



Marketing Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Dynamic Targeted Pricing in B2B Relationships

Jonathan Z. Zhang, Oded Netzer, Asim Ansari

To cite this article:

Jonathan Z. Zhang, Oded Netzer, Asim Ansari (2014) Dynamic Targeted Pricing in B2B Relationships. Marketing Science 33(3):317-337. <http://dx.doi.org/10.1287/mksc.2013.0842>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2014, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Dynamic Targeted Pricing in B2B Relationships

Jonathan Z. Zhang

Michael G. Foster School of Business, University of Washington, Seattle, Washington 98195, zaozao@uw.edu

Oded Netzer, Asim Ansari

Columbia Business School, Columbia University, New York, New York 10027 {on2110@columbia.edu, maa48@columbia.edu}

We model the multifaceted impact of pricing decisions in business-to-business (B2B) relationships that are governed by trust. We show how a seller can develop optimal intertemporal targeted pricing strategies to maximize profits over time while taking into consideration the impact of pricing decisions on short-term profit margin, reference price formation, and long-term relationships. Our modeling framework uses a hierarchical Bayesian approach to weave together a multivariate nonhomogeneous hidden Markov model, buyer heterogeneity, and control functions to facilitate targeting, capture the evolution of trust, and control for price endogeneity. We estimate our model on longitudinal transactions data from a retailer in the industrial consumables domain. We find that buyers in our data set can be best represented by two latent states of trust toward the seller—a “vigilant” state that is characterized by heightened price sensitivity and a cautious approach to ordering and a “relaxed” state with purchase behaviors that are consistent with high relational trust. The seller’s pricing decisions can transition buyers between these two states. An optimal dynamic and targeted pricing strategy based on our model suggests a 52% improvement in profitability compared with the status quo. Furthermore, a counterfactual analysis examines the seller’s optimal pricing policy under fluctuating commodity prices.

Keywords: business-to-business marketing; pricing; customer relationship management; hidden Markov models; channel relationships

History: Received: February 19, 2011; accepted: November 19, 2013; Preyas Desai served as the editor-in-chief and Peter Fader served as associate editor for this article. Published online in *Articles in Advance* April 1, 2014.

1. Introduction

The business-to-business (B2B) sector plays a major role in the United States and world economy. B2B transactions command more than 50% market share of all commerce within the United States (Dwyer and Tanner 2009, Stein 2013). Despite their obvious importance, B2B issues have received scant attention in the modeling literature within marketing. Only a small fraction (approximately 3.4%) of the articles published in the top four marketing journals deal with B2B contexts (LaPlaca and Katrichis 2009). Compared with other marketing decisions, the topic of pricing in B2B environments is particularly underresearched (Liozu 2012, Reid and Plank 2004). In this paper, we address this imbalance by developing an integrated framework for modeling the multiple impacts of pricing decisions in a B2B context and illustrate how our framework can aid sellers in implementing first-degree and intertemporal price discrimination for long-run profitability.

Pricing decisions in B2B contexts differ from those within business-to-consumer (B2C) environments across multiple dimensions. First, B2B settings are often characterized by product and service customization and by the reliance on personal selling to forge

and cement transactions. In many B2B situations, sellers can easily vary prices across buyers and can even change prices between subsequent purchases of the same buyer. In contrast, B2C retailers are often limited in their ability to fully target prices for individual consumers because of logistical and ethical concerns (Khan et al. 2009).

Second, B2B environments are generally characterized by long-term relationships between buyers and sellers (Morgan and Hunt 1994). The development of trust, commitment, and norms over repeated interactions can impact buyers’ attitudes, comfort levels, and price sensitivities over time (Dwyer et al. 1987, Morgan and Hunt 1994, Rangan et al. 1992). Pricing decisions, in turn, can play a vital role in developing, transforming, and sustaining such relationships (Kalwani and Narayandas 1995).

Third, transactions in B2B markets are more complex than those in B2C markets as business buyers typically make several interrelated decisions on a given purchase occasion. Specifically, B2B buyers not only choose what, when, and how much to buy but also decide how to buy. In B2B settings, buyers choose whether to ask for a price quote (offering the seller the opportunity to provide a price quote) or to order directly from the seller, without asking

for a price. Requests for price quotes allow sellers to observe demand and price sensitivity even when a sale is not made (i.e., when the seller makes a bid and the buyer rejects the bid). Such data are rarely observed in B2C settings (Khan et al. 2009).

Fourth, situational triggers can influence the decisions of buyers. For instance, price changes in commodity markets can impact purchasing decisions, thus necessitating the use of such external factors in modeling demand. Finally, decision makers (buyers and sellers) in B2B settings are often assumed to behave rationally (Reid and Plank 2004). Thus, one needs to consider whether behavioral pricing effects regarding reference prices and loss aversion (Kalyanaram and Winer 1995) are operant in the B2B domain (Bruno et al. 2012).

In summary, the above distinguishing features of B2B settings offer sellers significant pricing flexibility. In particular, the reliance on salespeople, the need for product/service customization, and the volatility of commodity prices make targeting and intertemporal customization of prices feasible and desirable. The adoption of sophisticated customer relationship management software and database capabilities is making such customization increasingly possible.

In this paper, we develop a modeling framework that incorporates the unique facets of B2B contexts and models the multiple buyer decisions on each purchasing opportunity in an integrated fashion. Specifically, we posit that the different aspects of buyer behavior are governed by a common latent state that represents the trust between the buyer and the seller. This latent state creates dependencies across the buyer's decisions (when to buy, how much to buy, whether to request a quote or order directly without a quote, and whether to accept or reject the quote). The level of trust can evolve over time as a function of the nature of interactions between the buyer and the seller and via the seller's pricing decisions. We use a multivariate nonhomogeneous hidden Markov model (HMM) to model how trust governs buyer decisions and how it evolves over the duration of the relationship as a function of pricing. In addition, our HMM framework accounts for buyer heterogeneity, and it incorporates internal and external (commodity) reference price effects and price endogeneity using a Bayesian version of the control function approach (Park and Gupta 2009, Petrin and Train 2010).

We apply our framework on longitudinal transaction data from an aluminum retailer that sells to industrial buyers. We identify two latent states of trust that are consistent with conceptual frameworks of buyer segmentation in the B2B literature (e.g., Rangan et al. 1992, Shapiro et al. 1987). These include (1) a *vigilant state* characterized by high buyer price sensitivity and a cautious approach toward ordering

and (2) a *relaxed state* that is characterized by more direct orders and lower price sensitivity. We also find strong evidence for asymmetric reference price effects such as loss aversion and gain seeking. Consistent with relationship life-cycle theory (Dwyer et al. 1987, Jap and Anderson 2007) and hedonic adaptation theory (Frederick and Loewenstein 1999), we find that buyers not only weigh price losses more than gains but also take longer to adapt to losses than to gains. We further provide empirical evidence for the rate of buyer–seller relationship migration over time and show how the seller can use prices to manage these relationship migrations profitably.

From a managerial perspective, we illustrate how the seller can use our model to compute optimal prices that are targeted for each transaction of each buyer so as to maximize long-term profits. Such an optimal pricing policy balances different short- and long-term perspectives. Although it is common to think of prices as having mainly short-term effects, prices are likely to have long-term effects in B2B markets because of the importance of buyer–seller relationships and because prices can impact trust. We find that the optimal dynamic targeted pricing policy can increase the seller's profitability by as much as 52% over that of the status quo. We also use a counterfactual analysis to examine the nature of the optimal pricing policy in the presence of a volatile aluminum commodity market. Changing commodity prices alter the seller's costs and the external reference prices of buyers. We find results that are consistent with the dual-entitlement principle (Kahneman et al. 1986)—it is optimal for the seller to pass on much of the cost increase to buyers when commodity prices increase, whereas it is optimal to “hoard” some of the benefits of a cost decrease when commodity prices drop.

In summary, our research advances the B2B pricing literature in several directions. On the methodological front, it offers a hierarchical Bayesian framework for relationship dynamics that weaves together a multivariate nonhomogeneous HMM, heterogeneity, and control functions. More important, on the substantive front, it offers B2B managers an approach to dynamically target prices. Our results showcase the effect of pricing decisions on the evolution of trust between buyers and sellers and illustrate how behavioral factors such as loss aversion and reference price, which are commonly ignored in what are traditionally considered to be “rational” purchasing contexts and in the commonly used cost-plus pricing approach are important for B2B pricing. We also offer insights about how the seller should react to volatile commodity prices.

The rest of the paper is organized as follows. Section 2 highlights the challenges and opportunities in

investigating pricing decisions in B2B settings. Section 3 describes the data from an industrial aluminum retailer. Section 4 outlines our modeling framework. Section 5 applies our modeling framework to the data. Section 6 describes the dynamic targeted pricing optimization based on the estimated model, and §7 concludes by discussing practical implications, theoretical contributions, and future directions.

2. Targeted Pricing Decisions in B2B Settings

The majority of the research on B2B pricing is conceptual and survey based (Johnston and Lewin 1996). Scant attention is given to quantitative pricing models (for an exception, see Bruno et al. 2012), perhaps because of conflicting views about the role and importance of prices relative to other attributes in B2B contexts (see Hinterhuber 2004, Lehmann and O'Shaughnessy 1974, Wilson 1994). B2B researchers, however, have intensively investigated the role of buyer–seller relationships in B2B markets and have offered various segmentation and targeting frameworks. We now review this literature and briefly discuss past research on reference prices.

2.1. Relationships in B2B Markets

Buyer–seller relationships can be described using a number of relational constructs such as trust, commitment, and norms (Dwyer et al. 1987). Morgan and Hunt (1994, p. 23) posit that trust, “the confidence in an exchange partner’s reliability and integrity,” and commitment, “an enduring desire to maintain a valued relationship” (Moorman et al. 1992, p. 316), are key elements that explain the quality of relationships and their impact on behaviors and performance. Palmatier et al. (2006) suggest that a composite construct called “relationship quality”—an amalgam of trust, commitment, and satisfaction—has a strong impact on objective performance. Similarly, Dwyer et al. (1987) suggest that relational variables such as trust, commitment, norms, and in general relationship quality increase as relationships progress to more positive states.

Dahl et al. (2005) show that activities that embody fairness enhance trust, whereas acts of unfairness, opportunism, and conflict negatively influence trust and commitment toward the seller (Anderson and Weitz 1992). Dwyer et al. (1987) and Jap and Anderson (2007) posit that negative actions, especially those that are perceived to be unfair, can cause the rapid deterioration of a relationship, with a low prospect of a rebound. Apart from seller actions, environmental uncertainty can also moderate relationship performance (Cannon and Perreault 1999). We rely on this research to model the impact of pricing decisions and the influence of uncertain commodity markets on the evolution of buyer–seller relationships.

2.2. Segmentation and Targeting in B2B Markets

Firmographics, such as customer size, industry, and customer location, are traditionally used for segmentation in industrial markets. Researchers, however, have also proposed segmentation based on buying behavior and relationship with sellers. Rangan et al. (1992) suggest that the weight given to price (relative to service) is an important driver of buyer heterogeneity. They use survey data to identify a segment of “programmed” business buyers who are less price sensitive and invest less in the buying process and a segment of “transactional” buyers who are more sensitive to price and are also more knowledgeable about the product because it is more important to their businesses. Shapiro et al. (1987) refer to these two segments as “passive” and “aggressive” buyers, and Matthyssens and Van den Bulte (1994) denote these as “co-operative” and “antagonistic.” Shapiro et al. discuss the merit of using different targeting strategies for each segment and the possible migration of buyers between these segments as a result of the seller’s targeting efforts.

In this research, we use transactional data to uncover evidence that supports the above dynamic segmentation framework. Consistent with the papers discussed above, we find that buyers at any given time can belong to either a *relaxed* or *vigilant* state, depending on their sensitivity to past prices, previous transactional outcomes, and sensitivity to market conditions. We then propose a targeting framework that leverages this information to migrate buyers between these two states.

The growing literature on targeting and customization is also relevant for our research. The empirical literature on targeting has focused mostly on *non-price* instruments. In B2B settings, marketing actions such as face-to-face meetings, direct mail and telephone contacts (Venkatesan and Kumar 2004), and dollar expenditure on marketing efforts (Kumar et al. 2011) have been investigated. Similarly, in B2C contexts, researchers have focused on marketing actions such as catalog mailing (Simester et al. 2006), coupons (Rossi et al. 1996), digital marketing campaigns (Ansari and Mela 2003), pharmaceutical detailing and sampling (Dong et al. 2009, Montoya et al. 2010), and promotions (Khan et al. 2009). Empirical research on *individually* targeted pricing has been relatively sparse, possibly because of the informational, logistical, ethical, and legal constraints that impact price discrimination in traditional (B2C) settings (Khan et al. 2009).

2.3. Reference Prices in Customer Buying Behavior

The notion that consumers rely on internal and external reference prices is well established within

marketing (Hardie et al. 1993, Kalwani et al. 1990, Kalyanaram and Winer 1995, Winer 1986). External reference prices (e.g., manufacturer's suggested retail price and prices of other brands) are generally observable and common to all customers, whereas internal reference prices are individual specific and are often constructed using the customer's observed prices on previous purchase occasions.

A large literature demonstrates the behavioral and psychological (Kalwani et al. 1990, Wedel and Leeflang 1998) as well as the rational and economic (Erdem et al. 2010) underpinnings of reference price effects. In the behavioral pricing literature, prices can be coded as either losses or gains relative to a reference price and can thus have an asymmetric impact on brand choice (Kalwani et al. 1990, Putler 1992), purchase timing (Bell and Bucklin 1999), or purchase quantity (Krishnamurthi et al. 1992).

Despite the voluminous literature, reference prices have found little application in B2B pricing models because B2B decision makers are presumed to be rational (Kalyanaram and Winer 1995). In a recent exception, Bruno et al. (2012) demonstrate that industrial buyers exhibit asymmetric reference price effects that are affected by the depth of interactions between buyers and sellers. We add to the sparse literature on B2B reference pricing and explore both internal (past prices) and external (commodity spot prices) reference prices. We also examine the possible long-term effects of reference prices in a B2B context. Next, we describe our data set and the business context in which the seller operates.

3. Data

Our data come from an East Coast aluminum retailer (seller) that supplies to industrial clients (buyers) who operate in its geographical trading area. The seller buys raw aluminum directly from the mills, cuts it according to the specifications provided by its small- to medium-sized industrial clients (e.g., machine shops, fabricators, small manufacturers), and then ships the product to them. The buyers use the product as a component in their own products or services. Hence, our seller is a value-add intermediary in the industrial consumable market, and our data set is typical of what is found in this B2B market.

The data set contains buyer-level information on purchase events over 21 months from January 2007 to September 2008. A *purchase event* begins with the need for a certain quantity at a given point in time. Given this need, the buyer either places a direct order, without asking for a price quote, or requests a price quote (usually via phone or fax). For example, a typical direct order may be received in the morning via a fax saying, "Send me four aluminum sheets, X inch

by Y inch and thickness of Z inch, by tomorrow afternoon." Direct orders are generally fulfilled immediately, and the buyer is charged a price determined by the seller. Alternatively, if the buyer requests a quote (i.e., places an "indirect order"), the firm bids for the buyer's business and can only "win" the business if the buyer accepts the quoted price. Thus, in our setting, purchase events include not only completed transactions but also lost transactions involving quotes that were not accepted. This allows for a better understanding of buyer price sensitivity.

The seller keeps a large number of stock-keeping units (SKUs) that are defined based on the shape, thickness, and customizable size of the aluminum. Furthermore, the wholesale cost of aluminum changes on a daily basis following the London Metal Exchange (LME). Therefore, as is typical in this industry, the seller does not maintain a price list and determines the price to charge or quote on a case-by-case basis. Because order quantities vary substantially and because of the large number of SKUs, the industry uses a common metric, "price per pound," to which the seller adds the cutting and delivery costs to arrive at a price for an order. As is typical of most customer relationship management data sets in B2B settings, our data set does not include information about the competition. However, it is likely that a buyer requests quotes from multiple vendors. Thus, unfulfilled indirect orders provide an indirect signal for a purchase that goes to the competition. Our buying process includes a first-level-auction "take it or leave it" quote process, in which the buyer requests a quote, the seller makes a bid, and the buyer decides whether to accept the bid. Conversations with management indicate that negotiation beyond the initial quote request, as well as customer returns, are rare in this business. Moreover, in our data, we observe that the price and quantity quoted by the seller are identical to the price and quantity reported on the final invoice for more than 99% of the orders, suggesting minimal negotiations. Our model needs to be extended to explicitly capture negotiation processes when applying it to other B2B domains in which negotiations are common (e.g., Milgrom and Weber 1982, Mithas and Jones 2007).

Our sample contains 1,859 buyers for whom we observe at least seven purchase events (quotes or orders) in the data period (see Tables 1 and 2 for

Table 1 Descriptive Statistics

Number of buyers	1,859
Overall number of observations (purchase events)	33,925
Proportion of direct purchases	0.53
Proportion of quotes that are accepted	0.47

Table 2 Descriptive Statistics per Buyer

	Mean	Std. dev.	Lower 10%	Median	Upper 90%
Total number of purchase events	23.6	20.8	8.0	16.0	52.0
Proportion of direct orders	55.6	24.6	21.4	57.1	87.5
Order amount for direct orders (US\$)	861	1,636	100	402	1,932
Order amount for quote requests (US\$)	1,724	3,445	165	667	3,770
Purchase event amount (US\$)	1,236	1,471	292	808	2,458
Quantity (lb.)	457	553	92	288	968
Interpurchase event time (weeks)	6.41	4.31	1.80	5.23	12.45

summary statistics of the data).¹ On average, a buyer in our sample engaged in 23.6 purchase events during the span of 21 months. Of these, 53% were direct orders on which no price quote was requested. The relatively large proportion of direct orders is consistent with the notion of “programmed” buyers described in Rangan et al. (1992) and Shipley and Jobber (2001). It is also consistent with the view that many buyers may have developed certain levels of trust and norms with the seller and would rely on these norms for speedy order fulfillment.

An average purchase event involves 457 lb. of aluminum, with an average price of \$3.24/lb. Table 2 shows that (1) direct orders tend to be smaller, suggesting the possibility that buyers are less price sensitive when ordering smaller quantities; (2) buyers are heterogeneous in terms of their metal needs and transactions with the firm; (3) buyers exhibit different propensities to order directly, implying variation in their attitudes and latent relationships with the firm; and (4) about half of the orders are direct, which result in a sale regardless of the price charged. This suggests that the firm may be tempted to charge “any” price on such direct orders. However, as we show later, such exploitative pricing behavior can have negative long-term consequences.

We now look at some model-free evidence to understand the relationship dynamics in our data and to motivate our modeling approach. Figure 1 is based on the group of buyers for whom we observe the complete history of interactions with the seller. The

Figure 1 Model-Free Evidence—Probability of Quote Request Over Time for New Buyers




Figure 2 Model-Free Evidence—Buying Behaviors Following a Gain or a Loss on a Direct Order




figure shows the probability of a quote request for the first six purchase events of these buyers. We see that, as expected, almost all buyers request a quote on their first order. However, this probability goes down for subsequent orders (i.e., buyers are more likely to order directly over time). This pattern is consistent with the view that most buyers start out with an “exploratory” or “transactional-only” attitude toward buying but then gradually build trust, commitment, and norms of interactions with the seller over repeated interactions.

Figure 2 shows an interesting pattern that captures how current pricing on a direct order impacts buyer behavior on the next purchase event. The figure shows that charging in a direct order a price that is higher than the average of the prices the buyer faced in the past² (interpreted as a loss) increases the likelihood of a quote request on the next purchase event

¹ From a substantive point of view, more than 92% of the buyers in our data have at least seven purchase events. Thus, this selection process does not have a significant impact on the representativeness of our sample. From a methodological point of view, we use this cutoff to ensure that our model is capable of capturing a rich set of relationship dynamics. To check for robustness, we also ran models with a larger sample that included buyers with three or more purchase events. The substantive results and their significance remain similar.

² As we describe later in more detail, we use the quantity-weighted average price the buyer observed in past purchase events as a reference price.

by about 50% (from 42% to 63%) and increases the probability of losing the next bid. This pattern implies that the seller needs to be careful in pricing direct orders because charging excessively above the reference price (hence violating the trust that the buyer places in the seller) can result in undesirable consequences on subsequent purchase events. We capture such considerations in our model by allowing the buyer's latent trust level to shift over time as a function of the prices received from the seller. We describe our modeling framework next.

4. Model

In this section, we model a sequence of purchase events for each buyer while taking into account the relationship states that evolve over time as a result of the seller's pricing decisions. A *purchase event* is characterized by four interrelated buyer decisions: (1) when to buy; (2) how much to buy; (3) whether to order directly without asking for a price quote, which always results in a purchase, or to request a quote, hence allowing the seller to bid for business; and (4) whether to accept the quote if a quote is requested. We can write the vector of observed behaviors for buyer i at purchase event j as $\mathbf{y}_{ij} = (q_{ij}, t_{ij}, b_{ij}, w_{ij})$, where q_{ij} is the quantity requested or ordered, t_{ij} is the time (in weeks) since the last purchase event (i.e., the interpurchase event time), and b_{ij} and w_{ij} are the binary quote request and quote acceptance decisions, respectively. The seller observes the marketing environment and buying and pricing history for buyer i at purchase event j before setting the unit price p_{ij} for the event.

To model buyer dynamics over repeated purchase events, we allow the buyer to transition between different latent behavior/relationship states of trust³ that differentially impact the four buying decisions. The seller's past pricing decisions may affect the buyer's transition between states. For example, as suggested by Figure 2, a buyer who is charged a high price may be more likely to transition from a "relaxed" or trusting state that is characterized by a high propensity to order directly and a low price sensitivity to a "vigilant" or evaluative state that reflects a higher propensity to request quotes and a higher price sensitivity.

We capture such dynamics using a multivariate nonhomogeneous HMM (Montoya et al. 2010, Netzer et al. 2008, Schweidel et al. 2011). In the HMM,

³ In what follows, we call the HMM latent states *states of trust*. However, this latent state could be interpreted more generally as a relationship quality state, which combines trust, commitment, and norms between the buyer and the seller (Palmatier et al. 2006). Because the states are inferred from secondary data, we remain agnostic about this distinction.

the joint probability of a sequence of interrelated decisions up to purchase event j , for buyer i , $\{\mathbf{Y}_{i1} = \mathbf{y}_{i1}, \dots, \mathbf{Y}_{ij} = \mathbf{y}_{ij}\}$, is a function of three main components: (1) the initial hidden state membership probabilities ψ_i ; (2) a matrix of transition probabilities among the buying-behavior states, $\Omega_{i,j-1 \rightarrow j}$; and (3) a multivariate likelihood of the buyer decisions conditional on the buyer's buying-behavior state, $L_{ij|s} = f_{is}(q_{ij}, t_{ij}, b_{ij}, w_{ij})$. We describe each of these components next.

4.1. Initial State Distribution

Let s denote a buying-behavior state ($s = 1, 2, \dots, S$). Let ψ_{is} be the probability that buyer i is in state s at time 1, where $0 \leq \psi_{is} \leq 1$ and $\sum_{s=1}^S \psi_{is} = 1$. We use $S-1$ logit-transformed parameters to represent the vector containing the initial state probabilities.

4.2. Markov Chain Transition Matrix

Ring and Van de Ven (1994) suggest that B2B relationships evolve with repeated buyer-seller interactions; the experience from each interaction can either elevate or upset a relationship. Consistent with this notion, we model the transitions between states as a Markov chain. Each element of the transition matrix ($\Omega_{i,j-1 \rightarrow j}$) can be defined as $\omega_{ijss'} = P(S_{ij} = s' | S_{ij-1} = s)$, which is the conditional probability that buyer i moves from state s at purchase event $j-1$ to state s' at purchase event j , and where $0 \leq \omega_{ijss'} \leq 1 \forall s, s'$, and $\sum_{s'} \omega_{ijss'} = 1$. Because the transition probabilities are influenced by the seller's pricing decisions at the previous purchase event $j-1$, we define

$$\omega_{ijss'} = \frac{e^{\mathbf{x}_{ij-1}' \boldsymbol{\gamma}_{is}}}{1 + \sum_{s=1}^{S-1} e^{\mathbf{x}_{ij-1}' \boldsymbol{\gamma}_{is}}}, \quad (1)$$

where $\mathbf{x}_{ij-1}$ is a vector of covariates (e.g., price or reference price) affecting the transition between states, and $\boldsymbol{\gamma}_{is}$ is a state- and buyer-specific vector of response parameters.

4.3. State-Dependent Multivariate Interrelated Decisions

The buyer makes the four interrelated decisions conditional on being in state s at purchase event j . These decisions, however, are unconditionally interrelated because they all depend on the buyer's latent state.

Given that buyer i is in a latent state $S_{ij} = s$ on purchase event j , we can factor the state-conditional discrete-continuous joint likelihood, $L_{ij|s}$, for the four interrelated behaviors as⁴

$$\begin{aligned} L_{ij|s} &= f_{is}(q_{ij}, t_{ij}, b_{ij}, w_{ij}) \\ &= f_{is}(q_{ij}, t_{ij}) \Pr_{is}(b_{ij}, w_{ij} | q_{ij}, t_{ij}). \end{aligned} \quad (2)$$

⁴ To avoid clutter, we describe first the model in the general distribution form and then outline the particular distributions and parameterizations that we used.

In the above, we assume that the joint decisions on timing and quantity stem primarily from the buyer's need for the product. Because these decisions occur prior to the decision to request a quote or order directly, they impact the latter set of decisions. The decision to accept or reject a quote (w_{ij}) occurs only when the buyer decides to request a quote rather than order directly from the seller ($b_{ij} = 1$), so we specify the joint probability of b_{ij} and w_{ij} as follows:

$$\begin{aligned} \Pr_{is}(b_{ij}, w_{ij} | q_{ij}, t_{ij}) \\ = \Pr_{is}(b_{ij} = 0 | q_{ij}, t_{ij})^{1-\delta_{ij}^b} \\ \cdot [\Pr_{is}(w_{ij} | b_{ij} = 1, q_{ij}, t_{ij})\Pr_{is}(b_{ij} = 1 | q_{ij}, t_{ij})]^{\delta_{ij}^b}, \end{aligned} \quad (3)$$

where δ_{ij}^b equals 1 if purchase event j for buyer i is a quote request and 0 otherwise.

In modeling the time between purchase events, t_{ij} , the last observation for each buyer, t_{ij}^* , is censored because of the fixed time horizon of the data set.⁵ Let $S(t_{ij}^*)$ be the survival function for the censored observation, and let δ_{ij}^c be a censoring indicator, which equals 1 if observation j for buyer i is censored and 0 otherwise. Accordingly, accounting for censoring and inserting Equation (3) into Equation (2), we can rewrite Equation (2) as follows:

$$\begin{aligned} L_{ij|s} = f_{is}(q_{ij}, t_{ij}, b_{ij}, w_{ij}) \\ = S_{is}(t_{ij}^*)^{\delta_{ij}^c} [f_{is}(q_{ij}, t_{ij})\Pr_{is}(b_{ij} = 0 | q_{ij}, t_{ij})^{1-\delta_{ij}^b} \\ \cdot [\Pr_{is}(w_{ij} | b_{ij} = 1, q_{ij}, t_{ij}) \\ \cdot \Pr_{is}(b_{ij} = 1 | q_{ij}, t_{ij})]^{\delta_{ij}^b}]^{1-\delta_{ij}^c}. \end{aligned} \quad (4)$$

Next, we describe the distributional assumptions for each of the four decisions.

4.3.1. Modeling Quantity and Time Between Events. We assume that the purchase event times follow a two-parameter log-logistic distribution (Kumar et al. 2008; Lancaster 1990, p. 44) because it flexibly accommodates both monotonic and nonmonotonic hazards. The probability density function (p.d.f.) and cumulative distribution function (c.d.f.) of the log-logistic are given by

$$\begin{aligned} f_{is}(t_{ij}) = \frac{e^{x'_{ij}\beta_{tsi} + \xi_{ij}^t} \alpha_s t_{ij}^{\alpha_s - 1}}{(1 + e^{x'_{ij}\beta_{tsi} + \xi_{ij}^t} t_{ij}^{\alpha_s})^2}, \\ F_{is}(t_{ij}) = \frac{e^{x'_{ij}\beta_{tsi} + \xi_{ij}^t} t_{ij}^{\alpha_s}}{1 + e^{x'_{ij}\beta_{tsi} + \xi_{ij}^t} t_{ij}^{\alpha_s}}, \end{aligned} \quad (5)$$

⁵ We do not model the first purchase event for each buyer and use it as an initialization period to account for left truncation and initialize the dynamics.

where $\alpha_s > 0$ is a shape parameter; β_{tsi} is a vector of coefficients for buyer-level, purchase event-specific covariates such as prices or reference prices; and ξ_{ij}^t represents an unobserved shock associated with the interpurchase event time. We assume that the random shock ξ_{ij}^t is correlated with the unobserved shock in the pricing equation to account for possible endogeneity (see §4.4).

We assume that quantities requested and/or ordered follow a log-normal distribution with p.d.f. and corresponding c.d.f. given by

$$\begin{aligned} f_{is}(q_{ij}) = \frac{\phi((\log q_{ij} - x'_{ij}\beta_{qsi} - \xi_{ij}^q)/\sigma)}{\sigma q_{ij}}, \\ F_{is}(q_{ij}) = \Phi\left(\frac{\log(q_{ij}) - x'_{ij}\beta_{qsi} - \xi_{ij}^q}{\sigma}\right), \end{aligned} \quad (6)$$

where β_{qsi} is a vector of coefficients for a set of buyer-level and purchase event-specific covariates that affect the mean quantity, ξ_{ij}^q is an unobserved random shock that is correlated with the unobserved shock in the pricing equation discussed below, σ is the scale parameter, and ϕ and Φ represent the p.d.f. and c.d.f. of the standard normal distribution, respectively.

4.3.2. Modeling Buyer Quote Request and Acceptance Decisions. Buyer i 's binary quote request decision on purchase event j , (b_{ij}) is governed by an underlying latent utility b_{ij}^* such that

$$b_{ij} = \begin{cases} 1 & \text{if } b_{ij}^* > 0 \text{ (indirect),} \\ 0 & \text{otherwise (direct).} \end{cases}$$

Similarly, conditional on a price quote, buyer i 's binary decision to accept or reject the quote on purchase event j , (w_{ij}), is driven by the latent utility w_{ij}^* such that

$$w_{ij} = \begin{cases} 1 & \text{if } b_{ij}^* > 0 \text{ and } w_{ij}^* > 0, \\ 0 & \text{if } b_{ij}^* > 0 \text{ and } w_{ij}^* \leq 0, \\ \text{unobserved} & \text{otherwise.} \end{cases}$$

We assume that each of the latent variables, b_{ij}^* and w_{ij}^* , are distributed logistic. Thus,

$$\begin{aligned} \Pr_{is}(b_{ij}^* < 0) = \frac{1}{1 + e^{x'_{ij}\beta_{bsi} + \xi_{ij}^b}} \quad \text{and} \\ \Pr_{is}(w_{ij}^* < 0) = \frac{1}{1 + e^{x'_{ij}\beta_{wsi} + \xi_{ij}^w}}. \end{aligned} \quad (7)$$

The vector of parameters β_{bsi} and β_{wsi} relate the quote request and quote acceptance latent utilities, respectively, to a set of covariates such as price, reference price, and time since the last order. The unobserved shocks ξ_{ij}^b and ξ_{ij}^w are associated with the quote request and quote acceptance decisions, respectively. These are correlated with the unobserved shock of the pricing equation, which is discussed subsequently.

4.4. The Control Function Approach to Price Endogeneity

We need to account for two potential sources of endogeneity. First, it is possible that the seller's pricing decisions are based on unobserved random shocks that also impact the buyers' decisions. For example, demand boosts and supply shortages can increase the prices that sellers charge. Such common economic shocks to both pricing and demand may be observed by buyers and the seller but remain unobserved to the researcher. In such a case, price will be correlated with the unobserved components (the ξ 's) of the four distributions. Second, the seller may set prices for each buyer individually by using its knowledge about each buyer's sensitivity. This again is private information that is not observed by the researcher. Ignoring endogeneity can result in misleading inferences about the price sensitivities of customers (Villas-Boas and Winer 1999). We therefore use a Bayesian analog of the control function approach to account for price endogeneity (Park and Gupta 2009, Rossi et al. 2005).

We express price as a function of an observed instrumental variable, z_{ij} , that is correlated with price but is uncorrelated with the unobserved factors that impact the four decisions. Specifically, we use the seller's wholesale cost (that is, the cost that the seller pays to the mills for the metal) as the instrumental variable z_{ij} to address the first source of potential endogeneity resulting from common unobserved economic shocks. This cost is observed by the seller but not by buyers. Conversations with the management team reveal that the salespeople observe the wholesale cost on their computer screens and rely heavily on it when setting the price. Wholesale prices have been commonly used as instruments for price (e.g., Chintagunta 2002). To address the second potential source of endogeneity, individual targeting, we use a buyer-specific random intercept in the pricing equation below. Formally stated, we have

$$p_{ij} = \lambda_{1i} + \lambda_2 z_{ij} + \mu_{ij},$$

where μ_{ij} represents unobserved factors that influence the pricing decision. We assume that μ_{ij} is distributed jointly bivariate normal with each of ξ_{ij}^l , $l \in \{t, q, b, w\}$ in Equations (5)–(7).

The bivariate normal distribution for each of the four decisions can be written as

$$f(\mu_{ij}, \xi_{ij}^l) \sim \text{MVN} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_p^2 & \rho_{pl} \\ \rho_{pl} & \sigma_l^2 \end{pmatrix} \right) \\ l \in \{t, q, b, w\}, \quad (8)$$

where σ_p^2 is the variance of μ_{ij} , σ_l^2 is the variance for the random shock ξ_{ij}^l , and ρ_{pl} is the covariance between μ_{ij} and ξ_{ij}^l .

Inserting Equations (5)–(8) into Equation (2), we obtain the likelihood of the four interrelated buyer decisions and the observed price, conditional on the buyer's state and the random shocks μ_{ij} and ξ_{ij}^l :

$$L_{ij|s} = f_{is}(q_{ij}, t_{ij}, b_{ij}, w_{ij}, p_{ij}) \\ = f_{is}(q_{ij}, t_{ij} | p_{ij}) \text{Pr}_{is}(b_{ij}, w_{ij} | q_{ij}, t_{ij}, p_{ij}) f(p_{ij}).$$

4.5. The HMM Likelihood Function

The likelihood of observing the buyer's decisions over J purchase events ($\mathbf{Y}_{i1}, \mathbf{Y}_{i2}, \dots, \mathbf{Y}_{iJ}$) can be succinctly written as (MacDonald and Zucchini 1997)

$$L_{ij} = P(\mathbf{Y}_{i1} = \mathbf{y}_{i1}, \dots, \mathbf{Y}_{iJ} = \mathbf{y}_{iJ}) \\ = \Psi_i \mathbf{M}_{i1} \mathbf{\Omega}_{i,1 \rightarrow 2} \mathbf{M}_{i2}, \dots, \mathbf{\Omega}_{i,J-1 \rightarrow J} \mathbf{M}_{iJ} \mathbf{1}', \quad (9)$$

where Ψ_i is the initial state distribution described in §4.1, $\mathbf{\Omega}_{i,j-1 \rightarrow j}$ is the transition matrix described in §4.2, $\mathbf{M}_{ij}$ is a $S \times S$ diagonal matrix with the elements $L_{ij|s}$ from Equation (4) on the diagonal, and $\mathbf{1}'$ is a $S \times 1$ vector of ones.

We restrict the probability of a quote request to be nondecreasing in the trust states to ensure the identification of the states. We impose the restriction $\tilde{\beta}_{b01i} \leq \tilde{\beta}_{b02i} \leq \dots \leq \tilde{\beta}_{b0Si}$ by setting $\tilde{\beta}_{b0Si} = \beta_{b01i} + \sum_{s'=2}^S \exp(\beta_{0s'i})$, $s = 2, \dots, S$. As both the intercepts and the response parameters are state specific, we impose this restriction at the mean of the vector of covariates by mean centering $\mathbf{x}_{ij}$. Finally, we scale the likelihood function in Equation (9) following the approach suggested by MacDonald and Zucchini (1997, p. 79) to avoid underflow.

4.6. Recovering the State Membership Distribution

We use filtering (Hamilton 1989) to determine the probability that buyer i is in state s at purchase event j conditioned on the buyer's history:

$$P(S_{ij} = s | \mathbf{Y}_{i1}, \mathbf{Y}_{i2}, \dots, \mathbf{Y}_{ij}) \\ = \Psi_i \mathbf{M}_{i1} \mathbf{\Omega}_{i,1 \rightarrow 2} \mathbf{M}_{i2}, \dots, \mathbf{\Omega}_{i,j-1 \rightarrow j} L_{ij|s} / L_{ij}, \quad (10)$$

where $\mathbf{\Omega}_{i,j-1 \rightarrow j}$ is the s th column of the transition matrix $\mathbf{\Omega}_{i,j-1 \rightarrow j}$, and L_{ij} is the likelihood of the observed sequence of joint decisions up to purchase event j from Equation (9).

5. Model Estimation and Results

In this section, we describe how we instantiate the above model in our application. We first present the rationale for our choice of variables and then interpret the parameter estimates.

5.1. Description of Variables

Internal reference prices and asymmetric reference price effects: We define the internal reference price for buyer i , at purchase event j , as a quantity-weighted average of the buyer's past observed prices (in dollars per pound)⁶

$$\text{reference_price}_{ij} = \frac{\sum_{k=1}^{j-1} \text{quantity}_{ik} \times \text{price_lb}_{ik}}{\sum_{k=1}^{j-1} \text{quantity}_{ik}},$$

where price_lb_{ij} is the price per pound observed by buyer i in purchase event j . The quantity weighting reflects that buyers attend closely to larger orders relative to smaller ones. To examine the differential effects of price increases and price decreases on buyer decisions, we incorporate asymmetric reference price effects using “gain” and “loss” variables:

$$\text{gain}_{ij} = \begin{cases} \text{reference_price}_{ij} - \text{price_lb}_{ij} & \text{if } \text{price_lb}_{ij} < \text{reference_price}_{ij}, \\ 0 & \text{otherwise;} \end{cases}$$

$$\text{loss}_{ij} = \begin{cases} \text{price_lb}_{ij} - \text{reference_price}_{ij} & \text{if } \text{price_lb}_{ij} > \text{reference_price}_{ij}, \\ 0 & \text{otherwise.} \end{cases}$$

External reference prices and the commodity market: Customers often use external reference prices when making buying decisions (Kopalle and Lindsey-Mullikin 2003, Mazumdar et al. 2005). Commodity prices serve as an obvious candidate for an external source of reference price in our setting. We expect that some industrial buyers would attend to this external source of reference price when making buying decisions. As the seller primarily sells aluminum products, we use the daily spot prices from the LME aluminum spot market to capture the impact of the fluctuations in the commodity market. We define lme_{ij} as the aluminum spot price (in thousand dollars per metric ton) on the LME at purchase event j for buyer i .

We theorize that, all else being equal, high LME prices at the time of purchase will make the seller's offered price appear relatively lower. This effect is likely to be stronger for buyers who have a low level of trust for the seller, leading them to consult the LME prior to purchasing. Additionally, high LME prices may result from an overall good market for aluminum and thus reflect increased demand.

⁶ We tested several alternative specifications of the reference price variable including simple average of past prices and time-weighted reference prices. The quantity-weighted reference price resulted in the best model fit. Incorporating time decay to the quantity-weighted reference price formulation did not result in significant improvement in the model's fit.

Economic volatility and the volatility in the commodity market: The extant literature in B2B relationship marketing suggests that buyers rely more heavily on relationships with sellers during periods of environmental uncertainty and volatility because strong relationships provide stability in an unstable environment (Fang et al. 2011, Jaworski and Kohli 1993, Palmatier 2008). To assess the impact of economic volatility, we define $\text{lme_volatility}_{ij}$ as the volatility of the aluminum spot prices, calculated as the standard deviation of the LME daily returns over the seven trading days prior to purchase event j for buyer i .

We theorize that the volatility in the spot market would have a negative effect on demand. This adverse effect of economic volatility will be attenuated for buyers who have a higher-quality relationship with the seller relative to those with lower levels of trust. We now describe how these variables are included in each of the buyer decision's components.

5.1.1. The State-Dependent Decisions. We include the following variables in the state-dependent components for the four decisions:

1. *Purchase event times:* We expect the timing of the purchase event to depend on the previous quantity because of inventory effects and on past internal reference price effects. Thus, $\mathbf{x}_{ij}$ in Equation (5) includes the covariates gain_{ij-1} , loss_{ij-1} , and quantity_{ij-1} .

2. *Quantity:* We expect that the price gain (loss) experienced on the previous purchase event and the current level and volatility of the commodity market to impact requested quantity. Thus, $\mathbf{x}_{ij}$ in Equation (6) includes the covariates gain_{ij-1} , loss_{ij-1} , lme_{ij} , and $\text{lme_volatility}_{ij}$.

3. *Quote request:* Given a particular relationship state, we expect that buyers in general will have a lower propensity to order directly when the quantity desired is large, when a long time has elapsed since the previous purchase, or when the market conditions are volatile. Furthermore, consistent with Figure 2, a perceived overcharge on the previous purchase event could increase the likelihood of requesting a quote. Thus, $\mathbf{x}_{ij}$ in Equation (7) includes t_{ij} , quantity_{ij} , gain_{ij-1} , loss_{ij-1} , lme_{ij} , and $\text{lme_volatility}_{ij}$.

4. *Quote acceptance:* We predict that the likelihood of accepting a quote will be higher when the quantity ordered is small, purchases are frequent, and the buyer experiences a price gain on the current purchase event. We also expect the gain and loss effects to be magnified for larger orders. Furthermore, the decision could be affected by the commodity market conditions. Thus, $\mathbf{x}_{ij}$ in Equation (7) includes t_{ij} , quantity_{ij} , gain_{ij} , loss_{ij} , $\text{gain}_{ij} \times \text{quantity}_{ij}$, $\text{loss}_{ij} \times \text{quantity}_{ij}$, lme_{ij} , and $\text{lme_volatility}_{ij}$.

5.1.2. The Nonhomogeneous Transition Matrix.

The gain or loss experienced on the previous purchase event can affect the buyer's evaluation (or reevaluation) of the relationship with the seller and can trigger a transition among relationship states, thereby affecting purchases in the long run. Specifically, we postulate that price losses (a perceived overcharge) may trigger a transition to a lower trust state. Similarly, a price gain on the previous order may trigger a transition to a higher trust state. Thus, x_{ij-1} in Equation (1) includes $gain_{ij-1}$ and $loss_{ij-1}$.

5.2. Heterogeneity Specification

Capturing cross-buyer heterogeneity facilitates targeting and allows for the proper accounting of reference price effects (Bell and Lattin 1998). Capturing heterogeneity is also crucial for empirically distinguishing dynamics from cross-buyer heterogeneity (Heckman 1981). Therefore, we allow the initial state membership, the transition matrix parameters, the coefficients in the four equations, and the intercept of the pricing control function equation to vary across buyers.

5.3. Estimation Procedure

We use a hierarchical Bayesian approach based on Markov chain Monte Carlo (MCMC) methods for inference. The inherent complexity of the HMM often leads to significant autocorrelation among the draws. We therefore use the adaptive Metropolis procedure in Atchadé (2006) to improve mixing and convergence. We use proper but diffuse priors for all parameters. Details of priors as well as full conditionals are available from the authors upon request. We use the first 18 months of data for estimation and the last three months for validation purposes. Our results are based on the last 250,000 draws from an overall MCMC run of a million iterations. Convergence was ensured by monitoring the time series of the MCMC draws.

5.4. Choosing the Number of States

We begin by determining the number of HMM states. We use the in-sample log-marginal density (LMD) and the deviance information criterion (DIC) to select the number of states. The models with one, two, three, and four states have LMD values of $-70,652$, $-62,781$, $-63,352$, and $-64,208$, respectively, and DIC values of $144,725$, $130,516$, $133,304$, and $135,921$, respectively. Thus, the model with two states has the highest support in terms of both LMD and DIC. We therefore move forward with this model.⁷

⁷ We also estimated a three-state HMM where the third state is an absorbing state with no purchase activity, capturing permanent defection. The LMD and DIC of that model were $-63,764$ and $133,421$, respectively. For completeness, we also tested the fit of a two-segment latent class model (LMD: $-94,462$, DIC: $187,234$) and

5.5. Model Fit and Predictive Ability

We compare the fit and predictive ability of the proposed model (*full model*) to that of six benchmark models. These differ from the full model with respect to (1) the extent of heterogeneity, (2) the degree of dynamics as captured by the HMM, (3) the accounting for price endogeneity, and (4) whether the buying state is modeled as observed or latent. We also compare our model to a recency-frequency-monetary (RFM) model (Benchmark 5) that is typically used in marketing to assess and predict customer relationships. Finally, we compare the full model to a latent class model in which the group membership is allowed to vary over time based on previous pricing (Allenby et al. 1999). The benchmark models are as follows:

1. *Benchmark 1*: In this model, all parameters are assumed to be invariant across buyers. A comparison of this model with the full model allows us to assess the importance of modeling heterogeneity.

2. *Benchmark 2*: This model ignores the two sources of dynamics present in the full model—i.e., the HMM specification and the reference price effects. We therefore estimate a single state (i.e., no HMM) model in which the reference prices are replaced with actual prices. In this model, the four decisions are independent. Comparing this “static” model to the proposed model allows us to assess the value of capturing relationship dynamics. Comparing this model with the one-state model from §5.4 allows us to assess the value of accounting for reference prices.

3. *Benchmark 3*: This model assumes that the prices are exogenous; otherwise, it is identical to the full model in all other aspects. Thus, this model does not have the Bayesian control function component. A comparison of this model with the full model can highlight the extent of price endogeneity and the perils of ignoring it.

4. *Benchmark 4*: This model uses an observed state variable instead of a latent relationship state. Among the four buying decisions, the quote request behavior is the most indicative of the relationship state. Hence, in this “simplified” observed state model, we deterministically assign state membership based on each buyer's quote request behavior on the previous purchase event instead of using the probabilistic latent relationship in an HMM.

5. *Benchmark 5*: This is the RFM model commonly used in B2C settings, where we model each of the

three-segment latent class model (LMD: $-89,321$, DIC: $179,442$) by restricting the transition matrix to a diagonal matrix. Thus, constraining the parameters of the transition matrix resulted in worse measures of performance.

four decisions using recency (interpurchase time), frequency (the buyer's historical average number of purchases per week), monetary value (the buyer's historical average of invoice prices), and the other covariates that are used in the full model.

6. Benchmark 6: This is a latent class model that uses time-varying class membership weights to capture dynamics. Instead of using an HMM, this model captures dynamics by specifying the latent state membership as a function of only the reference prices and not of the previous states.

We compare the fit and predictive ability of the seven models using the LMD and the DIC statistics on the calibration sample and the validation log-likelihood on the validation sample. We also assess the component-specific fit and predictive ability using the root mean squared error (RMSE) and the mean absolute deviation (MAD) between the predicted and observed values of the four outcome variables. In addition, we use hit rates for the binary quote request and conditional quote acceptance decisions within the calibration and validation samples. To account for uncertainty in MCMC, the prediction measures are averaged across MCMC runs, rather than using point estimate. We also assessed the performance of the model using posterior predictive checks (details of this analysis is available in the Web appendix available as supplemental material at <http://dx.doi.org/10.1287/mksc.2013.0842>).

Table 3 presents the model comparison statistics for the seven models. We see that the full model outperforms the benchmark models on both the component-specific and overall measures both in and out of sample. First, the relatively poor performance of the no-heterogeneity model results from the substantial buyer heterogeneity and suggests an opportunity for individually targeted pricing. Second, the results point to significant and latent relationship dynamics. Accounting for such dynamics in a holistic, latent fashion using an HMM, instead of using an observed state, improves the representation and prediction of buying behavior. We also find that accounting for price endogeneity results in only a marginal improvement in model performance. This is consistent with the reported use of a “cost-plus” pricing strategy by the seller and with the finding that a regression of price on wholesale cost yields an R^2 of 0.84. Although the RFM model has a reasonable fit and predictive ability in our application, it fits and predicts the data significantly worse than the proposed HMM. Furthermore, unlike the HMM, the RFM model cannot inform us about the impact of pricing decisions on the dynamics in buying behavior.

From a managerial perspective, the seller can relatively easily implement the Benchmark 4 model

based on the observed states. Although the performance of this benchmark is inferior to the full model, accounting for reference price effects and observed state changes based on bidding behavior is already a major step forward for most industrial consumable B2B companies (our company included) that still engage in cost-plus pricing and do not have any systematic way to perform individual price targeting. We explore this model further in §6.

5.6. Parameter Estimates

In this section we discuss the parameter estimates of the full model with two latent relationship states.

5.6.1. The HMM States. Table 4 contains the posterior means, standard deviations, and the 95% posterior intervals for the parameters of the full model. Recall that all covariates are mean centered, so the intercept of each equation captures the average response tendency. A comparison of the parameters across the two states indicates that buyers in State 2 are more likely to request a price quote but are less likely to accept the quote relative to buyers in State 1. Buyers in State 2 are more sensitive to reference price effects and exhibit stronger loss aversion in the interpurchase event time, quote request, and quote acceptance decisions. As we have theorized, buyers who are in State 2 at a given purchase event also tend to be more responsive to the commodity market (LME) as an external reference price. Interestingly, whereas economic volatility (as measured by the volatility of LME) has a negative impact for buyers in State 2, the impact is mitigated for buyers in State 1.

Overall, this multidimensional view of the two relationship states (see a summary of the two states in Table 5) implies that buyers in State 2 exhibit a more cautious approach toward buying from the seller, whereas buyers in State 1 appear more relaxed in their relationship with the seller. We therefore call State 1 the *relaxed state* and State 2 the *vigilant state*. As buyers transition to the relaxed state, possibly as a result of favorable past interactions with the seller, they simplify their buying decisions, become less focused on price losses, and are less concerned about the external economic environment. This simplified buying process is beneficial for both parties. For buyers, it saves resources on search and transaction costs. For the seller, buyers that are in a trusting state are a source of stable cash flow without the uncertainty of the quote request process. The relaxed state, therefore, represents a higher level of relationship quality when compared with the vigilant state. The average order quantity for customers in the relaxed state is smaller than in the vigilant state, which indicates that customers are more likely to be vigilant when the stakes are high.

Table 3 Model Selection and Predictive Validity

Heterogeneity Dynamics	Full model		Benchmark 1		Benchmark 2		Benchmark 3		Benchmark 4		Benchmark 5		Benchmark 6	
	In sample	Out of sample	Yes Dynamics (HMM + Reference price)	No Dynamics (HMM + Reference price)	Yes No dynamics	Yes No dynamics	Yes Dynamics (HMM + Reference price)	No Dynamics (HMM + Reference price)	Yes Observed state (last bid behavior) + Reference price	No Observed state (last bid behavior) + Reference price	Yes Dynamics (Reference price + RFM)	No Dynamics (Reference price + RFM)	Yes Reference price + Time-varying state membership	No Reference price + Time-varying state membership
Control function			Yes	Yes	Yes	Yes	No	No	No	No	Yes	Yes	No	No
LMD			-62,781	-90,584	-72,461	-72,461	-63,008	-63,008	-68,359	-68,359	-66,998	-66,998	-69,122	-69,122
DIC			130,516	181,623	148,811	148,811	131,197	131,197	140,388	140,388	137,351	137,351	141,025	141,025
Complexity pD			4,798	75	3,266	3,266	4,782	4,782	2,856	2,856	3,755	3,755	2,973	2,973
Validation LL			-16,025	-21,031	-16,856	-16,856	-16,073	-16,073	-17,985	-17,985	-17,045	-17,045	-17,320	-17,320
Hit rate														
Quote request	0.72	0.62	0.60	0.58	0.62	0.60	0.72	0.62	0.64	0.62	0.64	0.62	0.64	0.62
Bid acceptance	0.69	0.57	0.57	0.55	0.58	0.58	0.69	0.57	0.66	0.57	0.67	0.57	0.66	0.57
RMSE														
Quantity	0.91	0.99	1.66	1.67	1.05	1.04	0.91	0.99	1.04	1.03	1.03	1.03	1.01	1.01
Interpurchase time	15.57	19.23	24.31	27.50	20.33	21.39	15.60	19.27	20.13	21.18	19.93	20.97	19.53	20.55
Quote request	0.51	0.57	0.60	0.87	0.58	0.65	0.51	0.57	0.57	0.65	0.56	0.64	0.55	0.63
Bid acceptance	0.53	0.64	0.74	0.97	0.70	0.69	0.53	0.64	0.69	0.68	0.69	0.68	0.67	0.66
MAD														
Quantity	0.40	0.40	0.69	0.67	0.44	0.48	0.40	0.40	0.44	0.48	0.43	0.47	0.42	0.46
Interpurchase time	11.66	14.26	17.75	18.10	14.28	15.04	11.68	14.28	14.14	14.89	14.00	14.75	13.72	14.45
Quote request	0.30	0.36	0.43	0.53	0.31	0.38	0.30	0.36	0.32	0.37	0.33	0.38	0.32	0.37
Bid acceptance	0.39	0.47	0.60	0.73	0.39	0.46	0.39	0.47	0.39	0.46	0.41	0.49	0.40	0.48

Notes. Numbers in bold represent the best fit/predictive ability from among all the models. Complexity pD, the effective number of parameters of a Bayesian model; validation LL, log-likelihood in the validation sample.

Table 4 Parameter Estimates for the Four Decisions—Posterior Means, Standard Deviations, and 95% Posterior Intervals

State	Parameter	Mean	Std. dev.	Quantile		
				2.5%	97.5%	
(a): Quantity decision						
State 1	<i>Intercept</i>	−1.834	0.040	−1.912	−1.756	
	<i>gain(j − 1)</i>	0.004	0.021	−0.038	0.046	
	<i>loss(j − 1)</i>	−0.032	0.013	−0.058	−0.006	
	<i>lme(j)</i>	0.027	0.132	−0.231	0.285	
	<i>lme_volatility(j)</i>	−0.794	0.496	−1.766	0.178	
State 2	<i>Intercept</i>	−1.492	0.034	−1.558	−1.426	
	<i>gain(j − 1)</i>	0.127	0.019	0.089	0.165	
	<i>loss(j − 1)</i>	−0.094	0.014	−0.122	−0.066	
	<i>lme(j)</i>	2.947	0.550	1.869	4.025	
	<i>lme_volatility(j)</i>	−3.353	1.075	−5.461	−1.245	
(b): Interpurchase event time decision						
State 1	<i>Intercept</i>	0.994	0.030	0.936	1.052	
	<i>quantity(j − 1)</i>	−0.010	0.030	−0.068	0.048	
	<i>gain(j − 1)</i>	−0.067	0.014	−0.095	−0.039	
	<i>loss(j − 1)</i>	0.050	0.014	0.022	0.078	
State 2	<i>Intercept</i>	0.858	0.032	0.796	0.920	
	<i>quantity(j − 1)</i>	0.031	0.021	−0.011	0.073	
	<i>gain(j − 1)</i>	−0.016	0.016	−0.048	0.016	
	<i>loss(j − 1)</i>	0.040	0.010	0.020	0.060	
(c): Quote request decision (quote request vs. direct order behavior, where quote request = 1)						
State 1	<i>Intercept</i>	−1.103	0.035	−1.171	−1.035	
	<i>quantity(j)</i>	0.403	0.099	0.209	0.597	
	<i>interpurchase time(j)</i>	0.029	0.004	0.021	0.037	
	<i>gain(j − 1)</i>	−0.085	0.027	−0.137	−0.033	
	<i>loss(j − 1)</i>	0.094	0.024	0.046	0.142	
	<i>lme(j)</i>	−0.980	0.896	−2.736	0.776	
	<i>lme_volatility(j)</i>	1.184	1.077	−0.928	3.296	
	State 2	<i>Intercept</i>	1.040	0.082	0.880	1.200
		<i>quantity(j)</i>	2.119	0.587	0.969	3.269
		<i>interpurchase time(j)</i>	0.004	0.007	−0.010	0.018
<i>gain(j − 1)</i>		−0.340	0.015	−0.370	−0.310	
<i>loss(j − 1)</i>		0.721	0.017	0.687	0.755	
<i>lme(j)</i>		−1.850	0.601	−3.028	−0.672	
<i>lme_volatility(j)</i>		3.319	1.319	0.733	5.905	

Our empirical results are consistent with the theoretical literature in B2B relationship marketing, which posits that as relationships improve, buyers in the supply chain focus holistically on the relationship and on the long-term benefits that such a relationship provides (Dwyer et al. 1987). Relaxed buyers are also not easily affected by adverse external environmental changes because the relationship provides a safe harbor and a level of stability in an otherwise unstable business environment (Fang et al. 2011, Jaworski and Kohli 1993, Palmatier 2008). The two states we have identified have also been referred to in the B2B literature as “programmed” and “transactional” segments of buyers (Rangan et al. 1992), “passive” and “aggressive” buyers (Shapiro et al. 1987), and “co-

Table 4 (Cont'd.)

State	Parameter	Mean	Std. dev.	Quantile	
				2.5%	97.5%
(d): Quote acceptance decision (quote acceptance vs. rejection behavior, where accept = 1)					
State 1	<i>Intercept</i>	0.217	0.015	0.187	0.247
	<i>quantity(j)</i>	−0.212	0.033	−0.276	−0.148
	<i>interpurchase</i>	0.002	0.017	−0.032	0.036
	<i>time(j)</i>				
	<i>gain(j)</i>	0.206	0.015	0.176	0.236
	<i>loss(j)</i>	−0.182	0.016	−0.214	−0.150
	<i>quantity(j)</i>	0.081	0.018	0.045	0.117
	<i>× gain(j)</i>				
	<i>quantity(j)</i>	−0.049	0.016	−0.081	−0.017
	<i>× loss(j)</i>				
State 2	<i>lme(j)</i>	0.212	0.510	−0.788	1.212
	<i>lme_volatility(j)</i>	−0.351	1.077	−2.463	1.761
	<i>Intercept</i>	0.101	0.050	0.003	0.199
	<i>quantity(j)</i>	−0.525	0.047	−0.617	−0.433
	<i>Interpurchase</i>	−0.031	0.008	−0.047	−0.015
	<i>time(j)</i>				
	<i>gain(j)</i>	0.127	0.015	0.097	0.157
	<i>loss(j)</i>	−0.219	0.020	−0.259	−0.179
	<i>quantity(j)</i>	0.260	0.015	0.230	0.290
	<i>× gain(j)</i>				
	<i>quantity(j)</i>	−1.662	0.018	−1.698	−1.626
	<i>× loss(j)</i>				
	<i>lme(j)</i>	0.719	0.469	−0.201	1.639
	<i>lme_volatility(j)</i>	−1.311	1.145	−3.557	0.935

Notes. Posterior means and standard deviations are calculated across the MCMC draws. Estimates in bold indicate a significant effect (that is, 95% posterior interval excludes 0).

operative” and “antagonistic” buyers (Matthyssens and Van den Bulte 1994).

5.6.2. Using Pricing to Drive Buyer Dynamics.

Buyers can transition between the two states over time. The parameter estimates in Table 6 and their transition matrix representation in Table 7 illustrate these dynamics. The central matrix in Table 7 shows the transition matrix when the price equals the reference price (i.e., the buyer’s quantity-weighted average historical prices). One can see that the states are rel-

Table 5 Description of the Two HMM States

	Relaxed state	Vigilant state
Quote request probability (%)	24	83
Quote accept probability (%)	63	54
Average quantity ordered (lb.)	432	502
Interpurchase event time (weeks)	5.5	8.1
Average price elasticity ^a	1.4	3.1
Average loss aversion ratio ^a	0.89	3.01
Average sensitivity to LME ^a	0.8	6.2

^aPrice elasticity, loss aversion, and LME sensitivity are averaged across buyers and across the four decisions. Because the sign of the effect can vary across decisions (e.g., price elasticity is negative for quantity but positive for interpurchase event time), we average the absolute value of these measures.

Table 6 HMM and Distributional Parameter Estimates

				Quantile	
	Parameter	Mean	Std. dev.	2.5%	97.5%
Transition matrix					
State 1	<i>Intercept</i>	1.787	0.018	1.750	1.824
	<i>gain</i> ($j - 1$)	0.045	0.013	0.019	0.070
	<i>loss</i> ($j - 1$)	−0.090	0.013	−0.116	−0.065
State 2	<i>Intercept</i>	2.162	0.037	2.091	2.231
	<i>gain</i> ($j - 1$)	−0.434	0.015	−0.464	−0.405
	<i>loss</i> ($j - 1$)	0.612	0.016	0.581	0.643
Initial state membership probability	Vigilant (ψ)	0.683	0.029	0.626	0.737
Distributional parameters					
	Std. dev. for the quantity model, log scale (σ)	0.108	0.045	0.023	0.191
	State 1 shape parameter for interpurchase event time, log scale (α_1)	0.100	0.016	0.068	0.131
	State 2 shape parameter for interpurchase event time, log scale (α_2)	0.068	0.010	0.047	0.088

Notes. Posterior means and standard deviations are calculated across the MCMC draws. Estimates in bold indicate a significant effect (that is, 95% posterior interval exclude 0).

actively “sticky.” A somewhat alarming result for the seller is that the likelihood of a buyer dropping from the relaxed to the vigilant state (14.3%) is almost twice as high as the likelihood of moving from the vigilant to the relaxed state (7.7%). However, the seller can use its pricing policy to affect transitions between the states.

Comparing the left matrix in Table 7 with the central matrix, we see that experiencing a 10% price decrease (“gain”) in the previous purchase event increases the probability of moving from the vigilant to the relaxed state by almost 2% and also increases the chance of remaining in the relaxed state by 3%. This suggests that when buyers perceive that they are treated well, they are more likely to transition to or remain in a state of a more favorable relationship with the seller. In contrast, the right matrix in Table 7 shows that a 10% price increase (“loss”) on the previous purchase event increases the likelihood that the buyer will transition to the vigilant state by almost

Table 7 Posterior Mean of the Transition Matrix Across Buyers

	10% price decrease		Average price (reference price)		10% price increase	
	Relaxed ($j + 1$)	Vigilant ($j + 1$)	Relaxed ($j + 1$)	Vigilant ($j + 1$)	Relaxed ($j + 1$)	Vigilant ($j + 1$)
Relaxed (j)	0.886	0.114	0.857	0.143	0.789	0.211
Vigilant (j)	0.095	0.905	0.077	0.923	0.043	0.957

7% and also increases the likelihood of staying in the vigilant state by 3.4%. Thus, a price increase may have a long-term effect by transitioning the buyer to a (sticky) state of increased price sensitivity. It should be noted that Table 7 presents transition matrices computed at the posterior mean. Our MCMC estimation permits an accounting of posterior uncertainty for the transition matrices of each buyer. We leverage this heterogeneity in developing a targeted pricing policy in §6. The average loss aversion in the impact of reference prices on the transition between the states ranges from 1.5 to 2.5, consistent with the loss aversion ratios commonly reported in B2C applications.

5.6.3. Investigating Price Endogeneity. To assess the extent of endogeneity, we look at the correlations between the unobserved error in the price equation and each of the random shocks (ξ 's) for the four decisions. There are mildly negative correlations for the quantity and quote acceptance decisions and mildly positive correlations for the interpurchase event time and quote request decisions. These correlations are all in the expected direction and imply that (1) despite the predominant practice of cost-plus pricing, the seller does occasionally target price-insensitive (sensitive) buyers by charging them higher (lower) prices and that (2) unobserved shocks could influence both the seller's pricing decisions and the buyers' behavior. Although a failure to properly account for price endogeneity could result in overestimation of the effects of price gains and an underestimation of the impact of price losses, it should be noted that, overall, the differences in the price sensitivity estimates between the two models are relatively small, and accounting for endogeneity only results in a modest gain in our application. Estimates of the correlations between the unobserved error in the price equation and each of the random shocks for the four decisions, and a comparison of the parameter estimates are available in the Web appendix.

5.7. Disentangling the Short- and Long-Term Effects of Pricing

Assessing the marginal and integrative impact of price is not straightforward from the reference price coefficients in Tables 4 and 6 because price enters in our model in multiple places. We therefore numerically compute the short-run and long-run elasticities for each of the four decision components and for the HMM state membership probabilities. The elasticity is calculated for each decision variable using a one-time shock (price increase or a price decrease) of 10% in the unit price from the average price. We take a horizon of 20 simulated purchase events subsequent to the one-time shock to calculate the short-term and long-term elasticities. Following the one-time price shock, we

Table 8 Short- and Long-Term Price Elasticities of the Decision Components and the Vigilant State Membership

	Price increase			Price decrease		
	Short-term	Long-term	Total	Short-term	Long-term	Total
Quantity	−0.43	−2.73	−3.21	1.39	3.05	4.53
Interpurchase event time	1.22	2.34	3.57	−0.71	−0.62	−1.35
Quote request	0.73	4.91	5.72	−0.53	−1.59	−2.16
Quote acceptance	−1.19	−3.42	−4.65	0.38	0.97	1.37
Vigilant state membership	0.62	5.47	6.18	−0.31	−1.38	−1.72


set prices to the reference price level for the remaining 19 purchase events. The short-term elasticity captures the immediate impact (i.e., on the first purchase event), whereas the long-term elasticity captures the effect over the next 19 purchase events. These short- and long-term elasticities are reported in Table 8. In addition, Figure 3 illustrates the asymmetry in the percentage increases and decreases in each affected variable over the simulated 20 purchase events following the one-time 10% price increase or decrease.

Several insights can be gleaned from Table 8 and Figure 3. First, all price elasticities are in the expected direction. Second, the magnitudes of the long-term elasticities are generally much larger than the corresponding magnitudes of the short-term elasticities. On average, the short-term price elasticities are only 10%–53% of the total price elasticities. The average short-term quantity elasticities are consistent with those reported in the literature (Bijmolt et al. 2005, Jedidi et al. 1999, Tellis 1988). Third, factors that are directly related to the buyers' attitudes and relationship with the seller (i.e., quote request and vigilant

state membership) exhibit stronger long-term elasticities relative to other factors. Fourth, consistent with prospect theory (Kahneman and Tversky 1979), we find that price increases (losses) have a stronger impact than do price decreases (gains) for all decisions, except for quantity. The quantity decision, in contrast, exhibits gain seeking. Previous research in B2C has found gain seeking in the effect of reference prices on quantity when households face low inventory (Mazumdar et al. 2005). This result is consistent with buyers keeping a low inventory of aluminum and relying on the reseller to stock the material. Thus, consistent with Bruno et al. (2012), we find empirical evidence for asymmetric reference price effects for B2B buyers. Fifth, and perhaps most interestingly, Figure 3 shows that the negative effects of a price hike (loss) on all four decisions persist longer than the positive effects of price drop (gain). This result is consistent with the literature on B2B relationship life cycles that states that damage to the relationship is difficult to repair as it leaves "psychological scars" (Dwyer et al. 1987, Jap and Anderson 2007, Ring and Van de Ven 1994). This evidence for the longer persistence of the effects of negative price experiences is also consistent with *hedonic adaptation theory*, which states that individuals adapt faster to improvements than to deteriorations (Frederick and Loewenstein 1999). To the best of our knowledge, this is the first empirical investigation of the long-term effects of asymmetric reference price effects and the first demonstration of the hedonic adaption theory using actual transactional data.

Overall, these results imply that, in B2B contexts, researchers and managers that consider only the short-term effects of pricing can significantly and substantially underestimate the overall impact of pricing,

Figure 3 Duration of Asymmetric Price Effects on the Four Decision Components



as they ignore the impact of pricing on reference prices and on the latent relationships that drive long-term profitability. In the next section, we investigate how firms can use the model's parameters and leverage the resulting behavioral insights to dynamically target individual buyers.

6. Targeted and Dynamic Price Policy Optimization

In this section, we assess the separate and varied influences of the seller's pricing decisions on the behavior of buyers and consequently on the seller's medium to long-term profits. We first define the seller's profit from purchase event j of buyer i as

$$\text{Profit}_{ij} = (\text{Price}_{ij} - \text{Cost}_{ij})q_{ij}[(1 - \delta_{ij}^b) + \delta_{ij}^b\delta_{ij}^w],$$

where δ_{ij}^b is a binary indicator that equals 1 when buyer i requests a quote on purchase event j and is 0 for a direct order. The binary indicator δ_{ij}^w equals 1 when the quote is accepted and is 0 otherwise. The seller's objective is to maximize the medium to long-run profit (over T periods) across buyers. Therefore, for each buyer i , the seller sets the sequence of prices $\mathbf{p}_i$ to one that maximizes

$$\begin{aligned} \max_{\mathbf{p}_i} \quad & \sum_{j=1}^{J_i(\mathbf{p}_i)} \frac{\text{Profit}_{ij}(\mathbf{p}_i)}{(1+r)^{\sum_{k=1}^j t_{ik}(\mathbf{p}_i)}} \\ \text{s.t.} \quad & \sum_{j=1}^{J_i(\mathbf{p}_i)} t_{ij}(\mathbf{p}_i) < T, \end{aligned} \quad (11)$$

where $t_{ij}(\mathbf{p}_i)$ is the j th interpurchase event time and r represents the discount rate. As the interpurchase event times are influenced by the pricing decisions of the seller, the number of purchase events, $J_i(\mathbf{p}_i)$, is endogenously determined by the constraint that the sum of the interpurchase event times does not exceed the length of the planning horizon (T). We can obtain the seller's overall profit by summing the optimal profits across all buyers.

We conduct the optimization using a 15-month horizon. However, we evaluate the performance of our approach over the first nine months of the planning horizon to limit the impact of end-of-the horizon effects in the optimization. We therefore split the data into a calibration sample covering the duration from January to December 2007 and a holdout period of nine months ranging from January 2008 to September 2008.⁸ We then use the parameter estimates from

the calibration data set to conduct the price optimization. The optimization is performed for a representative sample of 300 buyers who experienced between 6 and 16 purchase events over the calibration period and an average of 10 purchase events over the nine months of our holdout data.

We discretize the continuous pricing decision on any purchase event and use a set of five buyer-specific price points that form the quintiles of the distribution of prices that the buyer experienced in the calibration period. We choose to stay within each buyer's historical range of experienced prices to avoid a drastic change in the price regime observed by each buyer, yet still account for the price variance experienced by each buyer.

We then use a combination of forward simulation and complete enumeration over all feasible price paths to obtain the set of optimal prices over the simulated purchase events of the buyer in the 15-month planning horizon. The optimization process is initialized for each buyer by setting the state membership probabilities and the reference price to their values at the end of the calibration period. The forward simulation then proceeds by generating a sequence of purchase events. Given the five feasible price points at each purchase event, there are $5^{J_{ik}}$ possible price nodes for buyer i within the k th random sequence. We account for different sources of uncertainty in computing the profits for each buyer. Given a vector of parameters for buyer i , we simulate 200 Monte Carlo sequences of purchase events for the buyer for each price path. These 200 sequences may differ in the number of purchase events, with the k th random sequence containing J_{ik} purchase events in the planning horizon. Each purchase event in the simulated sequence is characterized by the quantity, interpurchase event time, quote request decision, associated reference price, and latent state membership probabilities. We also account for parameter uncertainty by repeating the optimization procedure for each of the well-separated 100 draws from the MCMC posterior distribution of the parameters for each buyer. At each simulated purchase event, profits are computed by weighting the HMM latent state-specific profits by the state membership probabilities, and they are averaged over the 200 sequences and over the 100 draws from the MCMC posterior distribution to compute the average profit of each price path. We assume an annual discount rate of 12% (a weekly discount rate r of 0.22%).⁹ Full details of the price simulation

⁸ For the purpose of price policy simulation, we reestimate the proposed model on the first 12 months of data and perform simulation on the subsequent 9 months. The estimates do not differ substantially from those using 18 months of data, as reported in Table 4.

⁹ We tested the improvement in precision that can be gained from increasing the number of random sequence draws. We choose 200 draws as a good compromise between precision and computational time because it offers 8% improvement in profits over 100 draws, but it only underperforms 300 draws by 1%. It should also be noted that because of the dimensionality of the state space, our simulation serves as a heuristic for the optimal prices.

Table 9 Policy Performance Comparison Over Nine Months

	Policy						
	Current pricing	Individual dynamic pricing	Individual static pricing	Segment dynamic pricing	Aggregate static pricing	Individual dynamic pricing (myopic)	Observed state pricing
Average profit per customer (\$)	3,158	4,809	3,867	3,328	2,467	4,117	4,008
Mean price (\$/lb.)	3.24	3.40	3.28	3.44	3.41	3.28	3.44
Median price (\$/lb.)	2.87	3.30	3.11	3.23	N/A	3.11	3.28

and optimization are available from the authors upon request.

6.1. Price Policy Simulation Results

We now compare the performance of our proposed individually targeted dynamic pricing policy to that of the five competing policies:

1. *Individually targeted static pricing policy*: In this policy, a single price is determined for each buyer for the entire 15-month planning horizon. This policy leverages the heterogeneity in the model's estimates across buyers but ignores dynamics.

2. *Segment-targeted dynamic policy*: In this policy, only two optimal prices are determined—one price for each of the two states. The price at a given purchase event therefore depends on the HMM state membership at each purchase event.¹⁰

3. *Aggregate single-price static policy*: This policy chooses a single price for all buyers for the entire planning horizon. Thus, this policy ignores both heterogeneity and dynamics.

4. *Myopic individually targeted dynamic policy*: In this policy, the seller accounts for both the buyers' updated latent state membership and the heterogeneity in the buyers' response parameters. However, at each purchase event, the seller maximizes profits only for the current purchase event, as opposed to the entire planning horizon. Thus, the myopic dynamic policy considers only the short-term effect of pricing in each period.

5. *Observed state policy*: This policy corresponds to Benchmark 4 in §5.5. As it may be computationally difficult for the seller to infer the buyer's latent state, we investigate optimization based on observed proxies of the latent state—namely, whether the buyer requested a quote in the previous period. This policy can be thought of as a “simple” heuristic to our proposed policy.


6. *Current policy*: This is the seller's current pricing policy for the nine months.

A comparison of the results from the alternative policies highlights the marginal improvements

in profitability that stems from individual-level targeting, from dynamics pricing, and from adopting a long-term perspective. Table 9 shows how the seven price policies perform. The proposed policy yields the highest profits per buyer of \$4,809 over nine months. Leveraging heterogeneity, as given by the individually targeted static policy, generates a 57% improvement in profits when compared with the aggregate static policy (\$3,867 versus \$2,467). An additional 24% improvement results from dynamic targeting (\$4,809 versus \$3,867). This result is consistent with the findings of Khan et al. (2009), who highlight the potential gains from intertemporal targeting. Similarly, the proposed policy improves profits by 17% over the myopic policy by leveraging the dynamic impact of pricing and by adopting a long-term perspective. The proposed policy yields a 52% improvement in profit compared with the seller's current policy. Taking into account the entire customer base of the seller, this translates to a potential profit improvement of approximately \$4 million annually. Even employing the easier to use observed state policy can generate a 27% profit gain.

To compare the performance of the different policies over time, we plot the average monthly profits over the first nine months of the planning horizon in

Figure 4 Policy Performance Comparison Over Nine Months



¹⁰ In the segment policies, the latent state membership at each purchase event is determined by the state with the highest membership probability.

Figure 4. We see that the proposed policy carefully balances the interplay between several forces that govern buyers' buying behaviors: (1) charging lower prices to increase quote acceptance, (2) charging lower prices to keep buyers in the relaxed state or to transition them to it, (3) charging higher prices to increase margins, and (4) charging higher prices to keep buyers' reference prices high. The myopic policy, on the other hand, ignores points (2) and (4) and therefore charges lower prices than the proposed policy, aiming to convert quote requests into immediate sales (see Table 9). Although the myopic strategy leads to higher profits than the proposed policy in the first few months, charging lower prices results in lower reference prices and creates a downward pressure on the seller to continue offering lower prices, which results in a vicious cycle of decreasing prices and hence depressed profits. After the first three months, the proposed policy begins to outperform the myopic policy demonstrating the importance of using a long-term perspective when setting prices. These results suggest that in the world of B2B, where relationships are long term and sticky, myopia in price setting can be a slippery slope.

With regard to prices charged, the proposed price policy recommends a higher price for buyers in the relaxed state versus those that are in the vigilant state (\$3.63/lb. versus \$3.20/lb.). The current policy that is used by the seller, however, only mildly differentiates between buyers in the two latent states. Charging a higher price for buyers in the relaxed state can increase both the immediate profits and the reference prices in the long term. Although a higher price might increase the chance of transitioning these buyers into the vigilant state, the stickiness of the relaxed state and the already increased reference prices act as a "shield" against perceiving future prices as losses.


In summary, the superior profitability of our individual dynamic targeted price policy, relative to the current policy, stems from its ability to leverage (1) heterogeneity in price sensitivities, (2) differences across the latent states so that higher prices can be charged to the less price-sensitive customers in the relaxed state, and (3) the trade-off between the short-run and long-run dynamic effects of pricing.

6.2. Pricing in a Volatile Economic Environment—The Role of External Reference Price

In this section, we examine the optimal price strategies that the seller should use to manage volatile economic conditions. Specifically, we investigate to what extent the seller should pass through its cost increases when aluminum prices rise and whether the seller needs to reduce its prices when aluminum prices fall.

The aluminum prices on the LME fluctuated between US\$2,393 to US\$3,318 per tonne over the

Figure 5 Optimal and Current Unit Price per Pound by State, Under Different LME Regimes



duration of our data. A change in the commodity prices can lead to at least two opposing impacts on the profitability of the seller. On the one hand, aluminum prices influence the seller's cost of replenishment.¹¹ On the other hand, buyers use the price on the LME as an external reference price. It is therefore unclear, a priori, how the seller should change its pricing strategy in response to such fluctuations in the commodity market. We investigate this question empirically by computing and comparing our individual dynamic price policies under two scenarios: (1) a 20% increase in the LME prices over the actual LME prices in the nine months of the planning horizon and (2) a 20% decrease over the same period.¹²

Figure 5 shows how the seller should differentially adjust its prices for buyers in the relaxed and vigilant states. When the LME price increases by 20%, the seller should increase the unit prices by 12% for buyers in the vigilant state but by only 4.6% for buyers in the relaxed state. This result stems from the higher sensitivity to the LME prices for buyers who are in the vigilant state (see Table 4). Figure 5 also shows that when LME prices drop by 20%, it is optimal for the firm to "hoard" most of the cost savings and drop prices by only 2.5%–2.8% for buyers in both states. The rationale here is that lowering the price results

¹¹ The seller in our empirical application keeps as much as six months of inventory, so the relationship between the LME prices and wholesale prices of currently sold orders is relatively weak ($R^2 = 0.05$), but the LME impacts directly the wholesale cost of replenishment.

¹² Our use of a 20% shock in LME falls within the range of fluctuations observed during the data period. The maximum daily, monthly, and three-month LME fluctuations in our data were 5.6%, 22%, and 31%, respectively.

in a corresponding lowering of the internal reference price, and this can have long-term consequences for the seller's profitability.

This price strategy of passing on the cost increase and "hoarding" the benefit of a cost decrease is consistent with the *dual-entitlement principle* (Akerlof 1979, Kahneman et al. 1986, Okun 1981, Urbany et al. 1989). The dual-entitlement principle is based on the notion of perceived fairness and states that (1) firms are "justified," in the eyes of customers, to increase prices when costs increase, to protect firms' normal profits; and (2) firms do not need to lower prices when costs drop as customers' perceptions are mainly driven by their past reference prices. Although we do not model fairness directly, the external and internal asymmetric reference price effects in tandem with the latent trust state capture similar effects. To the best of our knowledge, this is the first paper to empirically demonstrate the dual-entitlement principle and to measure its extent in a B2B transaction setting.

In summary, this analysis demonstrates that the seller should pass on most of the cost increases, especially to buyers who are in the vigilant state and are paying close attention to external reference prices. The seller should also make the external reference price more salient to buyers (particularly those in the vigilant state) during inflationary periods. In contrast, cost decreases present a good opportunity for the seller to enjoy a period of heightened profitability by keeping prices at the same level, at least in the short run.

7. General Discussion

Understanding and managing the impact of pricing on buyer behavior and on evolving business relationships is critical for the long-run profitability of B2B sellers. In this paper, we present an integrative empirical framework for modeling different buyer decisions via a common latent and dynamic state of trust and capture the long-term effect of pricing decisions via a Bayesian nonhomogeneous HMM.

We generate several substantive insights in the underresearched area of B2B pricing. First, we empirically uncover two relationship states (vigilant and relaxed) and show how pricing decisions can affect the transition between these two states over time. Specifically, we find that appropriate pricing decisions can increase trust and relationship quality (in the form of the relaxed state). Such a state can, in turn, act as a "relationship shelter" by encouraging a simplified buying process and can mitigate the adverse effects of external volatility. Second, we not only find significant asymmetric reference price effects, i.e., price "losses" loom larger than "gains," but that it takes much longer for buyers to adapt to losses—a result that is in sync with the theoretical research

on B2B relationships life cycles and with the psychological research on adaptation. Third, we find that strong relationships facilitate a simplified buying process and act as a shelter against adverse economic environments.

We conduct a series of price policy simulations and demonstrate that the proposed dynamic targeted price policy can offer a 52% improvement in long-term profitability over the seller's current pricing policy. Furthermore, we show that the profitability of buyers in the relaxed state is almost twice as high as those in the vigilant state. The proposed policy balances two forces: (1) lowering prices to win business and to keep buyers in the relaxed state and (2) increasing prices to maximize margins and to avoid lowering of internal reference prices.

Our simulation results regarding the optimal pricing policy under volatile commodity prices indicate that the seller should pass on a part of the increased costs to buyers, but it should hoard most of the benefit when costs decrease. This active management can help the seller maintain existing levels of profitability during inflationary periods and enjoy increased profitability during deflationary periods.

More generally, this research offers B2B sellers a comprehensive decision framework to manage their buyer base using dynamic price targeting. As many B2B sellers routinely apply cost-based pricing strategies (Anderson et al. 1993), we demonstrate that there is substantial value in leveraging B2B relationships to implement valued-based, first-degree intertemporal price discrimination. However, it is important to note that not all buyers are relationship oriented—some customers will evaluate each deal as a unique transaction. The seller needs to realize this heterogeneity so that it does not overinvest (via its pricing actions) in buyers who may have little chance of migrating to a higher relationship quality state.

We now highlight some limitations of our work and propose directions for future research. First, we assume that buyers are not forward looking with respect to the seller's pricing decisions. One could extend our framework to incorporate buyers' expectations about future price changes (e.g., Lewis 2005) in applications where such expectations are likely to be important. Second, as is typical in B2B contexts, our data set (and the data available to the company management) does not include competitive information. Although quote requests and unfulfilled quotes provide indirect evidence for competition, future researchers can extend our framework to settings in which competitive pricing data are available. This would potentially uncover richer reference price formation and relationship development processes.

Third, given our focus on targeting and dynamics, we use data on buyers with at least seven purchase events. These were the vast majority (92% of all buyers) in our context. Therefore, our conclusions and results are most suited to these more frequent buyers.

Fourth, we focus on profit maximization for each buyer. Other objective functions that treat certain buyers to be of strategic importance (e.g., for consistent cash flow or because of the possible impact on other buyers) or objective functions that focus on requirements such as revenue maximization can be explored. Fifth, we focus on the impact of the pricing decision on the seller's profitability. Future research can investigate whether price-induced relationship dynamics differ from those that are triggered by other marketing actions (Kumar et al. 2011).

B2B and B2C business frameworks are not orthogonal to each other—rather, we think of these markets as a continuum with substantial overlap. Although our pricing framework focuses on B2B relationships, it is also appropriate for those B2C settings in which the customer buying process is composed of several inter-related decisions, where the firm has the opportunity to price discriminate to varying degrees and where long-term relationships play a big role.

Finally, in this paper we take an initial step toward studying the underexplored terrain of B2B pricing using a specific empirical application within the metal industry. We encourage future researchers to apply our framework to other B2B environments (e.g., those with more involved negotiations or with power asymmetry in the supply chain) to investigate the generalizability of our findings.

Supplemental Material

Supplemental material to this paper is available at <http://dx.doi.org/10.1287/mksc.2013.0842>.

Acknowledgments

The authors are grateful to the aluminum-trading company who provided the invaluable data and the managerial insights to our inquiries. The authors also thank the editor, the associate editor, the two anonymous reviewers, as well as the participants of seminars and conferences in which this work has been presented, for their constructive comments.

References

Akerlof GA (1979) The case against conservative macroeconomics: An inaugural lecture. *Economica* 46(183):219–237.

Allenby GM, Leone RP, Lichung J (1999) A dynamic model of purchase timing with application to direct marketing. *J. Amer. Statist. Assoc.* 94(446):365–374.

Anderson E, Weitz B (1992) The use of pledges to build and sustain commitment in distribution channels. *J. Marketing Res.* 29(1):18–34.

Anderson JC, Jain DC, Chintagunta P (1993) Customer value assessment in business markets: A state of the practice study. *J. Bus.-to-Bus. Marketing* 1(1):3–28.

Ansari A, Mela CF (2003) E-customization. *J. Marketing Res.* 40(2):131–145.

Atchadé YF (2006) An adaptive version for the Metropolis adjusted Langevin algorithm with a truncated drift. *Methodol. Comput. Appl. Probab.* 8(2):235–254.

Bell DR, Bucklin RE (1999) The role of internal reference price in the category purchase decision. *J. Consumer Res.* 26(2):128–143.

Bell DR, Lattin JM (1998) Shopping behavior and consumer preference for store price format: Why “large basket” shoppers prefer EDLP. *Marketing Sci.* 17(1):66–88.

Bijmolt THA, van Heerde HJ, Pieters RGM (2005) New empirical generalizations on the determinants of price elasticity. *J. Marketing Res.* 42(2):141–156.

Bruno HA, Che H, Dutta S (2012) Role of reference price on price and quantity: Insights from B2B markets. *J. Marketing Res.* 49(5):640–654.

Cannon JP, Perreault WD Jr (1999) Buyer-seller relationships in business markets. *J. Marketing Res.* 36(4):439–460.

Chintagunta P (2002) Investigating category pricing behavior at a retail chain. *J. Marketing Res.* 39(2):141–154.

Dahl DW, Honea H, Manchanda RV (2005) Three Rs of interpersonal consumer guilt: Relationship, reciprocity, reparation. *J. Consumer Psych.* 15(4):307–315.

Dong X, Manchanda P, Chintagunta P (2009) Quantifying the benefits of individual level targeting in the presence of firm strategic behavior. *J. Marketing Res.* 46(4):207–221.

Dwyer FR, Tanner JF (2009) *Business Marketing: Connecting Strategy, Relationships, and Learning*, 4th ed. (McGraw-Hill, Boston).

Dwyer FR, Schurr PH, Oh S (1987) Developing buyer-seller relationships. *J. Marketing* 51(2):11–27.

Erdem T, Katz ML, Sun B (2010) A simple test for distinguishing between internal reference price theories. *Quant. Marketing Econom.* 8(3):303–332.

Fang E, Palmatier RW, Grewal R (2011) Effects of customer and innovation asset configuration strategies on firm performance. *J. Marketing Res.* 48(3):587–602.

Frederick S, Loewenstein G (1999) Hedonic adaptation. Kahneman D, Diener E, Schwartz N, eds. *Foundations of Hedonic Psychology: Scientific Perspectives on Enjoyment and Suffering* (Russell Sage Foundation, New York), 302–329.

Hamilton JD (1989) A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* 57(2):357–384.

Hardie BGS, Johnson EJ, Fader PS (1993) Modeling loss aversion and reference dependence effects on brand choice. *Marketing Sci.* 12(4):378–394.

Heckman JJ (1981) Statistical models for discrete panel data. Manski CF, McFadden D, eds. *Structural Analysis of Discrete Panel Data with Econometric Applications* (MIT Press, Cambridge, MA), 114–178.

Hinterhuber A (2004) Towards value-based pricing—An integrative framework for decision making. *Indust. Marketing Management* 33(8):765–778.

Jap SD, Anderson E (2007) Testing a life-cycle theory of cooperative interorganizational relationships: Movement across stages and performance. *Management Sci.* 53(2):260–275.

Jaworski BJ, Kohli AK (1993) Market orientation: Antecedents and consequences. *J. Marketing* 57(July):53–70.

Jedidi K, Mela CF, Gupta S (1999) Managing advertising and promotion for long-run profitability. *Marketing Sci.* 18(1):1–22.

Johnston WJ, Lewin JE (1996) Organizational buying behavior: Toward an integrative framework. *J. Bus. Res.* 35(1):1–15.

- Kahneman D, Tversky A (1979) Prospect theory: An analysis of decision under risk. *Econometrica* 47(2):263–291.
- Kahneman D, Knetsch JL, Thaler RH (1986) Fairness as a constraint on profit seeking: Entitlements in the market. *Amer. Econom. Rev.* 76(4):728–741.
- Kalwani MU, Narayandas N (1995) Long-term manufacturer-supplier relationships: Do they pay off for supplier firms? *J. Marketing* 59(1):1–16.
- Kalwani MU, Rinne HJ, Sugita Y, Yim CK (1990) A price expectations model of customer brand choice. *J. Marketing Res.* 27(8):251–262.
- Kalyanaram G, Winer RS (1995) Empirical generalizations from reference price research. *Marketing Sci.* 14(3, Part 2):G161–G169.
- Khan R, Lewis M, Singh V (2009) Dynamic customer management and the value of one-to-one marketing. *Marketing Sci.* 28(6):1063–1079.
- Kopalle PK, Lindsey-Mullikin J (2003) The impact of external reference price on consumer price expectations. *J. Retailing* 79(4):225–236.
- Krishnamurthi L, Mazumdar T, Raj SP (1992) Asymmetric response to price in consumer choice and purchase quantity decisions. *J. Consumer Res.* 19(12):387–400.
- Kumar V, Venkatesan R, Reinartz W (2008) Performance implications of adopting a customer-focused sales campaign. *J. Marketing* 72(5):50–68.
- Kumar V, Sriram S, Luo A, Chintagunta PK (2011) Assessing the effect of marketing investments in a business marketing context. *Marketing Sci.* 30(5):924–940.
- Lancaster T (1990) *The Econometric Analysis of Transition Data* (Cambridge University Press, Cambridge, UK).
- LaPlaca P, Katrichis J (2009) Relative presence of business-to-business research in the marketing literature. *J. Bus.-to-Bus. Marketing* 16(1):1–22.
- Lehmann DR, O’Shaughnessy J (1974) Difference in attribute importance for different industrial products. *J. Marketing* 38(2):36–42.
- Lewis M (2005) A dynamic programming approach to customer relationship pricing. *Management Sci.* 51(6):986–994.
- Liozu S (2012) The need for more academic research in pricing. *Pricing Blog* (blog), May 10, <http://professionalpricingsociety.blogspot.com/2012/05/guest-post-need-for-more-academic.html>.
- MacDonald IL, Zucchini W (1997) *Hidden Markov and Other Models for Discrete-Valued Time Series* (Chapman & Hall, London).
- Matthyssens P, Van den Bulte C (1994) Getting closer and nicer: Partnerships in the supply chain. *Long Range Planning* 27(1):72–83.
- Mazumdar T, Raj SP, Sinha I (2005) Reference price research: Review and propositions. *J. Marketing* 69(4):84–102.
- Milgrom P, Weber R (1982) A theory of auctions and competitive bidding. *Econometrica* 50(5):1089–1122.
- Mithas S, Jones JL (2007) Do auction parameters affect buyer surplus in e-auctions for procurement? *Production Oper. Management* 16(4):455–470.
- Montoya R, Netzer O, Jedidi K (2010) Dynamic allocation of pharmaceutical detailing and sampling for long-term profitability. *Marketing Sci.* 29(5):909–924.
- Moorman C, Zaltman G, Deshpande R (1992) Relationships between providers and users of market research: The dynamics of trust. *J. Marketing Res.* 29(3):314–328.
- Morgan RM, Hunt SD (1994) The commitment-trust theory of relationship marketing. *J. Marketing* 58(3):20–38.
- Netzer O, Lattin JM, Srinivasan V (2008) A hidden Markov model of customer relationship dynamics. *Marketing Sci.* 27(2):185–204.
- Okun AM (1981) *Price and Quantities: A Macroeconomic Analysis* (Brookings Institute, Washington, DC).
- Palmatier RW (2008) *Relationship Marketing*, Monograph on Relationship Marketing (Marketing Science Institute, Cambridge, MA).
- Palmatier RW, Dant RP, Grewal D, Evans KR (2006) Factors influencing the effectiveness of relationship marketing: A meta-analysis. *J. Marketing* 70(4):136–153.
- Park S, Gupta S (2009) A simulated maximum likelihood estimator for the random coefficient logit model using aggregate data. *J. Marketing Res.* 46(4):531–542.
- Petrin A, Train K (2010) A control function approach to endogeneity in consumer choice models. *J. Marketing Res.* 47(1):3–13.
- Putler DS (1992) Incorporating reference price effects into a theory of consumer choice. *Marketing Sci.* 11(3):287–309.
- Rangan VK, Moriarty RT, Swartz GS (1992) Segmenting customers in mature industrial markets. *J. Marketing* 56(4):72–82.
- Reid DA, Plank RE (2004) *Fundamentals of Business Marketing Research* (Haworth Press, New York).
- Ring PS, Van de Ven AH (1994) Developmental processes of cooperative interorganizational relationships. *Acad. Management Rev.* 19(1):90–118.
- Rossi PE, Allenby GM, McCulloch R (2005) *Bayesian Statistics and Marketing* (John Wiley & Sons, New York).
- Rossi PE, McCulloch R, Allenby GM (1996) The value of purchase history data in target marketing. *Marketing Sci.* 15(4):321–340.
- Shapiro BP, Rangan VK, Moriarty RT, Ross EB (1987) Manage customers for profits (not just sales). *Harvard Bus. Rev.* 65(5):101–108.
- Schweidel DA, Bradlow ET, Fader PS (2011) Portfolio dynamics for customers of a multiservice provider. *Management Sci.* 57(3):471–486.
- Shipley D, Jobber D (2001) Integrative pricing via the pricing wheel. *Indust. Marketing Management* 30(3):301–314.
- Simester DI, Sun P, Tsitsiklis JN (2006) Dynamic catalog mailing policies. *Management Sci.* 52(5):683–696.
- Stein T (2013) The heat—for B2B marketing—is rising. *CMO.com* (May 2) http://www.cmo.com/articles/2013/5/2/the_heat_for_b2b_mar.html.
- Tellis G (1988) The price elasticity of selective demand: A meta-analysis of econometric models of sales. *J. Marketing Res.* 25(4):331–341.
- Urbany JE, Madden TJ, Dickson PR (1989) All’s not fair in pricing: An initial look at the dual entitlement principle. *Marketing Lett.* 1(1):17–25.
- Venkatesan R, Kumar V (2004) A customer lifetime value framework for customer selection and resource allocation strategy. *J. Marketing* 68(4):106–125.
- Villas-Boas JM, Winer RS (1999) Endogeneity in brand choice models. *Management Sci.* 45(10):1324–1338.
- Wedel M, Leeflang PS (1998) A model for the effects of psychological pricing in Gabor–Granger price studies. *J. Econom. Psych.* 19(2):237–260.
- Wilson EJ (1994) The relative importance of supplier selection criteria: A review and update. *J. Supply Chain Management* 30(3):34–41.
- Winer RS (1986) A reference price model of demand for frequently-purchased products. *J. Consumer Res.* 13(September):250–256.