A STRUCTURAL MODEL OF ORGANIZATIONAL BUYING FOR B2B MARKETS: INNOVATION ADOPTION WITH SHARE OF WALLET CONTRACTS

By

Navid Mojir and K. Sudhir

November 2021

COWLES FOUNDATION DISCUSSION PAPER NO. 2315

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

A Structural Model of Organizational Buying for B2B Markets: Innovation Adoption with Share of Wallet Contracts*

Navid Mojir
Harvard Business School

K. Sudhir
Yale School of Management

October 2021

*We thank the participants at marketing seminars at Chinese University of Hong Kong, Colorado, Cornell, Duke, HBS, HKUST, Illinois, Maryland, Minnesota, Notre Dame, Rochester, Syracuse, Texas A&M, UCSD, USC, UT Dallas, Virginia, the marketing camp at Dartmouth, the 2018 ISBM Academic Conference, the 2017 KU Leuven Manufacturer-Retailer Symposium and the 2017 Marketing Science Conference. This paper builds upon dissertation proposals that received the 2016 MSI Alden G. Clayton Award, the 2016 ISMS Dissertation Proposal Award and the 2015 ISBM Dissertation Proposal Award.

A Structural Model of Organizational Buying for B2B Markets: Innovation Diffusion with Share of Wallet Contracts

The paper develops the first structural model of organizational buying to study innovation diffusion in a B2B market. Our model is particularly applicable for *routinized exchange relationships*, whereby centralized buyers periodically evaluate and choose contracts, then downstream users order items on contracted terms. The model captures different utility tradeoffs for users and buyers while accounting for how buyer and user choices interact to impact user adoption/usage and buyer contracting. Further, the paper considers the dynamics induced by share of wallet (SOW) pricing contracts, commonly used in B2B markets to reward customer loyalty with discounts for buying more than a threshold share from a supplier. We assemble novel panel data on surgeon usage, SOW contracts, contract switching, and hospital characteristics. We find two segments of hospitals in terms of the relative power of surgeons and buyers: a buyer-centric and a surgeon-centric segment. Further, innovations diffuse faster in teaching hospitals and when surgeries are concentrated among a few surgeons. Finally, we answer such questions as: Should the marketer focus on push (buyer-focused) or pull (user-focused) strategies? Do SOW contracts hurt the innovations of smaller firms? Surprisingly, we find that the contracts can help speed the diffusion of major innovations from smaller players.

Keywords: Organizational Buying Behavior, healthcare marketing, B2B Markets, B2B Innovation, New Product Diffusion, New Product Adoption

INTRODUCTION

The business-to-business (B2B) sector in the United States has a transactional value of $10.7 trillion, roughly the same as the value of the business-to-consumer (B2C) sector. Further, B2B accounts for a staggering 90% of all e-commerce in the United States (Lilien 2016). Yet the share of academic research on B2B in leading marketing journals is very small at 3.4% (LaPlaca and Katrichis 2009). Though there is a long history of organizational buying models in marketing (e.g., Robinson et al. 1967, Webster and Wind 1972, Sheth 1973) and a rich literature testing various components of models using survey data (see integrative reviews in Lewin and Johnston 1996, Sheth 1996), econometric work using secondary data on organizational buying is scarce. Lilien (2016) attributes the relative scarcity and "misallocation of academic resources on B2C versus B2B" to three challenges: (i) Modeling organizational choice in B2B is more complex than individual choice, considering multiple stakeholders who differ in their objectives and priorities in B2B settings; (ii) Modeling B2B markets needs more institutional immersion and knowledge; and (iii) B2B panel data with choices of multiple stakeholders are harder to obtain.[1]

We study organizational buying and innovation diffusion in a duopoly B2B market: the U.S. hospital market for a surgical device. In studying this market, we address the three challenges mentioned above. First, we develop a tractable structural econometric model of organizational buying wherein multiple stakeholders with different objectives and preferences interact and influence each other to generate the buying outcomes within the organization. The model is particularly applicable in B2B buying modes characterized as *routinized exchange relationships* (RERs) (Grewal et al. 2015). Here centralized buyers *periodically* evaluate and choose contracts for price and service terms, and downstream users order items conditional on these terms during the period of contract. Second, we lay out various institutional aspects of buying in hospital markets (particularly medical devices), clarify which aspects are relevant for answering our research questions, and discuss how we account for the relevant aspects in our modeling approach. Such institutional

[1]Even with transactional choice data on B2B markets, past research (e.g., Bruno et al. 2012, Grennan 2013, Zhang et al. 2014) models the data as individual choices without considering the institutional reality of multiple stakeholders.

aspects include contract structure and the role of intermediary organizations for group purchasing common in health care markets. Finally, we assemble a rich and novel B2B panel data set on decisions made by different agents within each hospital. We obtain proprietary data on usage share of various products by surgeons and complement that information by hand-coding the terms and conditions of contracts available to each hospital plus the hospital administrator's contract choice. The data collection is critical to model and capture the multi-stakeholder nature of this B2B market.

In our hospital setting, suppliers use "share of wallet" (SOW) contracts, common in B2B markets, to induce and reward customer loyalty to suppliers.[2] Under SOW contracts, discounts are tied to achieving a certain threshold SOW level with the supplier.[3] In practice, customers contract with one supplier, committing to a certain wallet share and receiving a price discount for meeting the share threshold. Sellers use audits to verify share compliance. SOW discounts are more common than quantity discounts in B2B markets as quantity is typically outside the buyer's control. In business markets, quantity is usually a function of the total output of the firm and determined by strategic and tactical decisions made by senior management or other departments such as marketing and finance—not by the buyer. The buyer only decides from whom to buy products, and therefore the SOW.[4] We therefore model buyer choice of supplier to sign SOW contracts with and surgeon choice of product to use (from contracted and non-contracted suppliers).

Accounting for these two critical market features—namely, the multi-agent nature of B2B buying and routinized buying under SOW contracts, we answer the following questions. First, which

[2]The literature discusses share of wallet contracts in many B2B settings across a range of industries, where antitrust cases have drawn attention to this practice. Examples include contracts between: chip manufacturers and PC manufacturers, airlines and travel agents, tire manufacturers and retailers, and medical device suppliers and hospitals (e.g., see Majumdar et al. 2005). In general, the courts in the United States have upheld the practice, although, in Europe, courts have been less tolerant of discounts designed to induce loyalty, especially from large dominant firms.

[3]The economics literature (e.g., Calzolari and Denicolò 2013) uses the term "market-share discount" to describe "discounts that depend on the seller's share of a customer's total category purchase." They are also called "loyalty discounts" or "fidelity rebates" (see, for example, Tom et al. 2000; Majumdar and Shaffer 2009). We prefer the term SOW discounts relative to (i) market share discounts, because it clarifies that the discounts are for each customer's (not market) share, and (ii) loyalty discounts/fidelity rebates, as loyalty and fidelity can be operationalized in other ways than SOW.

[4]SOW contracts are also more practical, as suppliers can offer the same discount to large/small and high/low growth firms for achieving the same target SOW without perceptions of unfairness. Further, to enjoy the loyalty-inducing effect, the supplier does not need to predict the category demand for each hospital-quarter and update the quantity threshold in every period; however, the disadvantage is the additional data burden of using third parties to audit annual shares and/or use of field representatives to assess compliance with share thresholds.

stakeholder, the buyer or user, has more influence on the other in an organization's adoption and usage of an innovation? Is the relative influence of buyers (or users) the same across all hospitals? Second, given that buyers and users have different objectives and preferences, and the hospital product adoption and usage is the outcome of their interactions, when should suppliers focus on buyers (push tactics) versus users (pull tactics) to drive product adoption and usage? Third, given that theoretical research has shown that SOW contracts typically favor larger dominant suppliers (Majumdar and Shaffer 2009, Inderst and Shaffer 2010), do SOW pricing contracts suppress innovation adoption from smaller players? This question is important, not only for smaller firms, but also for regulators because medical device innovations disproportionately come from smaller firms (Roberts 1988). The question remains open because existing theoretical papers only consider unitary buyers and not the multi-stakeholder nature of B2B buying. For example, a smaller supplier with a superior but expensive innovation may gain user adoption by targeting price-insensitive users, leaving the hospital unable to meet the dominant firm's SOW threshold. This violation of SOW threshold of the dominant supplier can potentially trigger faster contract switching to the smaller innovator and may speed up innovation diffusion.

Overall, we make three key modeling contributions. First, our model captures the multi-stakeholder nature of a B2B buying center involving RERs in an internally consistent manner. To the best of our knowledge, there are no econometric models of organizational buying for B2B markets; the multi-stakeholder core of our framework can serve as a workhorse model for other B2B markets, especially in RER settings. As discussed, the structural model for a B2B market needs to consider the different (and potentially divergent) objectives and preferences of stakeholders and how they interact and influence each other to determine organizational choice. A key element of our contribution is how we model the buyer-seller mutual influence on each other in an internally consistent way. Although users can use products from both contracted and non-contracted suppliers, disutility arises due to the inconvenience of using off-contract products. This nudges users to choose the contracted product. But user dissatisfaction also impacts buyer utility and can restrain the extent of misalignment between buyer choice and user preference. A key innovation that we introduce in our model is the construct of "inclusive value loss" (IVL) to parsimoniously capture the total user dissatisfaction with a buyer's contract choice. The relative weight of the user-buyer influences on each other determines who has greater impact on product adoption in a given organization. Such disaggregate modeling of each stakeholder's choices and their influence on each other's utility helps with better predictions of organizational buying outcomes in counterfactual marketing environments.

Second, our empirical model captures the nuances of how SOW contracts affect an organization's choice in terms of product adoption, usage, and contracting. Consider a duopoly market with one large supplier and a smaller innovating supplier that launches a new product. Given the empirical context, we assume there is no other entry. The buyer has expectations about price evolution and usage share evolution based on contracts with the two suppliers. Buyer payoffs arise from both costs to the firm due to prices and user dissatisfaction with chosen contracts. At any given time t, the buyer evaluates the expected discounted sum of payoffs over time from staying with the current supplier or switching contracts to the innovating supplier based on the realized value of the state variables and expectations of future states. The buyer chooses the option with the higher payoff. If the buyer does not switch in period t, a decision must be made about switching contracts in period $t + 1$. Once the buyer switches to the innovating supplier in period t, the buyer "stops" and has no further contracting decision to make (as there is no additional entry). Therefore, we can characterize the buyer's decision as an optimal stopping problem.[5]

Finally, our model provides a structural framework to study the diffusion of *disposables* in a trial-repeat context. Thus far, structural models of new product diffusion have focused only on *durables* in B2C markets involving only single purchases (e.g., Song and Chintagunta 2003, Nair 2007). But the trial-repeat diffusion models (e.g., Hahn et al. 1994, Lilien et al. 1981) tend to be focused on B2C markets and are non-structural in that they merely fit diffusion curves. Further, the dynamics of innovation adoption in structural models for durables usually stem from the tradeoff

[5] There is no such dynamic tradeoff for surgeons. Consequently, we model surgeon trial and usage decisions as myopic, conditional on buyer contracting decisions. While surgeons collectively influence buyers, each surgeon has little direct impact on the buyer's utility, as any surgeon's share of surgeries in a hospital is small ($< 2\%$).

between buying a product now versus waiting for future lower price or higher quality. In contrast, for disposables in our trial-repeat context, the dynamics stem from the SOW pricing contracts and the multi-agent nature of buying. The buyer's decision when to switch contracts is not driven by falling prices but rather the wait for enough users (surgeons) to adopt the new product such that the total disutility from purchasing cost and user dissatisfaction is minimized.

Our key findings reveal two distinct hospital segments: a large "buyer-centric" segment and a smaller "user-centric" segment. For the focal innovation, the user-centric hospital segment drives initial adoption. We also find that innovation diffuses faster in hospitals with more concentrated surgeries among surgeons (i.e., few surgeons perform a large share of the hospital's surgeries). Using the model estimates, we performed two counterfactuals to answer our other research questions. The first counterfactual compares the differential effects of focusing on buyers by lowering price (i.e., push strategy) relative to improving the surgeon's perception of the innovation's value (i.e., pull strategy). We find that price reduction (push) has limited impact on the innovation diffusion as the dominant supplier's product still has significant price advantage; however, focusing on users (pull) is more effective—especially in the user-centric hospital segment where sufficient surgeon adoption triggers contract switching by the buyer to the innovator to obtain the SOW discount. After focusing initially on the surgeons (pull) until a threshold share is reached, it is optimal to shift the focus to buyers through prices (push) strategy. The counterfactual also reveals that the smaller innovator should focus on "major" innovations to make the pull strategy effective. Overall, by modeling the multi-stakeholder aspect of organizational buying, we obtain deeper insights about the proper tactics and timing of focusing on each stakeholder to maximize innovation adoption.

The second counterfactual addresses the question of whether SOW contracts hurt the diffusion of innovations introduced by smaller firms. As discussed before, SOW contracts require the consolidation of shares with one seller to obtain discounts. Though the conventional wisdom based on current theory suggests that this favors the larger, dominant firm, we find that SOW discounts *can* speed up the adoption of innovations from smaller firms if they have a superior innovation—even when it is more expensive. This diffusion acceleration is driven by hospitals in the user-centric segment because they adopt based on the innovation's superior benefits, rather than price, and are less influenced by buyers. Once the SOW threshold needed for the larger firm is breached, buyers accept the reality of the superior innovation and switch contracts. The contract switching further accelerates user adoption of the innovation.

The rest of the paper is organized as follows. After outlining the links to related literature and providing institutional background for the market, we describe details of the data and provide model-free evidence to support the modeling assumptions. We then present the model, the estimation approach, and details of model identification. This is followed by the estimation results and counterfactual analysis. We conclude with suggestions for future research.

RELATED LITERATURE

The paper is related to three key streams of literature: organizational buying and group decision-making, share of wallet pricing contracts, and new product adoption and diffusion.

Organizational Buying and Group Decision-Making

The pioneering conceptual models in organizational buying behavior (Robinson et al. 1967, Webster and Wind 1972 and Sheth 1973) spawned a rich literature focusing on concepts such as the buying center and decision making processes within the buying center, make versus buy, and modified rebuy. Grewal et al. (2015) provide a more recent taxonomy of B2B buying modes: (i) routinized exchange relationship (RER); (ii) transactional buying operations (TBO); and (iii) organic buying relationships (OBR). As discussed earlier, our model and empirical setting addresses the common RER buying mode. In his review and assessment of the organizational buying literature, Sheth (1996) notes that much of this research tends to be descriptive "survey research wherein a key informant is asked to either role play or recall incidents of specific organizational buying decisions" and calls for more analytical approaches that are common in B2C markets. Lilien (2016) explains the challenges of building and estimating econometric models of organizational buying. Our paper addresses this gap in the literature.

We note some exceptions that model organizational choice behavior using observational data. For example, to understand new product acceptance criteria, judgments, and choice, Rao and McLaughlin (1989) and Sudhir and Rao (2006) use contemporaneous data obtained from an internal decision-making form that buyers completed. Such internal organizational decision-making data are unusual and typically hard to obtain for researchers but can be valuable to understand organizational buying behavior. More commonly, researchers have used transactional data on choice in B2B markets and modeled them like B2C markets. For example, Bruno et al. (2012) and Zhang et al. (2014) use B2B transactional choice data to focus on reference price effects and latent, non-stationary buying, respectively. But these studies tend to be descriptive in nature and model buyer behavior as if choices were made by individuals.

In terms of studies on medical device markets, Laczniak (1979) employed Webster and Wind's model to analyze the buying behavior of medical devices in a sample of Midwestern hospitals using semi-structured, in-depth interviews and questionnaires. Moon and Tikoo (2002) used surveys of buyers and doctors to study commonalities and differences among buyers and users. Finally, Grennan (2013) uses observational data to estimate a structural model of bargaining between hospitals and suppliers but treats the negotiation as between unitary decision-makers. Grennan (2013) assumes that heterogeneity in bargaining power drives price variation. In contrast, we leverage our data on SOW contract availability through GPOs for different hospitals and find that SOW thresholds and usage share explain most of the price variation across and within hospitals in the context of our markets.

Our work is also related to the literature on group decision-making. One stream of research in this area measures the influence of different group members on the team's decision using stated preference or influence data, but these data have well-documented biases (e.g., Corfman 1989, Turk and Bell 1972). Others use conjoint methods to measure the preferences and relative influences of different group members. Arora and Allenby (1999) and Arora (2006) study preferences and influences of members within family units (e.g., husband-wife, parent-child) using individual and joint conjoint data. Aribarg et al. (2002) seek insight into the preference revision and concession process among parents and teenagers using a conjoint data collection procedure on individual and joint choices interspersed with an intermediate information exchange stage. In contrast, we use observational data about users' and buyers' choices within an organization over time to structurally model their preferences and influences. Unlike stated preference or conjoint data, our joint model of users and buyers reflects actual organizational choices in a real-world setting and thus illustrates how to extend revealed preference modeling used in B2C markets to B2B settings.

Share of Wallet Contracts

Share of wallet (SOW) pricing contracts are commonly used in B2B markets to encourage customer loyalty. Suppliers offer discounts to customers if they purchase above a particular threshold share in the category from the supplier. SOW discounts are preferred to quantity discounts in B2B markets to induce loyalty because quantity is driven by the downstream demand and usually outside the buyer's control, whereas the buyer controls SOW.

There is a small theoretical literature on SOW discounts in a manufacturer-retailer setting, but no empirical work. Mills (2010) studies how a retailer's downstream selling effort is affected by SOW discounts. Inderst and Shaffer (2010) find that SOW contracts favor a dominant supplier and can help increase its profit. Calzolari and Denicolò (2013) compare SOW discounts in competitive settings with exclusive contracts. These papers abstract away from the multi-stakeholder nature of the buying organizations, however, so there is little guidance in the current theory on equilibrium outcomes when users' decisions moderate buyers' contract choices. Further, this literature does not consider the dynamics of new product adoption and diffusion, which shed light on push-pull strategies over time. These two factors explain the divergence between our findings and existing results that SOW contracts favor dominant players.

New Product Adoption and Diffusion

The current paper is closely related to new product diffusion models and structural models of new product adoption. Early papers in this literature focused on fitting the diffusion curve at an aggregate level without attending to the micro-foundations of adoption. While most diffusion papers focused on first purchase (e.g., Fourt and Woodlock 1960, Mansfield 1961, and Bass 1969), a smaller literature extended these first purchase models to include trial and repeat purchases (e.g., Hahn et al. 1994, Lilien et al. 1981).[6]

Recently, structural models have been used to provide a micro-foundation for new product adoption of consumer durables, focusing only on first purchase diffusion. In these models, falling prices or increasing quality over time drives the diffusion (e.g., Song and Chintagunta 2003, Nair 2007). In contrast, prices and quality for consumables are typically fairly stable. The rising share of the consumable innovation over time (diffusion) is due to the growth in the "installed base" of users who have tried the product and are more likely to continue using it. Our work provides a structural framework for diffusion of consumables. Thus, our diffusion model is more similar to ongoing prescriptions by doctors (Hahn et al. 1994), whereby once a doctor or consumer has been acquired, ongoing prescriptions drive sales. There is a qualitative difference in that context, however, because a physician's prescriptions may vary by patient, whereas a surgeon typically uses the same product line for all surgeries.

Further, in the context of our category, the buyer's decision to switch contracts involves balancing the total cost to the hospital against disutility from user dissatisfaction. This balancing takes into account the buyer's expectation of not only future prices but also surgeons' usage behavior. Thus, beyond the differences in diffusion due to the consumable nature of the product category, the forward-looking buyer's tradeoffs are also novel to the literature on structural modeling of innovation diffusion.

INSTITUTIONAL DETAILS

We study a surgical device category involving instruments that help surgeons handle internal tissues and organs during surgeries. The category is sold through hospitals and has annual revenues of over $700 million. Two suppliers compete in the category. The market leader (Supplier A) has one

[6]Horsky (1990) is a notable exception in that diffusion is explained in terms of such underlying primitives as income, time use, and marketing mix variables such as price.

product line and 60% market share; the smaller Supplier B with the remaining 40% market share has one existing product line and introduced a new product line during the period of the study as easier-to-use and more efficient than the existing product lines. The paper focuses on the adoption and diffusion of this innovation from Supplier B.

Each product line has a set of tools (each tool is an SKU). Different surgeries require different subsets of tools from the product line; some tools are used more commonly than others. The tools are single-use disposable items that can be used for only one surgery. Surgeons have preferences between these product lines due to training and experience but may be willing to switch if they believe a superior product line is easier to use (with some training), more efficient or more effective (e.g., in reducing complications). As a physician preference item (PPI) category, surgeons can choose which product line to use for their surgeries. While it is technically feasible to use items across multiple product lines for different surgeries, enough differences exist in usage that in practice surgeons use only one product line at any point in time.

A hospital administrator (hereafter buyer) signs a contract with one of the two suppliers. Contracts involve three parties: suppliers, Group Purchasing Organizations (GPOs), and hospitals. Contracting terms are negotiated between the suppliers and GPOs and renegotiated every 1-3 years. Each hospital is typically a member of multiple GPOs. In each category, the buyer chooses a supplier and "best" contract available through its GPOs. Hospitals become members of GPOs to (i) obtain good contract terms (from GPO's consolidated buying power) and (ii) reduce contracting costs for supplies across a large number of categories (Schneller and Smeltzer 2006). Since any given category is a small share of the overall contracting value through GPOs, GPO membership can be treated as exogenous for a category.[7]

Each contract consists of three elements: an off-contract (or list) price, a discounted price, and a SOW threshold.[8] The SOW threshold is defined as the supplier's share of a hospital's total purchase in a category (all product lines and SKUs in the category) and can differ across suppliers; in our market, the larger Supplier A has a threshold of 80%, whereas the smaller Supplier B

[7]The focal category value of $700 million is 0.5% of the device market of $160 billion (Gravelle and Lowry 2015).

[8]In addition, discounts vary by spend thresholds. We describe these below and account for them in our estimation.

has a threshold of 60% (the same threshold for all GPO contracts). A hospital that signs up for a contract with a supplier *and* meets the SOW threshold pays the discounted price, otherwise it pays the higher off-contract price. The contracts also stipulate "service" levels for surgeon support. The support includes (i) surgeon assistance from supplier representatives with product usage in operating rooms (Moed and Israel 2017, O'Connor et al. 2016) and (ii) priority delivery of infrequently used tools (Montgomery and Schneller 2007). Such contract-tied support induces more trials by surgeons soon after the hospital switches contracts to the innovating supplier; the buyer can thus induce a "bump" in trials by contracting. Further, suppliers (especially the ones without contracts) provide greater support to surgeons with a higher share of surgeries at a hospital. This induces the high-share surgeons to try the innovation earlier.

There are nine GPOs and 201 hospitals in our sample. The average number of GPO affiliations per hospital is 2.6. But membership is concentrated among the top GPOs. For example, the top three GPOs for Supplier A have affiliations with 91%, 85%, and 69% of hospitals. The average number of hospital members for each of these GPOs is 164. Further, 94% of supplier contracts are concentrated among the top three GPOs. The off-contract price differences across GPOs are relatively small. On average, the larger Supplier A with higher SOW threshold of 80% offers larger discounts; the *average baseline* discounts in this category for the top 3 GPOs are 45%, 42%, and 41% respectively. The corresponding figures for supplier B with the 60% SOW threshold are 36%, 34%, and 33%. The discounts do vary over time.

As noted earlier, discounts also vary by hospital spend. Larger hospitals with greater spends receive higher discounts for meeting SOW thresholds. For example, Supplier B offers 42% discount through the largest GPO for $950K+ spend; it offers discounts of 37% and 39% for $750K+ spend through the second and third largest GPOs. But these spend thresholds are set such that smaller hospitals cannot obtain discounts meant for larger hospitals by merely shifting SOW. Effectively, spend criteria serve as a "hospital size proxy" for third-degree price discrimination.

Given the above background, we treat GPO memberships at each hospital as exogenous in the model because they are based on considerations beyond the category.[9] We also treat spend threshold based discounts as different contracts available to hospitals of different sizes during model estimation because choice among spend thresholds is outside the buyer's control.

DATA AND MODEL-FREE EVIDENCE

We assembled data in cooperation with various departments of Supplier B and from third-party providers. We obtained Supplier B's sales to each hospital for nineteen quarters from its marketing department. We procured data on each hospital's contract status and each GPO's contract terms and conditions from Supplier B's commercial analytics department. Supplier B's competitive analytics group provided estimates of each hospital's total sales potential in the category and GPO affiliations in addition to Supplier A's GPO contracts in the category. We then supplemented these data on prices, shares, and contracts with data on hospital and surgeon characteristics from third-party data providers, specifically HMS and IMS Health.

As this is a duopolistic market, given quarterly sales of Supplier B and the total "sales potential," we compute the quarterly shares/sales of Supplier A at each hospital. Each contract has information about the off-contract (i.e., list) price, SOW threshold, spend requirements, and discounts from each supplier. For Supplier B, we observe the actual prices paid, given its GPO contract and threshold. For Supplier A, we infer the price by the hospital in each quarter based on knowledge of hospital GPO memberships, who the hospital has contracted with, contract details on off-contract prices, and discounts in that quarter. We exclude 7 hospitals who directly negotiate contracts with each supplier and do not use GPO contracts. More details about the construction of Supplier A's price sequence are provided in the Web Appendix.

For hospital characteristics, we have data on (i) whether the hospital is a teaching hospital; (ii) number of surgeons in the category; and (iii) number of surgeries performed by each surgeon. We do not observe product line choices by individual surgeons, only the aggregate share of each product line in each hospital. To account for differential support for higher share surgeons, we

[9]We show further model-free evidence in support of this assumption later.

Table 1: Per-surgery prices by contract across hospital-quarters.

Product	SOW Contract	Mean($)	SD($)	Min($)	Max($)	NObs
Supplier A	Yes	582	165	431	962	3256
Supplier A	No	1196	287	647	1749	563
Supplier B old product	Yes	1258	436	411	2144	563
Supplier B old product	No	1770	201	1198	2198	3256
Supplier B innovation	Yes	1273	210	723	2283	474
Supplier B innovation	No	2070	204	1346	2411	2541

Note: The maximum price with contract is close to maximum price without contract because hospitals occasionally do not meet the share threshold for discount. This situation is relatively infrequent and has limited impact on the mean price with contracts, so the mean reflects discounted price.

Figure 1: Average per-surgery prices

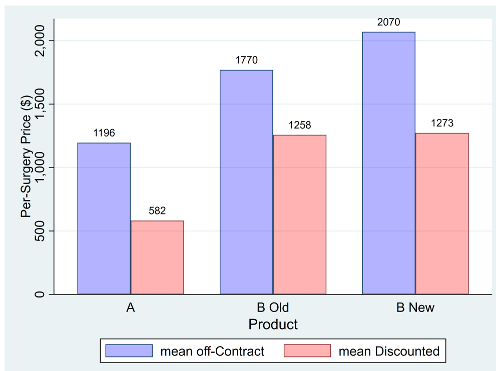

Note: This figure shows the average per-surgery price for each of the three products in the market. Each product has two prices: an off-contract price, which is the price that the hospital will pay if it does not have a contract with the supplier or does not meet the SOW criteria, and a discounted or contracted price, which is the price that the hospital will pay if it has a contract with the supplier and meets the SOW criteria for the contract.

label those whose share of total surgeries in the category in their hospital falls in the top decile among all surgeons in the sample as "high-share surgeons."

In all, we have 201 hospitals (34.3% teaching) across the US in our sample (we exclude 13 hospitals with either non-GPO contracts or incomplete data), collectively employing 17,239 surgeons. The average number of surgeons per hospital is 85 with a median of 70 and a standard deviation of 53. The average share of high-share surgeons in our sample is 17.4% with a median of 12.8% and a standard deviation of 15.7%. We have 19 quarters of data with 4 quarters before launch of innovation. Therefore, we have a balanced panel of 3819 (201×19) hospital-quarters for analysis.

Descriptive Statistics

We begin with descriptive statistics on prices and contracts to build intuition around the data. Table 1 shows the descriptive statistics of prices from both suppliers. Hospitals with contracts pay significantly lower prices for all three product lines. Figure 1 shows the average off-contract and discounted prices; the discounted prices are about 30% to 50% lower than the off-contract prices. The old product from Supplier A is the cheapest on both off-contract and discounted price.

Figure 2 shows the time series of average on- and off-contract prices for each of the three products. A few observations are noteworthy. First, both off-contract and discounted prices for Supplier A's product drop over time. The price drop starts a couple quarters before the launch of Supplier B's product, perhaps in anticipation of that launch. Second, although the off-contract price of Supplier A's product is initially higher than that of Supplier B's old product, the average contracted prices of the two products are initially about the same. Third, over time, as Supplier B's new product gains share, the supplier increases the off-contract prices of both its old and new products. Such a price increase might seem counter-intuitive at first, since for a novel technology a skimming pricing strategy could be more common in B2C markets. In this case, however, because of the decoupling of purchase and consumption, price is not the main driver of adoption; surgeons who adopt the new product will keep using it regardless of its price and surgeons who would like to switch to the new product will also consider the usage benefits rather than price. Therefore, raising the off-contract prices over time can encourage hospital buyers to switch contracts as their surgeons increasingly start using the new product. Fourth, Supplier B increases its old product's contracted price while keeping the contracted price of its new product roughly the same over time. Anecdotally, this strategy incentivizes buyers to encourage surgeons to use Supplier B's new product, as the supplier plans to shelve the older product for production efficiency. Finally, small price jumps in the graph are due to price re-negotiations with one or more GPOs over time.[10]

Next, we consider the cross-sectional price variation across hospitals. Figure 3 presents snap-

[10]Note that these figures present average prices across hospitals. The prices faced by each hospital is a step function with periodic changes due to GPO re-negotiations with suppliers, GPO switching, etc.

Figure 2: Average Per-Surgery Product Prices Over Time

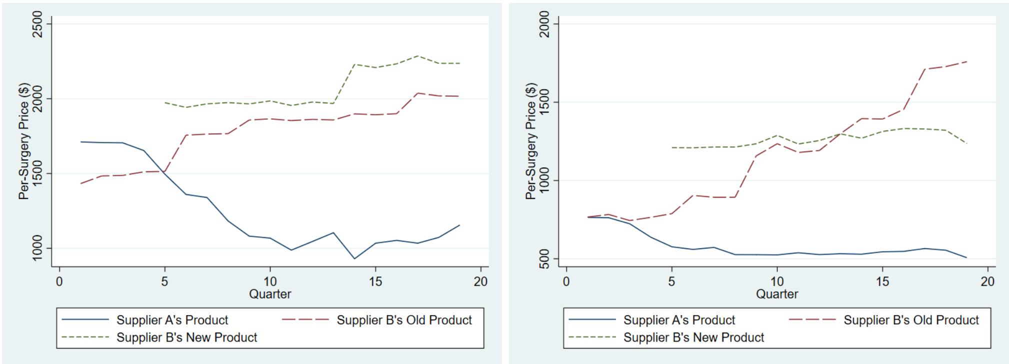

(a) Average Off-contract Prices Over Time (b) Average On-contract Prices Over Time

Note: This figure shows the average per-surgery price for on and off-contract hospitals over time for each of the three products in the market. The price of Supplier A's product drops in response to the launch of Supplier B's new product in the fifth quarter. Supplier A gradually increases its off-contract prices to encourage hospitals to switch their contracts. Supplier B also increases its discounted price for its old product to encourage its committed hospitals to use more of its new product as it plans to gradually shelve its old product due to production efficiency considerations.

Figure 3: Snapshot of Distribution of Per-Surgery Prices for Products at Q5 and Q19

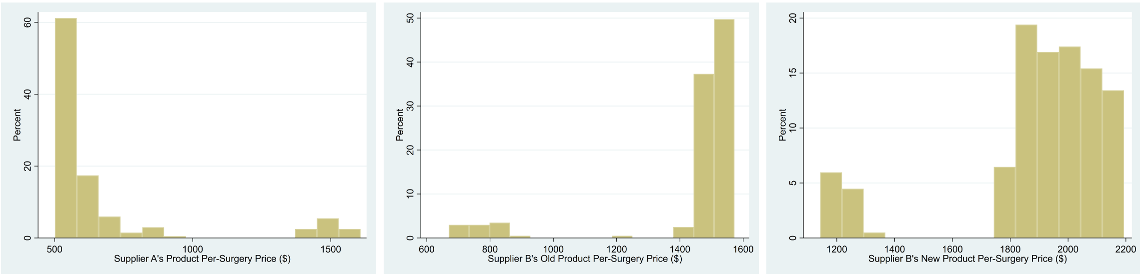

(a) Supplier A at Q5 (b) Supplier B-Old Product at Q5 (c) Supplier B–New Product at Q5

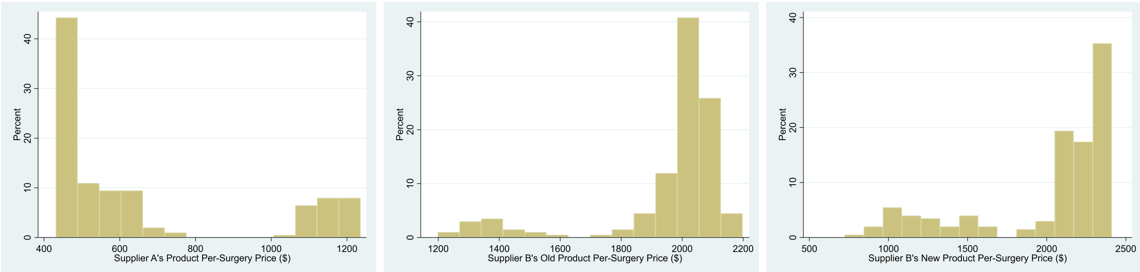

(d) Supplier A at Q19 (e) Supplier B–Old Product at Q19 (f) Supplier B–New Product at Q19

Note: The figures depict cross-sectional variation in prices of each of the three products in the market for two time periods: Quarter 5 (when Supplier B's new innovation was launched) and Quarter 19 (the last quarter). Each bar shows the percentage of hospitals in the sample that are charged a certain price in the corresponding quarter.

shots of per-surgery price distributions across hospitals for the three product lines at quarter five (Q5, when innovation was launched) and quarter 19 (Q19, last period). Each snapshot shows two clear groups: for a large group of hospitals in contract with Supplier A that meet the threshold share, the price for Supplier A's product is significantly lower than the price for Supplier B's products. For a smaller group of hospitals in contract with Supplier B, Supplier A's price is much higher than the price for Supplier B's products. The distributions of prices in other periods are similar. Around Q19, however, the number of hospitals charged a lower price by Supplier B increases. This change reflects that more hospitals switched to Supplier B closer to Q19.

Finally, we explore the realized distribution of shares for products from Supplier B across hospitals in Figure 4. Considering the duopoly structure of the market, the distribution of Supplier A's shares is the mirror image of Figure 4a. Most hospitals purchase less than 20% of their supplies in the focal category from Supplier B for majority of quarters (Figure 4a). This distribution is expected, given that for a large part of the data collection period, most hospitals are in contract with Supplier A, with its 80% share threshold. For a smaller number of hospital-quarters, Supplier B's share is above 60% (the threshold share for the supplier) with a high concentration around 60%. Examining the breakdown of Supplier B's share across its two products reveals that the supplier's new product pushes many of the hospitals' shares towards 20% (Figure 4c). Moreover, the new product is responsible for the concentration of hospitals around the 60% threshold, whereas both old and new products are responsible for hospitals that use higher shares. Hospitals with higher shares of Supplier B are those already committed to Supplier B before the launch of the new product or that switched to Supplier B earlier, meaning surgeons have had more time to start using the new product.

Figure 4: Distribution of Supplier B Product Shares

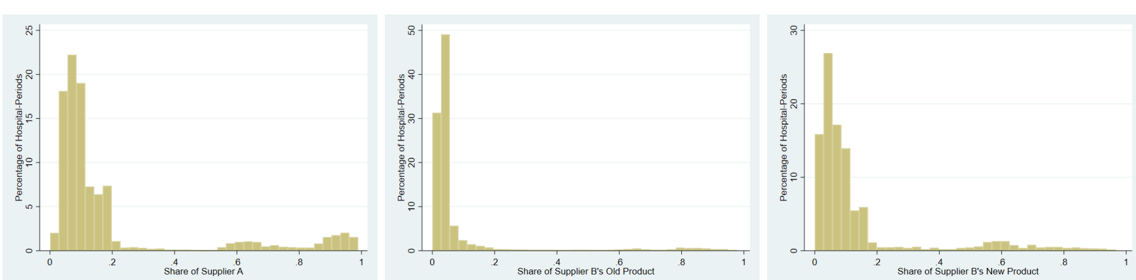

(a) Old + New (b) Old (c) New

Note: This figure depicts distribution of realized shares for Supplier B's products. Each bar shows the percentage of hospital-quarters for which the corresponding product constitutes a certain share in category.

Model-Free Evidence

We provide model-free evidence to support the following assumptions and features of the organizational buying and innovation adoption model: (i) users' influence on buyer contracting; (ii) the effect of hospital contracts (by buyers) on users; (iii) forward-looking behavior by buyers; and (iv) limited effect of GPO affiliations on price.

User Influence on Buyers Contracting: The solid line in Figure 5 shows the share of hospitals that contract with Supplier B over time. It reveals that even though the innovation is more expensive, it induces contract switching (the key buyer decision) over time to Supplier B. Prior to the new innovation, the share of hospitals contracting with Supplier B is around 10%. Supplier B introduces the new product in Q5, after which we observe a steady increase in the share of hospitals in contract with (or "committed to") the supplier; by Q19, it reaches over 23%.

Figure 5: Share of hospitals in contract with Supplier B over time

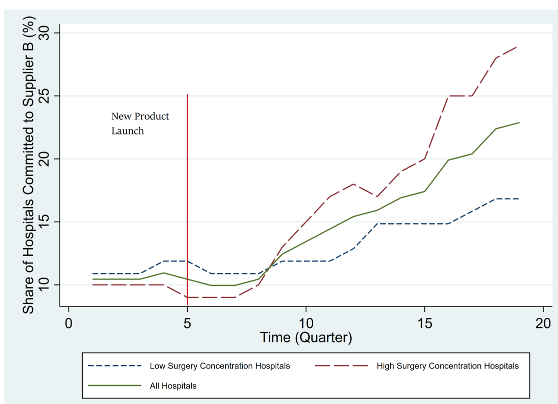

Note: This figure shows the share of hospitals in contract with Supplier B over time, overall, and by surgery concentration. Surgery concentration reflects the distribution of surgeries among surgeons; in a high surgery concentration hospital, a few surgeons perform a disproportionately large share of surgeries. After the launch of the new product by Supplier B, hospitals switch contracts to this supplier (the solid line in the middle). The contract switching is faster for hospitals with above-median surgery concentration (the top dashed line), compared to below-median ones.

To assess whether users (surgeons) influence contract switching, we explore if the contracting pattern varies by the extent to which surgeries are concentrated among fewer surgeons. The underlying assumption is that surgeons who perform a larger share of surgeries in a hospital have more clout. Therefore, in a hospital with highly concentrated surgeries, influential surgeons will sway contracting decision towards a more expensive, but preferred product. Accordingly, we compare

Table 2: Surgery concentration and contract switching

	(1) Contract Switching	(2) Contract Switching	(3) Contract Switching	(4) Contract Switching
Share of High-Share Surgeons	2.048*	1.890*		
	(0.922)	(0.948)		
Surgery Concentration (HHI)			2.442*	2.116*
			(1.048)	(1.073)
Supplier B Share	3.624***		3.556***	
	(0.438)		(0.437)	
Supplier B Old Product Share		0.663		0.634
		(0.718)		(0.721)
Supplier B New Product Share		3.260***		3.302***
		(0.517)		(0.514)
Constant	-6.544***	-5.117***	-6.444***	-5.074***
	(0.428)	(0.414)	(0.402)	(0.408)
Observations	3268	2548	3268	2548

Note: This table presents the estimates of a logit model of contract switching to Supplier B on hospital surgery concentration. The results show a positive relationship between probability of switching contracts and surgery concentration, measured by HHI or share of high-share surgeons, and controlling for overall share of Supplier B, or share of each of its products. Standard errors in parentheses. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

contract switching patterns across hospitals based on a median split of surgery concentration.[11] Figure 5 shows that, indeed, hospitals with greater surgery concentration switch faster; not only is the trajectory (slope) of contract switching to Supplier B greater for hospitals with greater surgery concentration, but the overall percentage of hospitals who contract with Supplier B is also higher (29% and 17% for higher and lower concentration hospitals, respectively), even though there is little difference between the two types of hospitals before introducing the innovation. To rule out the pattern reflecting faster attainment of threshold shares in hospitals with greater concentration, we estimate a logit model of contract switching controlling for Supplier A's share of surgeries in Table 2. We find that hospitals with greater surgery concentration have a higher probability of contract switching than those with lower concentration even after controlling for Supplier B's SOW from the hospital. This evidence suggests that more powerful surgeons have greater influence on contracting decisions, or, more broadly, surgeons do influence contracting decisions. This finding supports our accounting for surgeons' influence on buyers in the model.

Influence of buyer contracting on surgeon usage: Contracts provide superior support for surgeons and thus nudge them towards using the contracted supplier's product lines. Here we show

[11] We use Herfindahl-Hirschman Index (HHI) as a measure of concentration. To calculate HHI for each hospital, we added the squared average share of each surgeon from total surgeries in the category in the hospital: $HHI_h = \sum_{s_h \in S_h} (\frac{number\ of\ surgeries\ by\ s_h}{total\ number\ of\ surgeries\ at\ h})^2$. Using share of high-share surgeons at the hospital (defined earlier) as a measure of surgery concentration also yields similar results.

empirical support for the shift in surgeon usage towards contracted products.

Figure 6: Average share trajectory of innovation

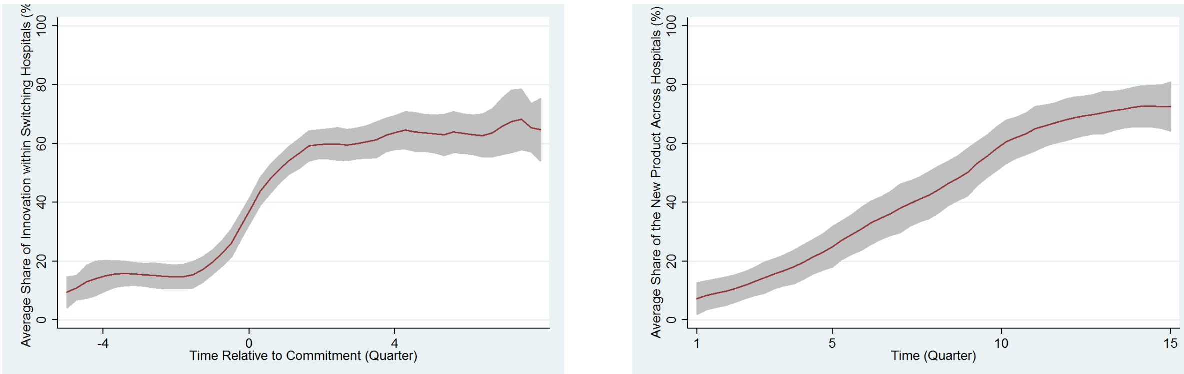

(a) Hospitals switching from Supplier A to B (Quarter 0: Switching Quarter)

(b) Hospitals already contracted with Supplier B before launch of innovation

Note: The two figures depict the trajectory of shares of the new product among hospitals that are: (a) initially not in contract with Supplier B, but switch contracts during analysis period; (b) already in contract with Supplier B before the launch of the innovation. The sharp jump in share of the innovation in panel (a) around the time of contract switching (time 0 on the x-axis) and its contrast with the gradual increase in share of the new product among hospitals that never switch contracts because they are already in contract with Supplier B in panel (b) indicates the effect of contracts on surgeons' usage behavior. Graphs based on 25 and 21 hospitals in left and right panels, respectively.

Figure 6a shows evolving innovation share among hospitals that switched from Supplier A to Supplier B during the analysis period. The quarter in which each hospital switched is coded as zero on the x-axis. Around the quarter in which the hospital switches to Supplier B, the share is about 20% (the maximum share without impacting Supplier A discounts), but over the next two quarters the share shifts dramatically towards the innovation, ultimately reaching over 60% (the threshold to obtain discounts for a hospital in contract with Supplier B). In contrast, Figure 6b is constructed using hospitals already in contract with Supplier B before the innovation launch. The x-axis in this figure represents time since the innovation launch. In Figure 6b, the change in share towards the innovation from the old product is much more gradual; it takes about 10 quarters to hit the 60% share. The contrast with Figure 6a suggests that contract switching by buyers influences surgeon product usage.

Forward-looking behavior by buyers: 73% of hospitals that switch to Supplier B do so before Supplier B's hospital share reaches 50% (well below the 60% discount threshold). Further, Supplier B's share does not hit 60% in these hospitals in the quarter after switching, though the

Table 3: Drivers of Price Variance for Different Products

	R^2			
Supplier A's product price	10.9%	74.1%	86.8%	87.9%
Supplier B's old product price	8.5%	86.5%	92.0%	92.1%
Supplier B's new product price	4.3%	76.1%	89.8%	90.0%
Included explanatory variables				
Time fixed effects	✓	✓	✓	✓
Contract & SOW threshold met*		✓	✓	✓
Category spend			✓	✓
GPO fixed effects**				✓

Note: This table presents the explanatory power (R^2) of various factors that could affect price for each of the three products in the market. The results show that the contract status and category spend explain most of the variation in price, whereas GPO affiliations explain very little.

* Whether hospital has contract with the supplier and has met the contract SOW threshold (60% for Supplier B or 80% for Supplier A).

** Other hospital characteristics, i.e., size and teaching status have no additional explanatory power as their effect is already absorbed by category spend and GPO affiliation.

switching accelerates the share shift. If buyers were myopic, they would not have switched early as switching increases their current cost. Switching is justified only if the buyers are forward-looking and consider that switching speeds up surgeons' usage of the new product.[12]

GPO affiliations and price: Finally, we explore the extent to which GPOs affect prices paid by hospitals beyond other factors. Table 3 presents the explanatory power (i.e., R^2) of five different factors driving price variation for each of the three products. Although temporal price changes explain minimal price variance, whether a hospital has a contract with a supplier and meets the threshold SOW has significant explanatory power, demonstrating the importance of modeling contract status in our context. Because suppliers differentiate among larger and smaller hospitals based on total spend beyond share, adding average category spend for each hospital unsurprisingly increases the explanatory power of regressions by about 14%. Critically, GPO fixed effects explain minimal additional variation ($< 1\%$). Overall, two factors justify abstracting away from modeling GPO memberships: (i) the limited explanatory power of GPO affiliations for the price, and (ii) the limited effect of any single category on GPO membership decisions (see earlier discussion on institutional details).

[12]Interviews with industry experts and hospital buyers also support forward-looking behavior in the contract switching process. A strategic sourcing director at a 10-hospital health system described their forward-looking approach thus: "We have a clinical value analysis specialist who would send surgeons a list of questions by email.... We then do a cost analysis for the next year based on surgeons' input and how they would react to different contracting decisions."

MODEL

The market consists of J suppliers, each offering an existing product, $j = 1...J$. Without loss of generality, we denote supplier J as introducing a new product $J + 1$ during the analysis period. In the market under consideration, there are only two suppliers, so $J = 2$. There is an existing product $j = 1$ from supplier 1, an existing product $j = 2$, and a new product $j = 3$ from supplier 2. Table 4 presents the complete list of the notation used in the model.

The customers are the hospitals, but two sets of decision makers impact demand at each hospital: the buyer and users (surgeons). Buyers make the contracting decision, i.e., deciding whether to continue with the current contract or switch suppliers (by incurring a contract switching cost), taking into account not just the costs of the product (what hospital pays) given the SOW contract, but also the aggregate user dissatisfaction that will arise from contracting with one supplier relative to another. The surgeons can use products from either the contracted or non-contracted suppliers, but contracted suppliers provide higher levels of service. Surgeons make the decision about which product to use for surgeries conditional on the buyer's contract choice. We assume that because the share of surgeries for any one surgeon at a hospital is relatively small, surgeon impact on buyer's choice is only in the aggregate; an individual surgeon's choice cannot have a direct impact on the buyer's contracting decision, so the individual surgeon treats the buyer's choice as exogenous. In any given period, each surgeon uses only one product for all their surgeries. Whereas all surgeons have tried the existing products, trying the new product for the first time involves a trial cost.

As explained in the section on institutional details, GPO membership and the SOW contracts available from the GPO are exogenous to the hospital buyer at this stage and outside the scope of the model. Let shr_{jht} be the share of supplier j (all its products in the category) from total category purchases of hospital h at time t. Given GPO membership and size of a hospital, the buyer has only one relevant contract from each supplier that yields the lowest price. We treat these contracts from each supplier as the choice set of the buyer. The relevant SOW contract from each supplier j for hospital h has an off-contract price of $p_{jht} = \overline{p}_{jt}$ when $shr_{jht} < \underline{shr}_j$ and a loyalty discounted

Table 4: Model Notation

Notation	Description
μ_{old}	Surgeon's utility from using Supplier B's old product
μ_{new}	Surgeon's utility from using Supplier B's new product
$\gamma\ (\gamma^H, \gamma^L)$	Utility from contract (estimated values for high, low-share surgeons, respectively)
τ	Incremental preference for the new product at teaching hospitals
σ_{s_h}	Surgeons' switching cost for surgeon type s
α	Buyer's price sensitivity
σ_b	Buyer's switching cost
λ	Buyer's marginal (dis)utility for surgeons' dissatisfaction with the contract
J	Total number of suppliers in the market
j	Index for suppliers (and products)
h	Index for hospitals
t	Index for time
t_0	Time period that the new product of Supplier B is launched
shr_{jht}	Share of supplier j at hospital h at time t
$\underline{shr}_j$	Threshold share for supplier j, above which hospitals get the discounted price
$\hat{shr}_{hjt}$	Predicted share of product j at hospital h at time t
p_{jht}	Price of product j at hospital h and time t
$\underline{p}_{jt}$	Discounted (on-contract) price of product j at time t
$\overline{p}_{jt}$	Price of product j at time t if off-contract or below SOW threshold
d_{ht}	Hospital h's buyer's choice of supplier at time t, i.e., equals j if the buyer signs a contract with supplier j
d_{hjt}	Dummy equivalent of d_{ht}; equals one if hospital h has a contract with supplier of product j at time t and zero otherwise
u_{s_hjt}	Utility of surgeon s in hospital h at time t, using product j
$\overline{u}_{s_hjt}$	The deterministic portion of the surgeon's utility
u_{ht}	Utility of the buyer in hospital h at time t
$\overline{u}_{ht}$	Deterministic part of the buyer utility
δ	The buyer's discount factor
θ_s^k	Surgeon model parameters for latent segment k
θ_b^k	Buyer model parameters for latent segment k
X_s	Surgeon s characteristics
X_h	Hospital h characteristics
X_{jt}	Characteristics of product j at time t (i.e., product time fixed effects)
π_k	A vector of probabilities that each hospital belongs to segment k
s_h	A surgeon at hospital h
S_h	Set of all surgeons at hospital h
$f(.)$	Utility from the product
β_{s_h}	Parameters defining surgeons' consumption utility for various product characteristics
$g(.)$	Utility from contract to the surgeon if product j is used with contract status d_{ht}
γ_{s_h}	Utility from contract for surgeon s_h
σ_{s_h}	Surgeon's trial cost for Supplier B's new product
ξ_{hjt}	Unobserved hospital product time specific shock to the surgeons' utility
I_{J+1}^j	Indicator function that equals one if $j = J + 1$, i.e., for the new product, and zero otherwise
$T_{s_h(t-1)}$	A dummy variable which is equal to one if surgeon s_h has tried Supplier B's new product by period $t-1$
ε_{s_hjt}	i.i.d surgeon-product-time shock to the surgeon's utility already tried Supplier B's new product up to time t and zero otherwise
$P_{s_hjt}^T$	Probability of choosing product j for surgeon s_h who has already tried the new product
$P_{s_hjt}^{NT}$	Probability of choosing product j for surgeon s_h who has not yet tried the new product
$P(T_{s_ht}) = 1$	Probability that surgeon s_h has already tried the new product the new product at time period t
w_{s_h}	Share of surgeries by surgeon s_h at hospital h
$IV_{hts}^{d_{ht}}$	Inclusive value for the surgeon utility
$IVL_{hts}^{d_{ht}}$	Inclusive value loss from contract decision d_{ht}
$EIVL_{s_ht\mid d_{ht}}$	Expected value of the inclusive value loss

price of $p_{jt} = \underline{p}_{jt} < \overline{p}_{jt}$ when $shr_{jht} \geq \underline{shr}_j$, where $\underline{shr}_j$ represents the threshold SOW to obtain discount from supplier j.

We now present a structural model of contracting by buyers at hospitals facing SOW contracts from J suppliers and downstream trial of innovation and ongoing usage by surgeons. The model accommodates observable and unobservable heterogeneity among hospital buyers and among surgeons within hospitals. Specifically, we allow for observed and unobserved heterogeneity among hospitals in the extent to which buyers can influence surgeons through their choice of contract. Similarly, we allow for observed heterogeneity among surgeons arising from surgery concentration, and unobserved heterogeneity among surgeons in terms of their trial cost, contract utility, and relative preferences for products.

Overall, we consider an infinite horizon model: Within each period, there are two stages of decision making by the buyer and the users. At the beginning of each period, the buyer makes a contracting decision with the supplier whether to continue the contract with the current supplier or switch to a competing supplier. Conditional on the contract, users (surgeons) decide which product to use. Of relevance to the multi-agent nature of the B2B model, the buyer's contracting decision considers not only the total financial cost to the hospital, but also user dissatisfaction with the choice of contracted supplier. As discussed earlier, given this tradeoff between costs to the hospital arising from SOW contracts and user dissatisfaction, the buyer's problem is an optimal stopping problem involving forward-looking behavior. The surgeon's choice problem is, by contrast, myopic. We next discuss the user and buyer choice models in turn.

Users (Surgeons)

We model the utility from product j at time t for a surgeon working at hospital h denoted as s_h, in terms of five elements. The first three capture key aspects of surgeon choice in the surgeon utility model. First, we capture the notion of the physician preference item (PPI) by allowing each surgeon to have a preferred product line used across all surgeries at any given time.[13] Second,

[13]While technically feasible for surgeons to use different products in different surgeries, in practice they do not. It is challenging to regularly alternate product lines given the variations in color coding of different tools across different

using a product line from a supplier with which the hospital has contracted is easier as the surgeon (i) is likely to experience fewer stockouts, especially for lower volume, rarely used tools, and (ii) will get better technical support in the operating room. Third, while surgeons have tried and are familiar with all the existing products on the market, a trial cost exists in using the new product line $J + 1$ for the first time. This trial cost can arise due to the surgeon's need to become familiar with the new product's different tools (i.e., SKUs), characteristics, and use cases. Since each period is a quarter, and a surgeon can master using the new product in less than a quarter, a surgeon's learning happens within a period. The details of surgeons' learning process are not visible to us; therefore, we do not model surgeons' learning.[14] Note that in considering these tradeoffs, the surgeon does not care about price.

In addition to these three sources of utility, we incorporate two unobservable shocks. First, we allow for an i.i.d. surgeon-product-time level shock from an extreme value distribution in every period to model surgeon product choice as a probability that can be integrated across surgeons to obtain the product's share in each hospital and period. Finally, as in the BLP model (Berry et al. 1995), we allow for an unobservable hospital-product-time level shock that helps match the model predicted share with the observed aggregate hospital share data in each time period. Equation (1) presents the surgeon's utility function where all the terms are defined in Table 4.

$$u_{s_h jt} = \underbrace{f(X_{s_h}, X_h, X_{jt}|\beta_{s_h})}_{\text{utility from product}} + \underbrace{g(X_s, X_h|d_{hjt}, \gamma_{s_h})}_{\text{utility from contract}} + \underbrace{I^j_{J+1}\sigma_{s_h}(1 - T_{s_h(t-1)})}_{\text{trial cost of new product}} + \xi_{hjt} + \varepsilon_{s_h jt} \tag{1}$$

As discussed earlier, the share of the innovation in all hospitals increases over time because the number of surgeons who have tried the product initially is zero and increases steadily over time. Consistent with institutional details, the utility framework of equation (1) accommodates this diffusion in two key ways. First, a hospital buyer switching and signing a contract with the innovation supplier can lead to a sudden increase in surgeon trial as it now becomes easier to use

product lines, and differences in how precisely they should be handled.

[14] If the analysis periods were shorter (say weekly) or learning required more than a quarter, the learning process may need to be modeled. This would require surgeon-level product choice data, which we do not observe.

the product (as shown in Figure 6a). Second, after a surgeon decides to try the innovation, ongoing usage does not involve the trial cost. This feature of increase in share differs from durable goods (e.g., cameras, video games) with one-time purchase. Thus, as more surgeons try the innovation, sales steadily increase.

In the estimated model, we allow for product fixed effects (μ_j) to capture surgeon preferences for each product line. Additionally, consistent with our institutional context, we allow surgeons in teaching hospitals to have a different preference (τ) for the innovation than at other hospitals. We model contract benefits as an incremental surgeon utility (γ_{s_h}) when using an on-contract product. This benefit of course depends on the buyer's contracting decision (d_{hjt}). When deciding whether to try the new product for the first time, surgeons will likely consider the immediate cost of using an unfamiliar product plus the value of having the new product as an option in the future. Therefore, σ_{s_h} captures the trial cost net the future option value of having tried the new product before.

All surgeons have tried the existing products and face no trial cost on those; they only face a trial cost on the innovation when they try it for the first time. We therefore model the trial process only for the innovation and track the evolving base of consumers who have tried the innovation.

$$\begin{aligned} &P(T_{s_h t} = 1) = 0 && \forall t \leq t_0 \\ &P(T_{s_h t} = 1) = P(T_{s_h (t-1)} = 1) + (1 - P(T_{s_h (t-1)} = 1)) P^{NT}_{s_h (J+1) t} && \forall t > t_0 \end{aligned} \tag{2}$$

Here, $P(T_{s_h t} = 1)$ represents the probability that surgeon s_h has tried the new product before or at time period t. t_0 represents the launch period of Supplier B's new product. We can write the probability of a product being chosen by a surgeon s_h in a given period t, conditional on trial versus no trial until the previous period $t-1$, denoted as $P^T_{s_h jt}$ and $P^{NT}_{s_h jt}$, respectively:

$$\begin{aligned} P^T_{s_h jt} &= \frac{\exp(u_{s_h jt} | T_{s_h (t-1)} = 1)}{\sum_{m=1}^{J+1} \exp(u_{s_h mt} | T_{s_h (t-1)} = 1)} && \forall j \in \{1, ..., J+1\} \\ P^{NT}_{s_h jt} &= \frac{\exp(u_{s_h jt} | T_{s_h (t-1)} = 0))}{\sum_{m=1}^{J+1} \exp(u_{s_h mt} | T_{s_h (t-1)} = 0))} && \forall j \in \{1, ..., J+1\} \end{aligned} \tag{3}$$

Using the above equations, we can then write the overall share of each product j at each hospital

h by summing over the probability of choosing j across the set of all surgeons (or surgeon types), S_h, broken down by their adoption status and weighted by the share of surgeries by each surgeon at each hospital, w_{s_h}:

$$\hat{shr}_{hjt} = \sum_{s_h \in S_h} w_{s_h}[P(T_{s_h(t-1)} = 1)P^T_{s_hjt} + P(T_{s_h(t-1)} = 0)P^{NT}_{s_hjt}] \tag{4}$$

The Buyer

In each time period, the buyer decides whether to stay with the same supplier or switch to another supplier by signing a new contract. For this decision, the buyer considers (i) the predicted purchasing costs to the organization; (ii) the cost of switching contracts to a new supplier; and (iii) the predicted surgeon unhappiness if contracting with a supplier not preferred by users. Operationalizing purchase costs and costs of switching is conceptually straightforward; our key innovation is on how we operationalize user unhappiness under alternative supplier contracts.

User disutility from buyer contracting: Let $d_{ht} = j$ denote that the buyer contracts with supplier j at time t; this translates to $d_{hjt} = 1$ when $j = 1...J-1$ and $d_{hJt} = 1$ and $d_{h(J+1)t} = 1$ when $j = J$ in equation (1) reflecting that supplier J supplies both products J and $J+1$. The user model in equation (1) allows for buyers to impact user behavior to favor products from suppliers that the buyer contracts with through the term $g(X_s, X_h|d_{hjt}, \gamma_{s_h})$. Yet getting users to adopt a product other than their preferred option reduces their utility. Because buyers and users influence each other, we consider a disutility for the buyer in proportion to the total loss in utility across all surgeons at the hospital due to the buyer's contracting choice.

We operationalize the effect of unhappiness of surgeon s_h due to the hospital h contracting with the surgeon's non-preferred supplier as the difference in the surgeon's expected utility when the contract is with the preferred supplier as opposed to the contracted supplier. If the contracted and preferred suppliers are the same, this metric should be zero. Given that the utility of the surgeon includes an extreme value error term (i.e., $u_{s_hjt} = \overline{u}_{s_hjt} + \varepsilon_{s_hjt}$), the expected maximum utility for surgeon s_h when the contract is with supplier j, i.e., $d_{ht} = j$ is the inclusive value of the surgeon's utility. Note this integration of the error shocks in operationalizing surgeon unhappiness is appropriate because buyer contracting choice is based on forecasts of surgeon unhappiness.

$$IV_{s_h t|d_{ht}} = E \max_k (u_{s_h kt|d_{ht}})) = \log(\sum_{k=1}^{J+1} \exp(\overline{u}_{s_h kt|d_{ht}})) \tag{5}$$

Therefore, the loss in utility for user s_h at time t due to buyer's contracting decision d_{ht} is the maximum difference in inclusive values ($IV_{s_h t}$) if the buyer contracted with any supplier ($d'_{ht} = 1, ..., J$). We denote this as the inclusive value loss (IVL):

$$IVL_{s_h t|d_{ht}} = \max_{d'_{ht}=\{1,...,J\}} (IV_{s_h t|d'_{ht}} - IV_{s_h t|d_{ht}}) \tag{6}$$

Note that in equation (6), zero is included in the set of values over which we take the maximum, because it includes the currently contracted supplier. So the value of IVL will always be non-negative. IVL captures only dissatisfaction with contracting decisions made by the buyer. Also, consider that surgeons can use products from any supplier regardless of the contract status, so a surgeon's unhappiness arises from the inability to obtain contract benefits for the preferred product. Given the above, the expected inclusive value loss ($EIVL$) for a surgeon is obtained by:

$$EIVL_{s_h t|d_{ht}} = P(T_{s_h(t-1)} = 1) IVL^{T}_{s_h jt|d_{ht}} + P(T_{s_h(t-1)} = 0) IVL^{NT}_{s_h jt|d_{ht}} \tag{7}$$

$EIVL$ integrates over the surgeon's adoption probabilities. Because IV already integrated out the surgeons' unobserved error shocks, $EIVL$ accounts for buyer's expectation of how surgeons' preferences evolve when choosing a contract at start of period t, before shares are realized.

We now specify the per-period utility of the buyer at hospital h at time t, conditional on making decision d_{ht} in terms of three components mentioned earlier (i.e., expected purchasing cost based on the buyer's prediction of surgeons' usage, switching cost, and expected surgeons' unhappiness captured through EIVL):[15]

[15]Note that the expected user dissatisfaction is based on the inclusive value loss (IVL) which integrates out surgeon shocks $\varepsilon_{s_h kt}$. This is appropriate because buyer contracting choice decisions are made at the beginning of period t

$$\begin{aligned} u_{ht}(X_{ht}, \epsilon_{ht}, d_{ht}; \theta) =& \underbrace{\alpha \sum_{j=1}^{J+1} p_{jt}(d_{hjt})(\hat{shr}_{jt}(X_{ht}, d_{ht}))}_{\text{expected purchase cost}} + \underbrace{\sigma_b I^{d_{ht} \neq d_{ht-1}}}_{\text{switching cost}} + \underbrace{\lambda \sum_{s_h \in S_h} w_{s_h} EIVL_{s_h t | d_{ht}}}_{\text{expected surgeon unhappiness}} + \varepsilon_{htd_{ht}} \\ =& \overline{u}_{ht}(X_{ht}, d_{ht}; \theta) + \varepsilon_{htd_{ht}} \end{aligned} \tag{8}$$

Note that the expected surgeon unhappiness includes a weighting based on the number of surgeries by a surgeon, allowing us to account for greater influence on the buyer from high-share surgeons than low-share surgeons. In addition to these components, we include an additive structural stochastic shock that captures factors observed by the buyer (and not the researcher) that affect a decision. These shocks are assumed to be independent across hospitals, time periods, and decisions as well as identically distributed, following Type I extreme value distribution.

As more surgeons (users) adopt the new product, the buyer must decide whether to continue the contract with the current firm or switch to the innovating firm. The buyer's objective is to maximize the total discounted utility over the infinite horizon and can be summarized in the following sequence problem:

$$\max_{\{d_{ht}\}_{t=1}^{\infty}} E \sum_{t=1}^{\infty} \delta^{t-1} u_{ht}(X_{ht}, \epsilon_{ht}, d_{ht}; \theta) \tag{9}$$

where u_{ht} represents the buyer's utility specified above. Here, θ represents the set of parameters for the utility function of the buyer and surgeons that will be estimated, and δ is the buyer's discount factor. In this equation, the expectation is over the set of unobserved state variables and also the evolution of observed state variables over time.

Value Functions

We can re-write the buyer's problem, presented in equation 9, as an infinite sequence of single-period decisions. We first define the (ex-post) value function as:

based on forecasts of future cost and user dissatisfaction.

$$V(X_{ht}, \epsilon_{ht}; \theta) = \max_{\{d_{ht}\}_{\tau=t}^{\infty}} E \sum_{\tau=t}^{\infty} \delta^{\tau-t} u_{h\tau}(X_{ht}, \epsilon_{ht}, d_{ht}; \theta) \tag{10}$$

and the choice-specific value function as the sum of present and discounted maximum future utilities conditional on the buyer's decision as:

$$\begin{aligned} v(X_{ht}, \epsilon_{ht}, d_{ht}; \theta) &= \overline{u}(X_{ht}, d_{ht}; \theta) + \delta E_{X_{ht+1}, \epsilon_{ht+1} | X_{ht}, \epsilon_{ht}, d_{ht}} V(X_{ht+1}, \epsilon_{ht+1}; \theta) + \varepsilon_{htd_{ht}} \\ &= \overline{v}(X_{ht}, d_{ht}; \theta) + \varepsilon_{htd_{ht}} \end{aligned} \tag{11}$$

Note that we follow Rust (1987) in assuming the conditional independence of unobserved stochastic shocks to write the expected value of future utilities as a function independent from $\varepsilon_{htd_{ht}}$ in the second line of the above equation. That is, we assume both observed and unobserved state variables, at any time period, are independent from unobserved state variables of the previous time period, conditional on observed state variables of previous period:

$$P(X_{ht+1}, \epsilon_{ht+1} | X_{ht}, \epsilon_{ht}, d_{ht}) = P(\epsilon_{ht+1} | X_{ht+1}) P(X_{ht+1} | X_{ht}, d_{ht})$$

Using the above notation, ex-post value function for each period could be written based on flow utility function and ex-post value function of the next period (Bellman Equation):

$$V(X_{ht}, \epsilon_{ht}; \theta) = \max_{d_{ht}} \{ \overline{u}(X_{ht}, d_{ht}; \theta) + \delta E_{X_{ht+1}, \epsilon_{ht+1} | X_{ht}, \epsilon_{ht}, d_{ht}} V(X_{ht+1}, \epsilon_{ht+1}; \theta) + \varepsilon_{htd_{ht}} \} \tag{12}$$

Defining $EV(X_{ht}; \theta) = E_{X_{ht+1}, \epsilon_{ht+1} | X_{ht}, \epsilon_{ht}, d_{ht}} V(X_{ht+1}, \epsilon_{ht+1})$ and considering properties of Type I extreme value distribution, we can write:

$$EV(X_{ht}; \theta) = \int_{X_{ht+1}} \log\{ \sum_{d_{ht} \in \{1,2,\ldots,J\}} \exp(\overline{u}(X_{ht}, d_{ht}, \theta) + \delta EV(X_{ht+1}; \theta)) p(dX_{ht+1} | X_{ht}, d_{ht}) \} \tag{13}$$

We will use this equation which defines EV in a recursive fashion for estimation.

Transition of State Variables

The set of observed state variables includes all the factors that affect each of the first three components of the utility specified in equation 8. That entails on- and off-contract prices for different products, observed characteristics of surgeons (S_h), hospital characteristics (κ_h), and the level of adoption of the new product. Contractual status in the previous time period is also included.

There are two relevant prices for each product in each hospital at any period: an off-contract price ($\overline{p}_{jht}$ or $p_{jht}(d_{hjt} = 0)$) and a discounted price ($\underline{p}_{jht}$ or $p_{jht}(d_{hjt} = 1)$). We denote each of these $2 \times (J+1)$ prices by p_{jht}. We assume prices follow a first order Markov process and the buyer has rational expectations about their evolution. More formally: $p_{jht} \sim Multinomial(1, \pi_{p_{jht-1}})$, where $\pi_{p_{jht-1}}$ is a vector of discrete probability distribution for p_{jht}, conditional on p_{jht-1}. The transition of the installed base of the new product is governed by equation 2. Formally, the set of all the state variables: $X_{ht} = \{\{p_{jht}(d_{hjt} = 0), p_{jht}(d_{hjt} = 1)\}_{j=1}^{J+1}, P^T_{s_h jt}, d_{ht-1}, \kappa_h, S_h\}$.[16]

ESTIMATION

Mojir and Sudhir (2021) combine the Expectation-Maximization (EM) algorithm approach introduced in Arcidiacono and Jones (2003) with the Mathematical Programming with Equilibrium Constraints (MPEC) approach developed in Su and Judd (2012) to estimate a dynamic structural model of consumer search with heterogeneous (latent class) preferences. In the MPEC approach, the dynamic structural model is formulated as a mathematical programming optimization problem of maximizing the likelihood function of observed agent choices. The Bellman equations of the dynamic structural model serve as constraints for the optimization problem. This approach has the advantage of not requiring contraction mapping to solve for the value functions in every stage of iteration, which speeds up estimation considerably.

The Arcidiacono and Jones (2003) EM approach introduces a computationally easier iterative

[16]Given that we observe only cross-sectional data on hospital characteristics, we assume that the distribution of surgeon characteristics at a hospital are stable over the period of analysis. Some surgeons may quit and be replaced during this period, but these transitions are likely to be secondary relative to the variables we emphasize. Also, considering the relatively long time intervals between the introduction of new products compared to the period of our study in this market, we assume that the buyer does not expect a new entry in the market.

method to estimate heterogeneous preferences by separating the estimation of utility preferences for each segment within an optimization step and describing the probability of belonging to a segment as a Bayesian update based on estimated utility preferences. The estimation approach is amenable to parallelization techniques, and we exploit these to speed up estimation.

We now describe how we adapt the procedure developed in Mojir and Sudhir (2021) to estimate the model of organizational buying. A key element of the organizational buying model is that there are multiple agents (surgeons and buyers in our setting) whose choices and interactions are considered in explaining organizational buying outcomes. Essentially, we estimate the surgeon model and buyer model parameters separately and iteratively, conditional on the parameters estimated for the other agent in the previous iteration or step. We begin by outlining the overall estimation algorithm and subsequently describing more details of the procedure.

The estimation algorithm is as follows:

1. Assume a certain number of segments. Start with initial values of buyer and surgeon model preference parameters for each segment and the segment probabilities for each hospital.

2. Use generalized least squares (GLS)[17] to estimate surgeon model parameters for each segment, θ_s^k, conditional on segment probabilities for each hospital.

3. Estimate buyer model parameters for each segment, using estimated surgeon parameters from step 2 and conditional probabilities of hospitals being in each segment.

4. Update the conditional probabilities of a hospital being in each segment, using estimates of the surgeon model from step 2 and buyer model estimates from step 3.

5. Go to step 2 and iterate until convergence.

We first explain the estimation procedure for the surgeon model in step 2. Next, we present the estimation for the buyer's dynamic model using MPEC approach in step 3. Finally, we discuss the calculations for conditional probabilities of segment membership in step 4 and conclude with

[17]We use GLS because the surgeon model has no endogenous variables (like price).

a brief discussion of identification. The EM algorithm allows parallelization and helps speed up step 3, estimating the dynamic buyer structural model with heterogeneity.[18]

The Surgeon Model

To estimate the surgeon model, we need to obtain values of ξ_{htj} in equation 1. The standard approach invovles solving the system of nonlinear equations through contraction mapping (Berry 1994). But to speed up estimation, we implement the estimation of the surgeon model as a Mathematical Program with Equilibrium Constraints (MPEC), following Dubé et al. (2012), by solving the following constrained optimization problem with the GLS objective function:

$$\min_{\theta_s^k} (\pi_k \circ \xi(\theta_s^k))' \Sigma^{-1} \xi(\theta_s^k)$$

subject to:

$$shr_{hjt} = \hat{shr}_{hjt} \qquad \forall h, t, j$$

$$P(T_{s_h t} = 1) = 0 \qquad \forall t \leq t_0$$

$$P(T_{s_h t} = 1) = P(T_{s_h (t-1)} = 1) + (1 - P(T_{s_h (t-1)} = 1)) P^{NT}_{s_h (J+1) t} \qquad \forall t > t_0$$

Here, θ_s^k represents the set of parameters for surgeons' model conditional on being in segment k, and π_k is a vector of conditional probabilities of being in segment k for each hospital. This vector of conditional membership probabilities gets updated in each iteration at step 4 of the estimation algorithm. $(\pi_k \circ \xi(\theta_s^k))$ is the element-wise (Hadamard) product of conditional membership probabilities and errors. Σ is the appropriate (GLS) weight matrix.

Note that the market share equations are imposed in the first constraint as in Dubé et al. (2012). What is novel here is imposing surgeons' adoption process as constraints to the optimization prob-

[18]Note that there are six prices to track: the off-contract and discount price for each of the three products. Given the need to discretize prices in solving the dynamic program, the state space explodes combinatorically with six prices with fine granularity making estimation computationally intractable. Hence, one has to use very coarse price discretization. Alternatively, we considered more granular discretization of only three off-contract prices, but treated the discount percentage relative to the off-contract price as constant for each hospital. After testing both approaches in initial estimation, we found that point estimates with both approaches are close, but the second approach yields tighter standard errors. The reported results in the paper are based on the second approach.

lem. This is equivalent to solving the nonlinear system of equations using contraction mapping to retrieve values of ξ and then using those values to solve the minimization problem iteratively.

We allow for observed heterogeneity among surgeons by categorizing them into high-share and low-share surgeons, based on their share of total surgeries in the category at each hospital. High-share surgeons may have different preferences for contract benefits. If a surgeon has an especially high share of surgeries, the off-contract supplier might provide a support representative just for that surgeon, in which case the surgeon would not care about using the off-contract product; however, if such accommodations are not available, it might be more costly for a surgeon to forgo the benefits of the contract. Therefore, the differential preference for contract benefits for high-share surgeons is an empirical question.

The Buyer Model

The buyer has two different options at each time period: sign a contract and commit to Supplier A, or sign a contract and commit to Supplier B (the innovator).[19] Using additive separability of the unobserved error shocks in the utility of the buyer, and assuming that it follows an i.i.d. Type I extreme value distribution, the probability of choosing each option based on the observed part of choice-specific value functions, conditional on estimated parameters for surgeons is:

$$P(d_{ht}|X_{ht};\hat{\theta}_s^k,\theta_b^k) = \frac{\exp(\overline{v}(X_{ht},d_{ht};\theta_b^k,\hat{\theta}_s^k))}{\sum_{j\in\{1,..,J\}}\exp(\overline{v}(X_{ht},j;\theta_b^k,\hat{\theta}_s^k))} \tag{14}$$

Here, θ_b^k is the set of parameters for the buyer model for segment k. We can write the conditional likelihood for each hospital using the above choice probabilities. This likelihood is conditional on

[19]It is sub-optimal to not commit to either supplier. In practice, it does not happen.

the estimated parameters for the surgeon model and also the hospital belonging to segment k:

$$\begin{aligned} L_{h|k;\hat{\theta}_s^k}(X_{h1}, X_{h2}, ..., X_{hT}, d_{h1}, d_{h2}, ..., d_{hT}; \theta_b^k) &= \prod_{t=1}^{T} P(X_{ht}, d_{ht}|X_{ht-1}, d_{ht-1}; \theta_b^k, \hat{\theta}_s^k) \\ &= \prod_{t=1}^{T} P(d_{ht}|X_{ht}; \theta_b^k, \hat{\theta}_s^k) P(X_{ht}|X_{ht-1}, d_{ht-1}; \hat{\theta}_s^k) \end{aligned} \tag{15}$$

The first equality assumes the Markovian property; the second equality arises from the conditional independence assumption. Transforming the likelihood into log form yields two additively separable components:

$$\log L_{h|k;\hat{\theta}_s^k} = \sum_{t=1}^{T} \log P(d_{ht}|X_{ht}; \hat{\theta}_s^k, \theta_b^k) + \sum_{t=1}^{T} \log P(X_{ht}|X_{ht-1}, d_{ht-1}; \hat{\theta}_s^k) \tag{16}$$

Note that the second term in the above equation does not depend on θ_b^k. Hence, to simplify the notation, from this point forward we drop the second term.

We formulate the estimation of the buyer's parameters as a constrained optimization problem. That is, instead of using contraction mapping approach to solve for the dynamic programming problem in each iteration, we will impose Bellman equations as constraints in maximizing the likelihood. In this case, choice probabilities would be written as functions of not only model parameters and state variables, but also value functions.

More formally:

$$\begin{aligned} &\hat{\theta}_b = \arg\max_{\theta_b} \sum_{k=1}^{K} \sum_{h=1}^{H} \sum_{t=1}^{T} P(k|X_h; \hat{\theta}_s^k, \theta_b^k, \hat{p}) \log P(d_{ht}|X_{ht}, \overline{v}, EV; \hat{\theta}_s^k, \theta_b^k) \\ &\text{subject to:} \\ &EV(X_{ht}, \hat{\theta}_s^k, \hat{\theta}_b^k) = \log\{\exp(\overline{v}(X_{ht}, d_{ht} = 0; \hat{\theta}_s^k, \hat{\theta}_b^k)) + \exp(\overline{v}(X_{ht}, d_{ht} = 1; \hat{\theta}_s^k, \hat{\theta}_b^k))\} \\ &\forall k \in \{1, ..., K\}, \forall h \in \{1, ..., H\}, \forall t \in \{1, ...T\} \end{aligned} \tag{17}$$

where $P(k|X_h; \hat{\theta}_s^k, \hat{\theta}_b^k, \hat{p})$ represents the probability of hospital h being in segment k, conditional on observed state variables, set of parameters, and unconditional probabilities of being in each segment, $\hat{p}$ (i.e., segment size). θ_b includes parameters for buyers in all segments, $\theta_b = \bigcup\limits_{k=1,\ldots,K} \theta_b^k$.

Note that in equation 17, the maximand could be written as the sum of multiple terms, each being a function of parameters for one segment. Given that there is no dependence in parameters across segments at this stage, instead of maximizing over all the parameters, we can maximize over parameters of each segment separately. This design speeds up the estimation because the sub-problems are faster to solve and can be solved in parallel.

Conditional Probabilities of Segment Membership

In Step 2 of the estimation procedure, we calculate the probability that each hospital is a member of each segment, conditional on a set of parameters and segment sizes. Following Arcidiacono and Jones (2003), we can write:

$$P_h(k|X_h;\hat{\theta}_s^k,\hat{\theta}_b^k,\hat{p}) = \frac{\hat{p}_k L_{h|k}(X_h;\hat{\theta}_s^k,\hat{\theta}_b^k)}{\sum\limits_{m=1,\ldots,K} \hat{p}_m L_{h|m}(X_h;\hat{\theta}_s^m,\hat{\theta}_b^m)} \tag{18}$$

Where $\hat{p}_k$ will be updated in each iteration based on:

$$\hat{p}_k = \frac{1}{H}\sum_{h=1}^{H} P_h(k|X_h;\hat{\theta}_s^k,\hat{\theta}_b^k,\hat{p}) \tag{19}$$

Since the estimates of the surgeon parameters in the first step of the process are input for estimating the buyer parameters in the second step, we need to consider the estimation error for surgeon parameters when calculating the standard errors of the buyer model estimates. Instead of deriving an analytical expression for the standard errors of our estimates, we use bootstrapping. We take 20 samples, with replacement, from our original set of hospitals. Each sample has the same number of hospitals as our original set. We then estimate the model for each sample and calculate the standard errors for our estimates using:

$$se(\hat{\theta}) = \sqrt{\tfrac{1}{n_b-1}\sum_{e=1}^{n_b}(\hat{\theta}_e-\bar{\theta})^2}$$

where n_b represents the number of samples used for bootstrapping, $\hat{\theta}_e$ is the estimated value of parameters for sample e, and $\bar{\theta}$ is the mean of estimated values of parameters across all the samples.

Identification

We present an informal discussion of identification of the model parameters starting with the user model. Unlike the common diffusion models of durable goods, wherein each consumer makes one purchase and exits the market, we model a disposable category involving trial, adoption, and ongoing product usage. As product quality tends to be relatively stable and surgeons do not directly consider price in choosing a product, the diffusion over time arises from the heterogeneity in the surgeon's timing of trial and the buyer's contract switching (through its step function impact on surgeon utility). Therefore, the gradual increase in share of the new product over time when there has been no contract switching allows us to identify the surgeon trial cost separately from product preferences. Further, the stable differences in long-term shares across the three products after contract switching and user adoption have stabilized identify the relative overall preferences for the three products.

The benefit from contracts for surgeons is identified by the change in product share after hospitals switch contracts. The effect of contract switches on shares varies by share of high-share surgeons at the hospital; this identifies the differential benefits of contract by surgeon type. Cross-sectional variation in observed hospital characteristics (e.g., teaching hospital) identifies differential preferences based on these characteristics.

Since surgeons' product choices in each period are conditional on the buyer's contract choices in that period, surgeon model parameters are identified independently of buyer utility parameters. Given surgeon model parameters, each hospital's purchasing cost and surgeon dissatisfaction and their evolution are known. Variation in hospitals' purchasing costs helps identify the buyer's price sensitivity. This variation arises from differences in hospitals' product share compositions, and differences in prices that hospitals face (mainly due to the status of compliance with SOW thresholds, hospital size, and changes in contracts after periodic re-negotiations between suppliers and GPOs). The surgeon dissatisfaction (i.e., IVL) varies across hospitals and time depending on which supplier is under contract, the composition of surgeons, and the share of surgeons who have tried the new product in each hospital. This variation identifies the coefficient of IVL (i.e., λ). Finally, contract switching cost (σ^b) is identified from the timing of contract switching, considering the net gain from switching given both purchasing cost and user dissatisfaction from the surgeon model estimates.

RESULTS

We begin by reporting the model estimates, then discuss the counterfactual analyses.

Model Estimates

Table 5 presents the estimates of the structural model with two segments. All parameters are statistically significant either at the 1% or 5% level, except the buyer's aversion to surgeons' dissatisfaction (i.e., λ) in Segment 1, which is significant at 10% level. We estimated the model with one, two, and three segments; the two-segment buyer model fits best based on AIC and BIC.

Table 5: Model estimates

		Segment 1 / Buyer-Centric	Segment 2 / User-Centric
User (Surgeon)			
Pref for Supp B Old	μ_{old}	-0.218***	-1.261***
		(0.0211)	(0.0974)
Pref for Supp B New	μ_{new}	1.129***	0.730***
		(0.1344)	(0.0618)
Teaching Hospital	τ	0.749***	0.162***
		(0.1734)	(0.0424)
Low-Share Contract Utility	γ^L	3.381***	2.253***
		(0.3442)	(0.4518)
High-Share Contract Utility	γ^H	2.350***	0.915**
		(0.8789)	(0.4472)
Trial Cost	σ_{s_h}	-1.919***	-1.529***
		(0.6504)	(0.3822)
Buyer			
Price	α	-0.003***	-0.008***
		(-0.0001)	(-0.0016)
Contract Switching Cost	σ^b	-6.748***	-3.636***
		(0.2331)	(0.2623)
LIV	λ	-0.087*	-0.297**
		(0.0459)	(0.132)
Segment Size		68%	32%
		(3.09)	(3.03)

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The first segment makes up more than two-thirds of the sample. We begin with discussion of the surgeon utility function. Segment 1 has higher preference for the new product μ_{new} =1.129 relative to zero for Supplier A's product, but Segment 2 values the innovation relative to Supplier B's old product much more (μ_{new} relative to μ_{old}). Additionally, users in teaching hospitals value the innovation more given the positive τ.[20]

As expected, low-share surgeons of both segments have a higher utility for contract (i.e., low-share surgeons value the contract benefits more, $\gamma^L > \gamma^H$ for both segments). Further, the contract utility is greater in Segment 1 (i.e., estimated values for both γ^L and γ^H in Table 5 are larger for Segment 1 compared to Segment 2), consistent with the notion that overall users have less power in hospitals belonging to Segment 1 than in Segment 2. Finally, the trial cost has a negative coefficient as expected.

Buyers in Segment 2 are more price sensitive, whereas hospitals in Segment 1 are more reluctant to switch contracts (larger magnitude of σ_b) and less affected by user dissatisfaction (smaller magnitude of λ). Overall, these differences in λ, σ_b, γ^L and γ^H all suggest that buyers have more influence in Segment 1 relative to Segment 2.

Table 6: Segments' observed behavior

Segment Behavior	Buyer-Centric	User-Centric
Probability of contract switching	2.9%	42.2%
Membership probability conditional on contract switching	12.9%	87.1%
Teaching hospital	35.3%	31.3%
Percentage of high-share surgeons	13.3%	28.8%
Surgery concentration (HHI)	0.0926	0.1557

Note: This table presents key dimensions of each segment's observed behavior. To calculate the numbers presented in this table, we computed the average value for each variable across all the hospitals weighted by each hospital's probability of being a member of each segment over the total number (i.e., sum of membership probabilities) of hospitals in the segment. The membership probabilities are calculated using equation 18 in step 3 of the estimation algorithm.

Table 6 compares the observed behaviors of the two latent segments. Given that the innovation from Supplier B is superior in quality but more expensive, buyers who care more about price are more likely to stay with Supplier A, but surgeons who care about product superiority are more likely to switch to the new supplier. Therefore, we expect the user-centric segment to switch to the

[20]The differential preference parameter for high and low share surgeons was insignificant, so we omitted it from the final specification.

innovation faster. In fact, the probability of switching for the user-centric segment is much higher at 42.2%, and negligible at 2.9% for the buyer-centric segment.[21] As anticipated, conditional on switching, the probability of being in the user-centric segment is 87.1%. In terms of observable characteristics, teaching hospitals have a larger share in the buyer-centric segment, but as expected, the concentration of surgeries measured in either percentage of high-share (i.e., superstar) surgeons or HHI index is significantly higher in the user-centric segment.

Counterfactual Analysis

We consider two counterfactuals. The first seeks to reveal how push-pull marketing tactics targeted towards buyers versus users in an organizational buying setting impact market outcomes. The second investigates how SOW discounts affect the diffusion of innovations; specifically, we assess the conventional wisdom that SOW contracts always favor dominant firms, even when smaller firms launch superior products.[22]

Marketing Mix and Push-pull Tactics

An important marketing question in organizational buying contexts is whether and when sellers should focus on the push tactics targeted to buyers versus pull tactics targeted to end-users (surgeons). This counterfactual explores the differential impact of push-pull tactics. We operationalize the pull tactic with a 10% increase in the coefficient of user preference for supplier B's new product (i.e., μ_{new}). We operationalize the push tactic targeted towards the buyer with price cuts, specifically a 10% price cut for Supplier B's new product. We simulate innovation diffusion under these counterfactuals and compare against the benchmark with estimated user preferences and observed prices. The results are presented in Figure 7.[23]

[21] These numbers for switching decisions simulated by the model based on the estimated parameters are 40.2% and 2.2%, respectively. The closeness of the observed and simulated moments indicate that the model fits well.

[22] In the Web Appendix, we report an additional counterfactual to isolate the role of user "dissatisfaction" (through IVL) on buyer. We do this by setting the coefficient of IVL to 0, thus only price and the direct effect of buyer contracts on surgeon choice are present.

[23] In practice, the pull tactic of enhancing the perceived value of the product among surgeons may be implemented through advertising, product demonstrations, or directly enhancing product quality. For push tactics, cutting prices is often used as buyers are responsible for managing procurement costs, but buyer-valued service metrics such as lead times and ease of ordering could also be used. We recognize that counterfactual changes will be accompanied by price changes from the same supplier or competitors. We show that our qualitative conclusions related to SOW contracts

Figure 7: Change in diffusion of new product with price cuts and product quality improvement

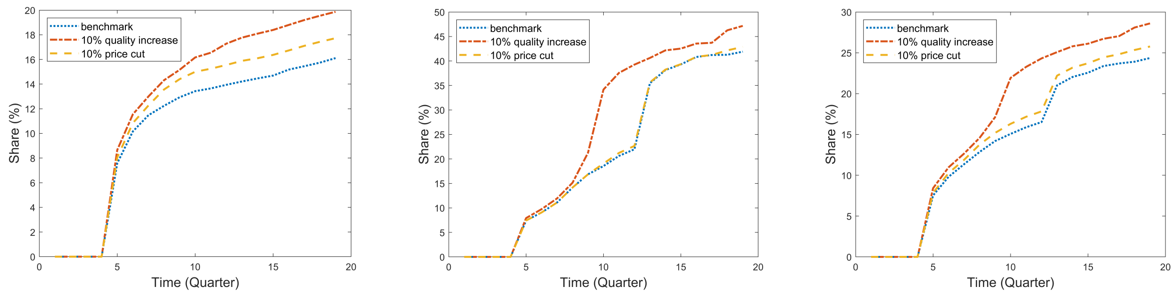

(a) The buyer-centric segment (b) The user-centric segment (c) The market

Note: This figure shows the results of a simulation for the diffusion of Supplier B's new product for: (1) the benchmark; (2) 10% increase in the quality of the new product; and (3) 10% reduction in the price of the new product. A 10% quality improvement (pull strategy) increases the diffusion speed much more than a 10% price cut (push strategy).

To understand the results, it is useful to recognize that buyers are more sensitive to price cuts, whereas surgeons are more sensitive to product improvements. We begin our discussion with the buyer-centric segment. Interestingly, although buyers care directly about price, both price cut and preference increase lead to faster and higher penetration in the buyer-centric segment. This is because even when Supplier B offers a price cut, Supplier A's product is still cheaper and preferred by the buyer. So, the buyer has little incentive to switch to Supplier B; at most hospitals, Supplier B's shares do not exceed the 20% threshold even 15 quarters after the innovation launch. In the few hospitals where usage exceeds 20% (mostly higher surgery concentration and teaching hospitals), the discount speeds up contract switching to Supplier B, leading to an increase in average shares with price cuts compared to the benchmark. In contrast, with the preference increase, the 20% threshold is met more often due to the pull from the surgeons. Thus, if a smaller player has to succeed in buyer-centric hospitals that received SOW discounts, increasing user product preference is more effective; price cuts are not that effective when the innovation is more expensive.

Next, we discuss the user-centric segment. Interestingly, here the price discount has almost no effect on innovation diffusion, as buyers have little ability to change surgeon behaviors, and the price cut does not increase surgeon usage. Therefore, the SOW discount has little impact on contract switching, and overall price discounts do not impact diffusion at all. In contrast, the

remain robust with various price response scenarios in the Web appendix.

improvement in product preference leads to greater adoption among surgeons, and their choices lead to accelerated contract switching among hospitals in this segment.

Overall, we conclude that an innovating small supplier selling a relatively higher quality, more expensive product in markets with SOW contracts will find it better to invest first in pull tactics that increase product preference (e.g., surgeon education, more user benefits) rather than push tactics (e.g., price discounts or buyer-focused service). Further, the pull strategy is more effective in the user-centric segment. Within the user-centric segment, it is important early on to seed the diffusion and generate the pull effect by marketing to surgeons to increase their adoption. But once close to the SOW contract threshold of 20%, adopting push tactics that persuade buyers to switch contracts helps accelerate diffusion.[24] The counterfactual illustrates how our organizational buying model can inform which marketing tactics to adopt for which stakeholder and when. Consider especially that although the user-centric segment has higher price sensitivity, without the pressure from users to switch, the price cut has lower impact than on the buyer-centric segment.

Share of wallet discounts and the diffusion of innovation

We next explore how share of wallet discounts impact innovation diffusion, especially those from smaller firms. As discussed, this counterfactual is interesting given that extant theory research concludes that SOW contracts favor dominant firms. It is also motivated by the smaller Supplier B's concern that SOW discounts are intrinsically anti-competitive and suppress innovations from smaller firms, which could warrant regulatory actions or remedies.[25]

We simulate the new product diffusion under two counterfactual scenarios: (i) reducing the SOW thresholds for both suppliers by 10% (from 80% to 70% and 60% to 50% for Supplier A and B, respectively); (ii) setting SOW thresholds to zero. In both scenarios, we allow additional services (e.g., better stocking, support for product use during surgery) under a contract as before;

[24] To isolate the roles of contract switching by buyers versus the increased adoption by surgeons on accelerated diffusion when quality is increased, we analyzed a counterfactual whereby surgeon adoption adjusts in response to the innovation while contract switching remains the same as in the data. We find that without contract switching, diffusion acceleration is minimal, suggesting that whereas surgeon adoption is necessary to seed, contract switching is critical for accelerated diffusion beyond the 20% threshold. These results are available in the Web Appendix.

[25] Buyers may place frictions on surgeon adoption of the smaller player's superior innovation to avoid losing the larger firm's SOW discount on entire threshold share of usage. This raises the concern of whether SOW contracts stifle the smaller player's innovation through unfair status quo advantages to larger firms.

thus, the buyer and supplier could still influence the demand for different products in the hospital through the choice of the contract. If the semi-exclusive nature of SOW discounts slows down the diffusion of innovations from smaller players, we would expect to see the adoption of the new product of Supplier B speed up when we lower the SOW threshold or eliminate SOW discounts.

We allow for contract discounts to change when changing SOW thresholds. Maintaining the original off-contract prices, we conduct a grid search over a focal supplier's discounts, conditional on the other supplier's discount choice, and vice versa iteratively until convergence, to arrive at the new equilibrium discount. Not surprisingly, discounts are lower for both suppliers with lower share commitments. The larger Supplier A, with the higher discount threshold of 80%, reduces the discount by 30% (20%) for the no-SOW (low SOW) threshold case. The corresponding discounts for Supplier B with the less stringent threshold of 60% is 10% (5%).[26]

Recall that Supplier A's product is cheaper, and all else being equal the buyer would prefer Supplier A. Hence the buyer-centric segment is likely to be more aligned with Supplier A. Further, though Supplier B's innovation is higher priced, it is considered superior and preferred by surgeons who do not consider price. Hence the innovation is likely to succeed more in the user-centric segment. Figure 8 presents the diffusion of the new product under the two counterfactual scenarios and a benchmark with the SOW thresholds in the market (80% for A and 60% for B). We highlight two key results. First, contrary to the smaller Supplier B's concern that SOW discounts hurt its innovation, the SOW discounts with tighter thresholds speed up diffusion; in fact, the more exclusive the contract, the faster the diffusion. Second, despite their price sensitivity, SOW thresholds and discounts have little impact on innovation diffusion in the buyer-centric segment. The faster diffusion occurs due to the user-centric segment.

Clearly, buyers prefer Supplier A's product even without SOW discounts (due to lower off-contract price). Therefore, buyers in the buyer-centric segment try to suppress the innovation's

[26]We describe the details of our grid search to find adjusted prices under new SOW thresholds in the Web Appendix. In the appendix, we also empirically show that the change in the SOW threshold is the primary driver of our results, with price changes not having much bite. We note that a complete ban on SOW discounts could potentially result in changes in aspects of the market equilibrium other than prices. Nonetheless, the consistency between the results of the two simulations (i.e., making the contracts less exclusive, or completely nonexclusive) reassures us that the semi-exclusive aspect of SOW contracts can indeed help smaller suppliers with major innovations.

Figure 8: The effect of share of wallet contracts on diffusion of the new product.

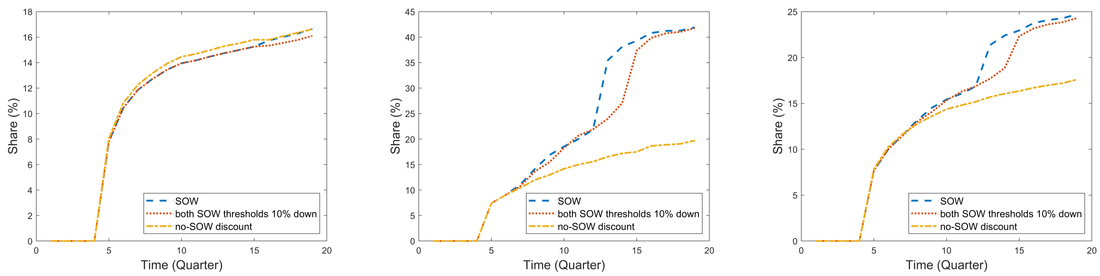

(a) The buyer-centric segment (b) The user-centric segment (c) The market

Note: The figure shows the simulation results for (i) the benchmark scenario with observed SOW thresholds; (ii) the 10% symmetrically reduced SOW thresholds scenario; and (iii) no-SOW threshold scenario. Prices are adjusted to reflect the change in SOW thresholds. Making the contract less exclusive by lowering the SOW threshold slows down the diffusion of innovation.

diffusion by contracting with Supplier A; they can persuade surgeons to continue using Supplier A's product because contract benefits are highly valued by surgeons in the buyer-centric segment. Hence share of Supplier B's innovation never exceeds 20% at these hospitals, so there is little difference in the innovation diffusion trajectory either with or without SOW discounts in this segment.

For the user-centric segment, where surgeons have greater influence and contract benefits are less valued, more surgeons are willing to try the innovation, and buyers are less able to use contracts to restrict them. Forward-looking buyers switch contracts as Supplier B's share reaches the SOW non-compliance threshold for Supplier A and the hospital cannot receive discounts from Supplier A; contract switching speeds up surgeon trial of the innovation and helps the hospital reach the Supplier B's threshold faster. For both the benchmark and reduced SOW threshold cases, the innovation diffusion accelerates at two points compared to the no-SOW threshold case. Around Q7, hospitals with higher surgery concentration (and some teaching hospitals) reach the non-compliance threshold for Supplier A; forward-looking buyers switch contracts to reach the Supplier B's threshold faster. For the benchmark case, other hospitals pass the 20% threshold (the non-compliance threshold for Supplier A) at Q12, and switch contracts. For the lower SOW threshold case, this passing of non-compliance threshold is delayed by two quarters, as the buyer now has more room before hitting the Supplier A's non-compliance threshold (30% instead of 20%). Nonetheless, the maximum share after 19 quarters for both benchmark and lowered SOW threshold cases is the same. In the no-SOW case, however, the buyer does not have much incentive to switch contracts, because the hospital pays Supplier A's very low discounted prices regardless of the supplier's share, as long as it has a contract with the supplier. This lack of incentive for buyers leads to fewer hospitals switching contracts to Supplier B, and explains the more gradual diffusion of the new product under the no-SOW scenario. Overall, with SOW thresholds, Supplier B reaches about 40+% share of sales in the user-centric segment and about 25% of market share. Our result contrasts with the extant conclusions in the theoretical literature based on unitary buyer models that SOW contracts always favor dominant firms; it clarifies the importance of accounting for the multi-stakeholder nature of organizational buying to generate the correct insights around SOW contracts due to its different impacts on users and buyers.

CONCLUSION

We develop a structural model of organizational buying behavior accommodating the preferences and interactions between multiple stakeholders (users and buyers). It is particularly applicable in routinized exchange relationship (RER) contexts, where a buyer periodically evaluates and chooses contracts and downstream users choose between on- and off-contract products, conditional on the buyer's decision. The model accounts for the effect of buyers' contracting decisions on user trial and ongoing usage of the innovation. Similarly it allows for user influence on buyers. A key idea is how we account for user influence on the buyer through an internally consistent utility-based construct of user dissatisfaction with the buyer-chosen contract ("inclusive value loss"). We contribute to the innovation diffusion literature by introducing a novel structural model for diffusion of disposables in a trial and repeat context, in contrast to past structural models of diffusion focused on one-time durable goods purchases. Finally, the paper introduces an empirical analysis of SOW contracts—a common contract in B2B markets.

In terms of findings, our paper sheds insights on the effective use of push-pull tactics in B2B markets. For example, it illustrates how a smaller innovator with a major innovation that is liked by users can use pull tactics to build early market share and then shift to push tactics once it has breached the SOW threshold. Further, in contrast to conventional wisdom that SOW contracts always favor larger firms, we show that in multi-stakeholder settings (involving a buyer and users), SOW contracts can help accelerate innovation diffusion even from smaller firms as long as they are major innovations that are strongly preferred by end-users. This finding should give regulators pause on treating SOW contracts as per se anti-competitive and antithetical to innovation in markets where small manufacturers drive innovation.

We conclude with a discussion of limitations and suggestions for future work. Although our main contribution is the development of an empirical framework to model organizational buying behavior under SOW contracts, the application is in one category: a disposable surgical device product sold to hospitals and used by surgeons. Future work should extend the model and analysis for other product categories across a range of B2B markets to yield richer insights about how different user, buyer, and influencer interactions shape organizational buying outcomes. Similarly, it would be useful to assess the generalizability of our conclusions about SOW contracts and their effect on innovation diffusion in settings with other market structures.

As our model applies to buying modes involving routinized exchange relationships (RERs), we hope future research can employ our framework in other RER contexts; some examples include procurement from original equipment manufacturers (OEMs) by downstream producers or service contracting (e.g., IT maintenance contracts). The buying center in these cases usually consists of multiple stakeholders with diverging preferences and priorities (e.g., design engineers and sourcing directors or employees and IT departments) and there is an ongoing use of suppliers/service providers based on the user-specific needs as they arise. Further, with novel technologies such as Internet of Things (IoT) and cloud-based services, the usage data beyond purchase and contracting choice in B2B markets are becoming easily available (e.g., Mojir and Yucaoglu 2021). Overall, we hope our modeling framework in combination with easier data availability aid addressing the challenges of econometric modeling of B2B markets (raised in Lilien 2016) and thus catalyze rapid growth of B2B research.

REFERENCES

Arcidiacono P, Jones JB (2003) Finite mixture distributions, sequential likelihood and the em algorithm. *Econometrica* 71(3):933–946.

Aribarg A, Arora N, Bodur HO (2002) Understanding the role of preference revision and concession in group decisions. *Journal of Marketing Research* 39(3):336–349.

Arora N (2006) Estimating joint preference: A sub-sampling approach. *International Journal of Research in Marketing* 23(4):409–418.

Arora N, Allenby GM (1999) Measuring the influence of individual preference structures in group decision making. *Journal of Marketing Research* 36(4):476–487.

Bass FM (1969) A new product growth for model consumer durables. *Management science* 15(5):215–227.

Berry S, Levinsohn J, Pakes A (1995) Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society* 841–890.

Berry ST (1994) Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics* 242–262.

Bruno HA, Che H, Dutta S (2012) Role of reference price on price and quantity: insights from business-to-business markets. *Journal of Marketing Research* 49(5):640–654.

Calzolari G, Denicolò V (2013) Competition with exclusive contracts and market-share discounts. *American Economic Review* 103(6):2384–2411.

Corfman KP (1989) Measures of relative influence in couples: A typology and predictions for accuracy. *ACR North American Advances* 16.

Dubé JP, Fox JT, Su CL (2012) Improving the numerical performance of static and dynamic aggregate discrete choice random coefficients demand estimation. *Econometrica* 80(5):2231–2267.

Fourt LA, Woodlock JW (1960) Early prediction of market success for new grocery products. *The Journal of Marketing* 31–38.

Gravelle JG, Lowry S (2015) *The Medical Device Excise Tax: Economic Analysis*.

Grennan M (2013) Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review* 103(1):145–77.

Grewal R, Lilien GL, Bharadwaj S, Jindal P, Kayande U, Lusch RF, Mantrala M, Palmatier RW, Rindfleisch A, Scheer LK, et al. (2015) Business-to-business buying: Challenges and opportunities. *Customer needs and Solutions* 2(3):193–208.

Hahn M, Park S, Krishnamurthi L, Zoltners AA (1994) Analysis of new product diffusion using a four-segment trial-repeat model. *Marketing Science* 13(3):224–247.

Horsky D (1990) The effects of income, price and information on the diffusion of new consumer durables. *Marketing Science* 9(4):342–365.

Inderst R, Shaffer G (2010) Market-share contracts as facilitating practices. *The RAND Journal of Economics* 41(4):709–729.

Laczniak GR (1979) An empirical study of hospital buying. *Industrial Marketing Management* 8(1):57–62.

LaPlaca PJ, Katrichis JM (2009) Relative presence of business-to-business research in the marketing literature. *Journal of Business-to-Business Marketing* 16(1-2):1–22.

Lewin JE, Johnston WJ (1996) Business and industrial marketing: past, present and future. *Journal of Business & Industrial Marketing* 11(1):7–16.

Lilien GL (2016) The b2b knowledge gap. *International Journal of Research in Marketing* 33(3):543–556.

Lilien GL, Rao AG, Kalish S (1981) Bayesian estimation and control of detailing effort in a repeat purchase diffusion environment. *Management science* 27(5):493–506.

Majumdar A, Bishop S, Ridyard D, Williams I, Akgun U (2005) Selective price cuts and fidelity rebates. Technical report, Economic discussion paper.

Majumdar A, Shaffer G (2009) Market-share contracts with asymmetric information. *Journal of Economics & Management Strategy* 18(2):393–421.

Mansfield E (1961) Technical change and the rate of imitation. *Econometrica: Journal of the Econometric Society* 741–766.

Mills DE (2010) Inducing downstream selling effort with market share discounts. *International Journal of the Economics of Business* 17(2):129–146.

Moed BR, Israel HA (2017) Device sales representatives in the operating room: do we really need or want them? a survey of orthopaedic trauma surgeons. *Journal of Orthopaedic Trauma* 31(9):e296–e300.

Mojir N, Sudhir K (2021) A model of multipass search: Price search across stores and time. *Management Science* 67(4):2126–2150.

Mojir N, Yucaoglu G (2021) Borusan CAT: Monetizing Prediction in the Age of AI (A). *Harvard Business School Case 521-053* .

Montgomery K, Schneller ES (2007) Hospitals' strategies for orchestrating selection of physician preference items. *The Milbank Quarterly* 85(2):307–335.

Moon J, Tikoo S (2002) Buying decision approaches of organizational buyers and users. *Journal of Business Research* 55(4):293–299.

Nair H (2007) Intertemporal price discrimination with forward-looking consumers: Application to the us market for console video-games. *Quantitative Marketing and Economics* 5(3):239–292.

O'Connor B, Pollner F, Fugh-Berman A (2016) Salespeople in the surgical suite: relationships between surgeons and medical device representatives. *PLoS One* 11(8):e0158510.

Rao VR, McLaughlin EW (1989) Modeling the decision to add new products by channel intermediaries. *Journal of Marketing* 53(1):80–88.

Roberts EB (1988) Technological innovation and medical devices. *New Medical Devices: Innovation, Development, and Use*, 35–47 (Washington (DC): National Academy of Engineering, Institute of Medicine).

Robinson PJ, Faris CW, Wind Y (1967) *Industrial buying and creative marketing* (Allyn and Bacon).

Rust J (1987) Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica: Journal of the Econometric Society* 999–1033.

Schneller ES, Smeltzer LR (2006) *Strategic management of the health care supply chain* (Jossey-Bass).

Sheth JN (1973) A model of industrial buyer behavior. *Journal of marketing* 37(4):50–56.

Sheth JN (1996) Organizational buying behavior: past performance and future expectations. *Journal of Business & Industrial Marketing* 11(3/4):7–24.

Song I, Chintagunta PK (2003) A micromodel of new product adoption with heterogeneous and forward-looking consumers: Application to the digital camera category. *Quantitative Marketing and Economics* 1(4):371–407.

Su CL, Judd KL (2012) Constrained optimization approaches to estimation of structural models. *Econometrica* 80(5):2213–2230.

Sudhir K, Rao VR (2006) Do slotting allowances enhance efficiency or hinder competition? *Journal of Marketing Research* 43(2):137–155.

Tom WK, Balto DA, Averitt NW (2000) Anticompetitive aspects of market-share discounts and other incentives to exclusive dealing. *Antitrust Law Journal* 67(3):615–639.

Turk JL, Bell NW (1972) Measuring power in families. *Journal of Marriage and the Family* May 1:215–222.

Webster FE, Wind Y (1972) *Organizational buying behavior* (Prentice Hall).

Zhang JZ, Netzer O, Ansari A (2014) Dynamic targeted pricing in b2b relationships. *Marketing Science* 33(3):317–337.

WEB APPENDIX

A1. Constructing Supplier A Price Sequence

In this appendix, we describe how we construct the price sequence for Supplier A. We use multiple sources of data to deduce prices charged by Supplier A to each hospital each quarter: GPO affiliations of hospitals; price contracts of GPOs with Supplier A; each hospital's spend with Supplier A (calculated based on hospital's number of surgeries); and Supplier A's share from each hospital's business in the category, which equals 1 - Supplier B's share.

For Supplier B, we observe the letter of commitment (LOC) signed by each hospital that informs us which GPO contract is operational for the hospital. Since we do not observe the actual contract for Supplier A, we assume that hospital uses the GPO contract that would give it the lowest price, from the set of GPOs with which it is affiliated. This assumption is reasonable, because this is consistent with how hospitals choose GPO contracts with Supplier B. Further, given the limited variation across prices among different GPOs, the data are not very sensitive to this assumption, even if there are some violations. To clarify this price deduction process for Supplier A in each hospital-quarter, we further provide a hypothetical example that covers all scenarios observed in the data.[27]

Consider hospital H affiliated with GPOs G1, G2, and G3. In quarter 1 (Q1), these GPOs offer prices $\overline{p}_1$, $\overline{p}_2$, and $\overline{p}_3$, if H is not in contract with Supplier A. If H is in contract with Supplier A and complies with the threshold share of the contract, it would pay $\underline{p}_1$, $\underline{p}_2$, and $\underline{p}_3$. If $\underline{p}_1 < \underline{p}_2 < \underline{p}_3$ and H is committed to Supplier A (we know this is the case because we do not observe H having any contracts with Supplier B), and if Supplier A's share of its business is above 80%, H will pay $\underline{p}_1$.

In Q2, however, suppose Supplier A's share of hospital H's business drops to 65%, below the required share in Supplier A's contract. Because G2 has the lowest off-contract price among all three GPOs (i.e., $\overline{p}_2 < \overline{p}_1$), H will use its affiliation with G2 to pay $\overline{p}_2$ in the second quarter. In

[27]The hypothetical example is used to illustrate all possible scenarios and how we address them when they occur. It is important to recognize that in most cases, a hospital's relevant GPO remains the same for the entire duration of the data. The scenarios discussed here are relatively rare.

Q3, suppose H develops a new affiliation with a new GPO, G4. G4's on-contract price is slightly below G1 (i.e., $\underline{p}_4 < \underline{p}_1$), and Supplier A's share of H's business in the category is once again above 80%. Therefore, H will pay $\underline{p}_4$ for Supplier A's product in the third quarter. In Q4, suppose G3 renegotiates its contract with Supplier A, which offers the lowest price among the four GPOs. Then H will use its affiliation with G3 to pay $\underline{p}_3$ for the fourth quarter. In Q5, suppose Supplier A's share of H's business drops to below 80% and continues to drop, following the launch of the new product by Supplier B. H will then pay $\overline{p}_2$ starting in Q5.

We note also that suppliers offer different prices based on aggregate category spend to discriminate by hospital size. Therefore another hospital, H_1, of a different size with different category spend can pay a different price even if it has same GPO affiliations and share trajectories.

A2. Decomposing the Role of Users vs. Buyers in Counterfactual 1

Here we disentangle the effect of users vs. buyers on the increased diffusion of improved (perceived or actual) product quality by examining the outcomes when the buyers' contracting decisions are fixed and only surgeons adjust their usage. Figure 9 depicts the diffusion of the new product under the two scenarios. Although users do have a small impact in increased diffusion speed for the better product, most of the impact arises from the buyers' contract switching. Thus buyer contract switching that provides a boost to surgeon utility turns out to be critical even when surgeons experience higher utility from the higher quality product.

Figure 9: The share of buyer vs. users of increased diffusion from quality improvement.

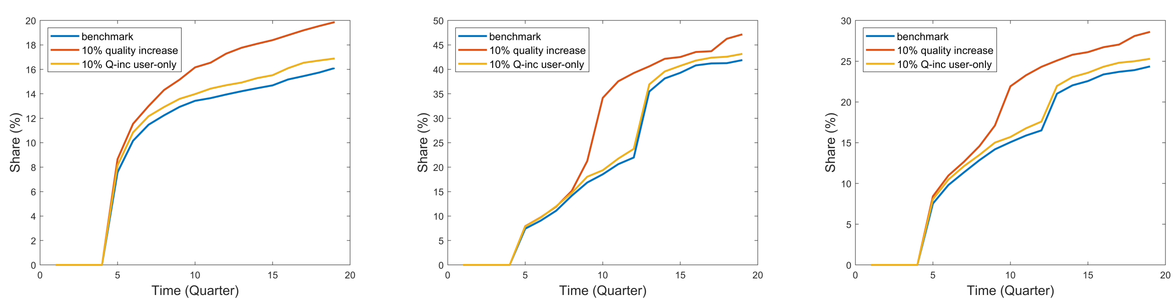

(a) The buyer-centric segment (b) The user-centric segment (c) The market

A4. Price Adjustments Under lower or no-SOW Thresholds

With lower SOW thresholds, we recognize that price discounts will adjust and be smaller in equilibrium. To keep the computational burden of finding the new price equilibrium manageable, we keep the list (off-contract) price fixed and assume that empirically observed discounts in each quarter will adjust by a fixed supplier-specific multiplier across all quarters. This preserves the shape of the price trajectory and other correlations in prices across suppliers over time.

We conduct a grid search to obtain the supplier-specific price multiplier in equilibrium assuming that each supplier is profit maximizing.[28] For the grid search, we keep the prices of Supplier A constant and calculate the profit of Supplier B for different multipliers. We choose the multiplier that yields maximum profit for Supplier B. Keeping Supplier B's multiplier at that (maximizing) level, we calculate Supplier A's profit for various Supplier A multipliers and find the multiplier that maximizes Supplier A's profit. We repeat the above steps until neither supplier finds it optimal to change its multiplier, given the other supplier's multiplier.

Figure 10: Counterfactual with no SOW threshold; with and without equilibrium price adjustment.

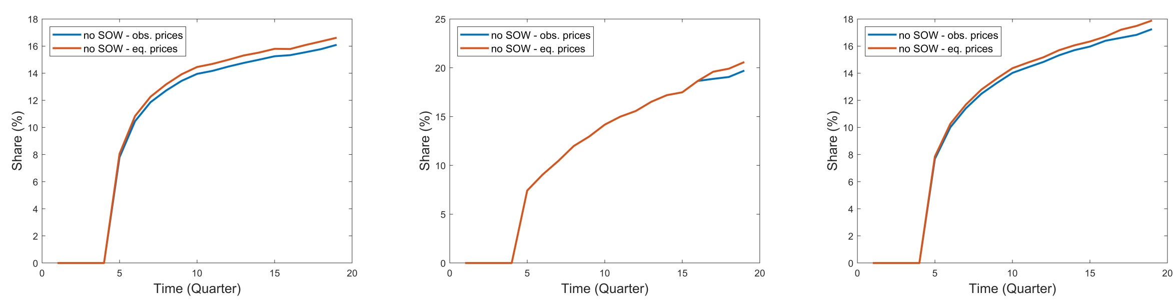

(a) The buyer-centric segment (b) The user-centric segment (c) The market

To assess the relative importance of changes in SOW threshold and prices on diffusion speed, we compare the counterfactual outcomes with only SOW threshold changes and no price adjustments. Figure 10 compares the diffusion of the new product under no-SOW scenario with and without price adjustment. Equilibrium price adjustments have limited effect on the diffusion of the new product under the no-SOW scenario. Thus our qualitative conclusions about changes in speed

[28] To compute profits, we infer supplier costs based on gross margins in the 10K statements of each supplier and assume that margins are the same across all products of each supplier.

Figure 11: The result of various adjustments to price of the new product of Supplier B as its quality improves by 10%.

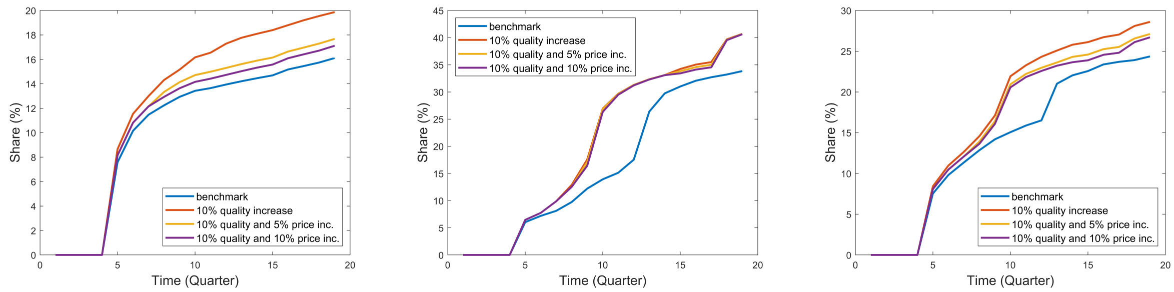

(a) The buyer-centric segment (b) The user-centric segment (c) The market

of diffusion of the innovation from the smaller supplier can be attributed to the changes in SOW thresholds rather than price changes.

A5. Robustness of the Push-Pull Counterfactual to Price Adjustments

In this section, we present the results of the push-pull counterfactual analysis under various assumptions of price changes. It is conceivable for Supplier B to increase the price of its new product as it improves the quality of the product by 10%. To check the sensitivity of our results to such potential price adjustments, we simulated the diffusion of the new product under three different scenarios, besides the benchmark (i.e., estimated level of quality). First, we assumed there is only a 10% increase in the quality of Supplier B's new product without any price adjustments (i.e., what is presented in the counterfactual analysis of the paper). Second, we assumed that Supplier B would increase the price of its new product by 5% to account for the fact that it has had a 10% quality improvement. Third, we assumed a 10% price increase for the new product of Supplier B as an adjustment for its (10%) higher quality. Figure 11 presents the results. As we see in the Figure, price adjustments have a limited effect on the diffusion of the new product overall, with almost no effect on the user-centric segment, which is the segment that drives the diffusion.

Moreover, it is also feasible that Supplier A will change the price of its product in response to changes in quality and price of Supplier B's new product. To explore the sensitivity of our results to such price changes, we simulated the diffusion of the new product under similar scenarios as

Figure 12: The result of various adjustments to price of the new product of Supplier B as its quality improves by 10%. Here we change Supplier A's prices in response to Supplier B's quality and price changes.

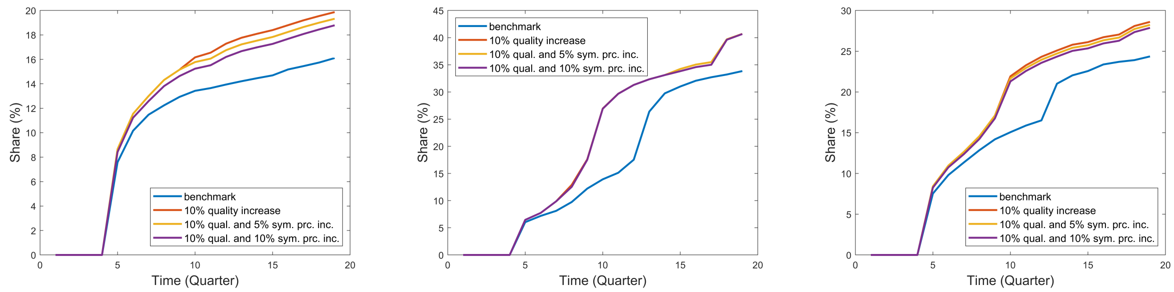

(a) The buyer-centric segment (b) The user-centric segment (c) The market

before, but this time made the same adjustment to the price of Supplier A's product as we did to Supplier B's new product (e.g., increased its price by 5% or 10% accordingly). As Figure 12 shows, price adjustments would have even a more limited effect on the diffusion of the new product if we consider potential price responses from Supplier A.

Finally, it is also conceivable that under the push scenario (i.e., 10% price cut of the new product of Supplier B), Supplier A would respond by adjusting its prices. To explore the sensitivity of our results to such potential price adjustments, we simulated the diffusion of the new product of Supplier B under the push scenario but dropped the price of Supplier A's product by 1) 5%, 2) 10%, and 3) 15%. As Figure 13 shows, such price adjustments have very limited impact on the overall market response, considering the limited role that price plays in diffusion of the new product in this market, in general. In fact, the scenario without any price adjustments (presented in the paper) is the farthest from the benchmark.

Overall, the results of our robustness checks demonstrate that our conclusions about quality improvements having a much larger impact on the diffusion of the new product are robust to assumptions about both Supplier B's potential price adjustment to account for the quality improvement, and Supplier A's potential price responses. We further simulated the diffusion of the new product under other price adjustment assumptions (e.g., changing the price of Supplier A's product in the opposite direction, changing the price of Supplier B's old product to various degrees). Our

Figure 13: The result of various adjustments to price of the Supplier A's product in the push scenario (i.e., 10% price cut of the new product of the Supplier A).

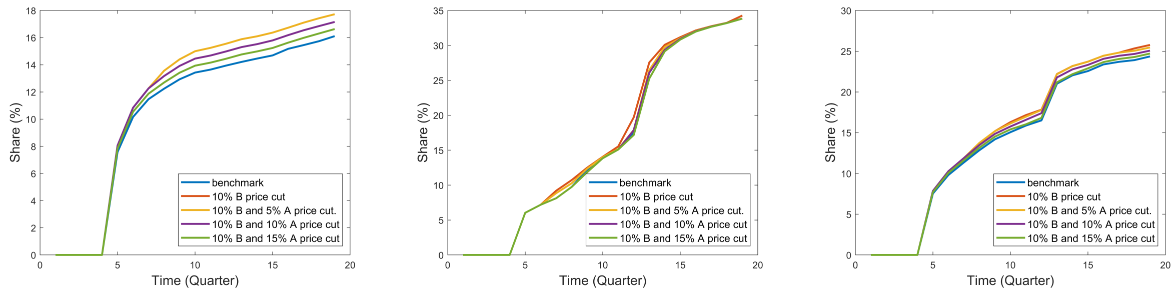

(a) The buyer-centric segment (b) The user-centric segment (c) The market

results are generally robust to such changes and our conclusions about the superior influence of pull vs. push strategy remain intact.

A6. Isolating the Effect of User Dissatisfaction

In this section we present the results of a counterfactual analysis in which we try to isolate the effect of user dissatisfaction on Supplier B's new product diffusion. To eliminate the effect of user dissatisfaction, we set the coefficient on IVL, i.e., λ to zero and simulate the diffusion of Supplier B's new product. Figure 14 shows the results of the simulation across the three segments.

Setting λ to zero reduces the diffusion speed both in the buyer-centric segment and overall. Note that users affect the buyer's contracting decision in two ways. First, users' product choices determine the overall share of each product in the hospital, therefore the price of each product, and consequently, the total purchase cost. Second, users voice their dissatisfaction with the lack of proper support for a product, which could be the result of the buyer's contracting decision. By setting λ to zero, we have eliminated the effect of the second element, hence the reduction in the buyer's contract switching incentive and diffusion of the Supplier B's new product.

Interestingly, setting λ to zero does not affect the diffusion speed in the user-centric segment. Note that in the user-centric segment compared to the buyer-centric segment, users value the contract benefits much less (i.e., γ^L and γ^H for this segment is much smaller compared to the buyer-centric segment); the buyer has much smaller leverage to influence users' choices in this segment.

Additionally, the buyer in the user-centric segment is more price sensitive. Therefore, in this segment the users' influence on the buyer through price overpowers their influence through voicing their dissatisfaction, resulting in setting $\lambda = 0$ having a negligible effect on the buyer's contract decisions and consequently the new product diffusion.

Figure 14: Isolating the effect of user dissatisfaction on diffusion of the new product. Share of Supplier B's new product over time for each segment and for the market.

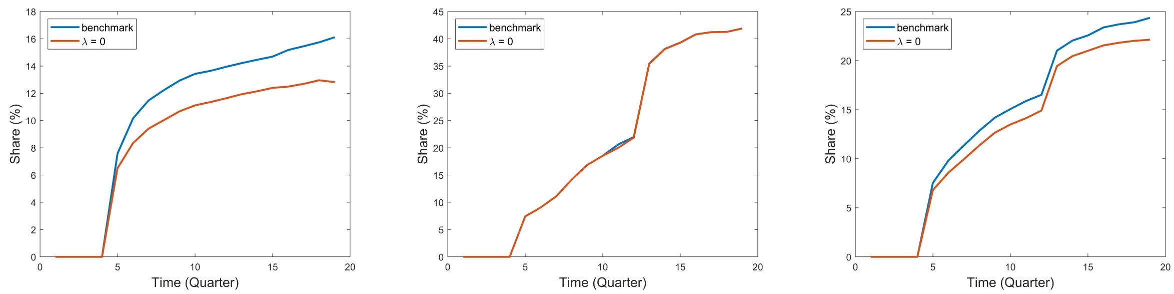

(a) The buyer-centric segment (b) The user-centric segment (c) The market