curb expenditure, prefers lower r^c (greater negotiated discounts) and lower r^d (smaller per-case payments). Consumers prefer lower r^c to reduce prices through pass-through and higher r^d to increase coverage and reduce out-of-pocket payments. Hospitals, driven by margins, also favor lower r^c and higher r^d.

The sharpest conflict is over r^d: the insurer's cost-minimizing preference for a lower reimbursement opposes the preferences of consumers and hospitals. By contrast, preferences over r^c are more aligned—everyone benefits from lower input costs—though the division of the gains between patients and hospitals depends on competitive conditions and pricing conduct.

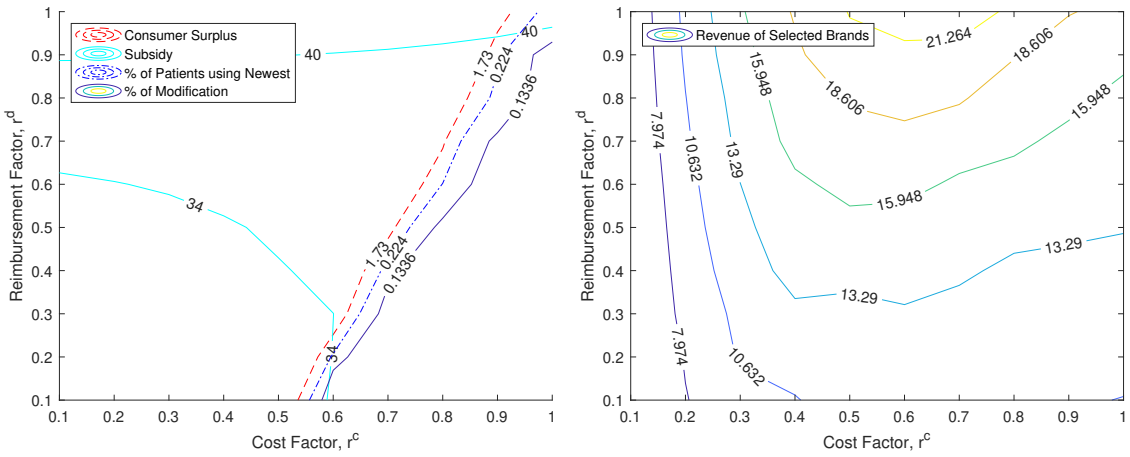
For selective contracting to be feasible, targeted manufacturers must also prefer participation over the status quo. Although their costs are unobserved, their revenue under selective contracting reflects a trade-off between price and volume: deeper discounts reduce per-unit revenue but increase quantity via exclusivity and lower patient prices. Because market share expands when non-targeted brands are excluded, the revenue effect of r^c is generally non-monotonic. Manufacturers often prefer an interior discount that balances volume gains against price cuts.[6]

To characterize the joint effects of the policy levers, we simulate outcomes on a grid with $r^d \in [0.1, 1]$ and $r^c \in [0.1, 1]$ (using increments of 0.1 for both r^c and r^d shown in the figures). Figure 8 plots iso-curves for four outcomes—consumer surplus, hospital portfolio-modification probability, total NHI expenditure, and the share using the newest-generation DES—for markets with $N = 2$, $N = 10$, and $N = 20$.

Consistent with Table 10, consumer surplus increases when reimbursement is more generous (higher r^d) and wholesale costs are lower (lower r^c), corresponding to the upper-left regions of each panel. Total NHI spending falls as both levers are tightened (lower-left). Using the status quo (Original in Table 9) as a benchmark, the contour lines identify policy combinations that weakly improve patient and hospital outcomes while keeping NHI expenditure at or below baseline. In the $N = 2$ market (Figure 8(a)), the status-quo reference values are consumer surplus of 1.73 million NTD, newest-DES use of 22.40%, a portfolio-modification probability of 13.36%, and total NHI expenditure of 43.97 million NTD. Regions to the "north-west" of the consumer-surplus, newest-DES, and portfolio-modification iso-curves and to the "south-west" of the 43.97 million NTD expenditure iso-curve yield improvements for patients and hospitals without higher public spending.

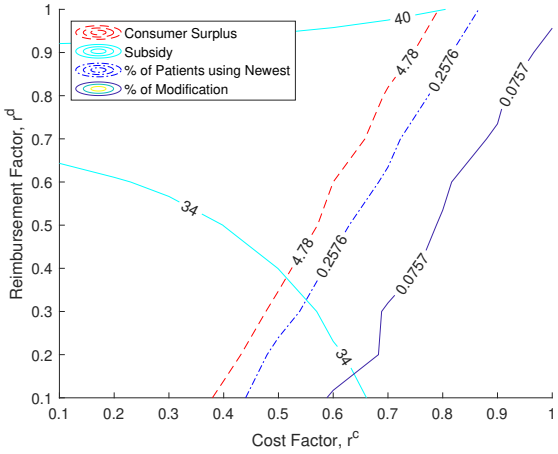
Figure 8(b) shows iso-revenue curves for targeted manufacturers in the $N = 2$ market. Under moderate discounts (lower r^c) combined with maintained reimbursement (higher r^d), exclusivity-

[6] In practice, the NHI could operationalize this scheme by announcing coverage for a limited set of DES brands and inviting manufacturers to bid on r^c. The participation constraint we study is a necessary condition; an auction mechanism could further improve the NHI's terms.

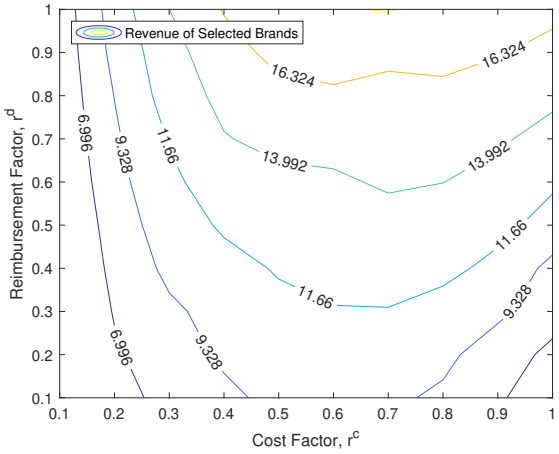


(a) NHI, Patient, Hospital for $N = 2$

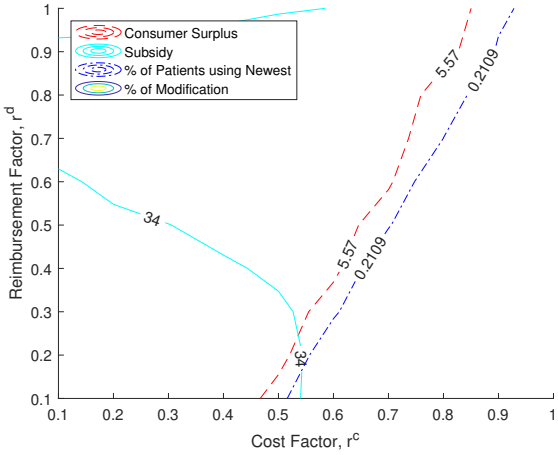
(b) Manufacturer for $N = 2$



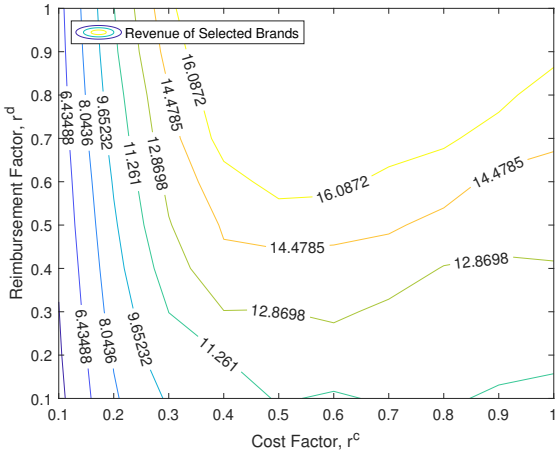
(c) NHI, Patient, Hospital for $N = 10$



(d) Manufacturer for $N = 10$



(e) NHI, Patient, Hospital for $N = 20$



(f) Manufacturer for $N = 20$

Figure 8: Policy Outcomes under Selective Contracting: Varying (r^c, r^d)

driven market-share gains and demand expansion can raise revenue above the status quo. Taken together, panels (a) and (b) indicate that selective contracting can attain a "quadruple-win"—benefiting

patients, hospitals, the NHI (through cost control), and targeted manufacturers—while non-targeted manufacturers lose share.

Similar patterns obtain in the $N = 10$ and $N = 20$ markets (Figures 8(c)–(f)). Using status-quo benchmarks from Table 6—for $N = 10$: consumer surplus 4.78 million NTD, newest-DES use 25.76%, portfolio modification 7.57%, NHI expenditure 43.97 million NTD; for $N = 20$: 5.57 million NTD, 21.09%, 5.14%, and 43.97 million NTD—Figures 8(c) and (e) delineate (r^c, r^d) regions that improve patient and hospital outcomes relative to the status quo while respecting the NHI budget constraint. Figures 8(d) and (f) likewise show sizable sets where targeted manufacturers' revenue weakly exceeds baseline, reinforcing the potential for a quadruple-win in more competitive markets as well, subject to the choice of r^c and r^d.[7]

Alternative Design #3: Patient Coupon Program

The third policy instrument addresses distributional concerns through targeted demand-side subsidies. While the previous approaches manipulate supply-side incentives through reimbursement rates or input costs, patient coupons directly reduce out-of-pocket burdens for vulnerable populations. We model a means-tested program providing 30,000 NTD coupons ($r_{\text{coupon}} = 0.5$) to patients in the lowest income decile. These coupons are redeemed at the point of care, reducing patient payments without altering hospital reimbursements. This design preserves market pricing mechanisms while selectively expanding access for those most constrained by affordability.

Figure 9 illustrates the distributional challenge motivating this intervention. The figure shows stark income-based disparities in technology access under the status quo. Patients in the lowest income decile (annual income approximately 0.686 million NTD) have DES utilization rates of just 19.87%, compared to 28.80% for higher-income patients—a 45% gap in access. This disparity is even more pronounced for newest-generation technologies, reflecting the interaction between income constraints and preferences for innovation documented in our demand estimates. Such baseline inequities raise both efficiency and equity concerns, as clinical need does not necessarily correlate with ability to pay.

[7] In the $N = 20$ market, all (r^c, r^d) combinations shown yield portfolio-modification rates above the 5.14% status-quo benchmark, so that iso-curve is omitted from Figure 8(e). Concentrating reimbursement on targeted DES increases the profitability of newer products and strengthens adoption/upgrade incentives across the depicted parameter range.

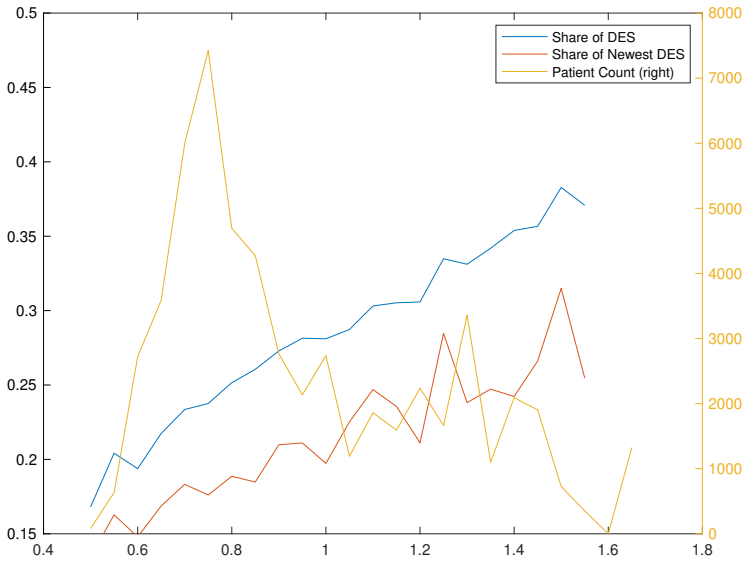


Figure 9: Distribution of Income and DES Uptake under the Status Quo (N=20)

Table 11 quantifies the program's equilibrium effects across market structures. The coupon successfully narrows access gaps: DES utilization among eligible low-income patients increases by 30-33% across all market structures (from 23.34% to 30.33% in duopoly, from 23.99% to 31.49% with moderate competition, and from 19.87% to 26.42% in the fully competitive market). This substantial improvement occurs without generating significant spillovers to ineligible patients, whose utilization rates remain virtually unchanged.

The absence of significant price responses reveals an important feature of targeted subsidies. In our model, each hospital sets one price per product for all patients, and thus in theory it is possible that a hospital may raise the list price of DES facing all patients when the patient coupon helps to increase some patients' willingness to pay for DES. However, low-income patients are more price sensitive than high-income patients and intensive hospital competition may undermine the profitability of price increase. Our simulation suggests that the list price change is slightly positive when $N = 2$, zero when $N = 10$, and slightly negative when $N = 20$ (the status quo). In any case, changes in average DES price remain minimal (less than 0.1%) across all market structures.

Similarly, portfolio modification rates are essentially unchanged in concentrated and competitive markets (13.36% when $N = 2$ and 5.14% when $N = 20$), with only marginal increases in moderately competitive settings (from 7.57% to 7.65% when $N = 10$). This pattern suggests that the incremental demand from 10% of patients is insufficient to alter hospitals' technology adoption

Table 11: Market Outcomes of the Patient Coupon Program

	$N = 2$ Hospitals		$N = 10$ Hospitals		$N = 20$ Hospitals	
Coupon Value Factor (r_{coupon})	0	0.5	0	0.5	0	0.5
Patient-Related Outcomes:						
Consumer Surplus (million NTD)	1.73	1.75	4.78	4.81	5.57	5.59
Hospital-Related Outcomes:						
Hospital Profit (million NTD)	26.18	26.48	28.47	28.99	25.67	26.03
Avg. Patient-Paid Price (thousand NTD)	65.72	65.73	60.40	60.40	59.46	59.42
Portfolio Modification Rate (%)	13.36%	13.36%	7.57%	7.65%	5.14%	5.14%
DES Usage Rate (%):						
Low-Income Patients	23.34%	30.33%	23.99%	31.49%	19.87%	26.42%
Other Patients	28.52%	28.52%	32.80%	32.97%	28.80%	28.82%
NHI-Related Outcomes:						
NHI Hospital Reimbursement (million NTD)	43.97	43.97	43.97	43.97	43.97	43.97
Total Coupon Cost (million NTD)	0	1.97	0.00	2.09	0.00	1.77
Total NHI Expenditure (million NTD)	43.97	45.94	43.97	46.06	43.97	45.74

Notes: The table compares outcomes with no patient coupon ($r_{\text{coupon}} = 0$, status quo) versus a coupon program ($r_{\text{coupon}} = 0.5$). Eligible low-income patients (lowest 10% decile) receive a coupon valued at $0.5 \times 60\text{k} = 30\text{k}$ NTD, which directly offsets out-of-pocket payment for DES. Consumer Surplus, Hospital Profit, NHI Hospital Reimbursement, Total Coupon Cost, and Total NHI Expenditure are in millions of NTD; Avg. Patient-Paid Price is in thousands of NTD; Portfolio Modification and DES Usage Rates are percentages.

calculus, particularly when margins are either protected by market power or eroded by competition.

From a fiscal perspective, the program's targeted nature ensures cost containment. Total government expenditure increases by only the direct coupon costs: approximately 1.8-2.1 million NTD across market structures, representing less than 5% of baseline NHI spending. This modest fiscal burden achieves meaningful equity improvements, reducing the access gap between income groups by nearly one-third. However, the program's limited market-wide effects raise questions about whether addressing distributional concerns alone is sufficient when the underlying market failures in innovation diffusion remain unresolved.

6 Conclusion

This paper examines how hospital intermediaries shape the diffusion of medical innovations in Taiwan's drug-eluting stent market. Our structural analysis reveals fundamental trade-offs that market forces alone cannot resolve. Competition reduces prices but weakens hospitals' incentives

to adopt frontier technologies, while patient preferences for innovation enable rent extraction in concentrated markets. Neither perfect competition nor monopoly maximizes welfare; intermediate market structures optimally balance price discipline against innovation incentives. Our policy simulations demonstrate that selective contracting can achieve a "quadruple-win"—benefiting patients, hospitals, contracted manufacturers, and the government—through vertical coordination that creates value rather than merely redistributing surplus. Targeted patient subsidies complement this approach by addressing distributional concerns, though with limited market-wide effects.

These findings have broader implications for healthcare markets where intermediaries control technology access. Traditional policy approaches focusing solely on upstream R&D incentives or downstream competition 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, 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, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt. 2005. "Competition and Innovation: an Inverted-U Relationship." *The Quarterly Journal of Economics*, 120(2): 701–728.

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.

Bergman, Alon, Matthew Grennan, and Ashley Swanson. 2021. "Medical Device Firm Payments To Physicians Exceed What Drug Companies Pay Physicians, Target Surgical Specialists: Study examines payments medical device firms pay to physicians." *Health Affairs*, 40(4): 603–612.

Bergman, Alon, Matthew Grennan, and Ashley Swanson. 2022. "Lobbying Physicians: Payments from Industry and Hospital Procurement of Medical Devices." National Bureau of Economic Research w29583, Cambridge, MA.

Berry, Steven T., and Joel Waldfogel. 1999. "Free Entry and Social Inefficiency in Radio Broadcasting." *The RAND Journal of Economics*, 30(3): 397.

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

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*. Vol. 2 of *Handbook of Health Economics*, , ed. Mark V. Pauly, Thomas G. Mcguire and Pedro P. Barros, 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.

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

Ching, Andrew T., Tülin Erdem, and Michael P. Keane. 2013. "Invited Paper—Learning Models: An Assessment of Progress, Challenges, and New Developments." *Marketing Science*, 32(6): 913–938.

Collard-Wexler, Allan, Matthew Grennan, and Andrew Steck. 2024. "Regulating New Product Testing: the FDA vs. the Invisible Hand." *Working Paper*.

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

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

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 Thousands of Cost-Effectiveness Studies." *working paper*.

Garthwaite, Craig, Christopher Ody, and Amanda Starc. 2022. "Endogenous quality investments in the U.S. hospital market." *Journal of Health Economics*.

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).

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.

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

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

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

Mankiw, N. Gregory, and Michael D. Whinston. 1986. "Free Entry and Social Inefficiency." *The RAND Journal of Economics*, 17(1): 48.

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.

McClellan, Mark B., David T. Feinberg, Geisinger, Peter B. Bach, Paul Chew, Omada Health, Patrick Conway, Nick Leschly, Bluebird Bio, Greg Marchand, Michael A. Mussallem, Edwards Lifesciences Corporation, and Dorothy Teeter. 2017. "Chandra." *NAM Perspectives*, 7(3).

Nexon, David, and Stephen Ubl. 2010. "Implications Of Health Reform For The Medical Technology Industry." *Health affairs (Project Hope)*, 29: 1325–9.

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

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

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.

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

Yurukoglu, Ali, Eli Liebman, and David B. Ridley. 2017. "The Role of Government Reimbursement in Drug Shortages." *American Economic Journal: Economic Policy*, 9(2): 348–382.

1

Appendix A Details on 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 $C(M_{ht}, M_{h,t-1})$ 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,

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

Update hospital h's portfolio to the maximizer,

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

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 B Willingness to Pay and Willingness to Travel Calculation

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 Table 3, Column 4.

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):

$$\text{WTP}_{\text{avg}} = -\frac{0.636 + (-0.029 \times 0.915) + (-0.175 \times 1.007) + (0.040 \times 0.766) + (0.001 \times 65.304)}{-0.025 + (-0.004 \times 0.915) + (0.004 \times 1.007) + (-0.004 \times 0.766) + (-0.001 \times 65.304)} \tag{8}$$

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:

$$\text{WTT}_{\text{avg}} = -\frac{0.636 + (-0.029 \times 0.915) + (-0.175 \times 1.007) + (0.040 \times 0.766) + (0.001 \times 65.304)}{-0.095 + (0.005 \times 0.915) + (0.026 \times 1.007) + (-0.009 \times 0.766) + (0.0001 \times 65.304)} \tag{9}$$

This yields approximately 5.03 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 C Demand Elasticity Calculation

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 mixed logit parameter estimates reported in Table 3.

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:

$$\eta_{ijt} = \beta_i^p \cdot p_{jt} \cdot (1 - \hat{P}_{ijt}) \tag{10}$$

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:

$$\beta_i^p = \beta_0^p + \beta_1^p \cdot \text{Charlson}_i + \beta_2^p \cdot \text{Income}_i + \beta_3^p \cdot \text{Male}_i + \beta_4^p \cdot \text{Age}_i \tag{11}$$

The empirical distribution of calculated elasticities demonstrates uniformly elastic demand throughout the market. As illustrated in Figure A.1, 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.

Substantial heterogeneity emerges along both product and patient dimensions. Table A.1 reveals systematic variation across manufacturers, with Cordis products exhibiting markedly higher price sensitivity (mean elasticity of -2.31) relative to established market leaders such as Abbott and Medtronic (both approximating -2.0). These differential elasticity patterns illuminate why selective contracting arrangements generate asymmetric welfare effects across manufacturers, with higher-elasticity brands experiencing more pronounced quantity responses to exclusion.

Patient-level heterogeneity proves even more striking. The comprehensive evidence in Tables A.2, A.3, and A.4 establishes clear gradients in price responsiveness: elasticity magnitudes increase monotonically with comorbidity burden (Charlson Index) and patient age, while exhibiting a strong negative correlation with income level. Most notably, patients in the lowest income decile display price sensitivities approximately 50% greater than their highest-income counterparts, highlighting the distributional burden imposed by the current cost-sharing arrangement. This systematic heterogeneity in price response provides the economic foundation for understanding how carefully designed targeted interventions—whether through selective contracting or patient-specific

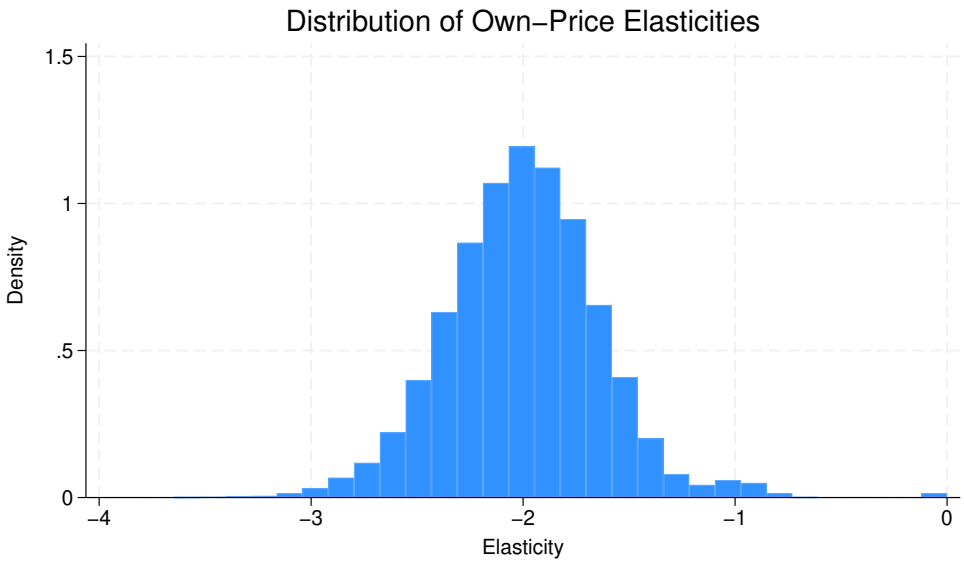


Figure A.1: Elasticity Histogram

Notes: The histogram shows the distribution of individual patient-level own-price elasticities for DES. Most elasticity estimates fall within the -3 to -1 range, centered around a mean of approximately -2, and the distribution is fairly symmetric.

subsidies—can simultaneously enhance allocative efficiency and improve equity outcomes, as our counterfactual analyses demonstrate.

Beyond brand-level differences, the analysis reveals systematic variations in price sensitivity correlated with patient-specific characteristics. As detailed in Tables A.2, A.3, and A.4, patients with greater health burdens, those with lower incomes, males, and older individuals tend to exhibit more elastic demand. Specifically, Table A.2 shows a clear positive relationship between the Charlson Comorbidity Index and the magnitude of price elasticity, indicating that patients with more comorbidities are more responsive to price changes. Similarly, Table A.3 demonstrates that patients in lower income brackets are, on average, more price-sensitive than their higher-income counterparts. Finally, Table A.4 indicates that males exhibit slightly higher price sensitivity compared to females, and that price elasticity tends to increase with patient age. These demographic and health-related patterns in price responsiveness are important for understanding the distributional consequences of pricing and reimbursement policies.

Table A.1: Elasticity by Brand

Brand	Mean Elasticity	Std. Dev.	Count
Abbott	-1.998	0.376	7,206
Bio Sensor	-1.964	0.365	2,002
Cordis	-2.306	0.313	3,004
Medtronic	-2.021	0.298	9,083
OTHER	-1.873	0.302	2,404
Boston Scientific	-1.861	0.395	5,662

Notes: The table reports mean own-price elasticity, standard deviation, and observation counts for each DES brand. Cordis exhibits the highest average price sensitivity (mean elasticity of -2.306), while Boston Scientific and OTHER have the least sensitive demand on average.

Table A.2: Elasticity by Charlson Index

Charlson Index	Mean Elasticity	Count
0	-1.840	11,766
1	-2.038	13,000
2	-2.243	3,700
3	-2.422	681
4	-2.577	162
5	-2.736	52

Notes: The table shows mean own-price elasticity and patient counts grouped by Charlson Comorbidity Index. Price sensitivity increases monotonically with the Charlson Index, indicating that patients with more comorbidities are more responsive to price changes.

Table A.3: Elasticity by Income Category

Income Category	Mean Elasticity	Count
1 (Lowest)	-2.106	2,164
2	-2.077	2,540
3	-2.061	2,599
4	-2.038	3,098
5	-2.032	3,042
6	-2.000	3,017
7	-2.024	3,044
8	-1.947	3,000
9	-1.897	3,287
10 (Highest)	-1.876	3,570

Notes: This table presents mean own-price elasticity and patient counts across ten income categories. Patients in the lowest income category exhibit the highest price sensitivity (elasticity of -2.106), with a general trend of decreasing sensitivity as income rises.

Table A.4: Elasticity by Gender and Age Quartile

Panel A: By Gender		
Gender	Mean Elasticity	Count
Female	-1.910	6,758
Male	-2.024	22,603
Panel B: By Age Quartile		
Age Group	Mean Elasticity	Count
Quartile 1	-1.893	7,026
Quartile 2	-1.966	7,891
Quartile 3	-2.035	7,760
Quartile 4	-2.103	6,684

Notes: Panel A reports mean elasticity by gender, showing males (-2.024) are slightly more price-sensitive than females (-1.910). Panel B shows mean elasticity by age quartile, indicating that price sensitivity increases with age, with the oldest quartile being the most price-sensitive (-2.103).