Dynamic Targeted Pricing in B2B Settings

Jonathan Z. Zhang*
Oded Netzer
Asim Ansari

December 2011

Jonathan Zhang is an Assistant Professor at the Michael G. Foster School of Business, University of Washington (email: zaozao@uw.edu). Oded Netzer is the Phillip H. Geier Jr. Associate Professor of Business at Columbia Business School (email: on2110@columbia.edu). Asim Ansari is the William T. Dillard Professor of Business at Columbia Business School (email: maa48@columbia.edu).

Abstract

We model the multifaceted impact of pricing decisions in B2B contexts and show how a seller can develop optimal inter-temporal targeted pricing strategies to maximize long-term customer value. We empirically model the B2B customer's purchase decisions in an integrated fashion. In order to facilitate targeting and to capture the short and long-term dynamics of B2B customer purchasing, our modeling framework weaves together in a hierarchical Bayesian manner, multivariate copulas, a non-homogeneous hidden Markov model, and control functions for price endogeneity. We estimate our model on longitudinal transactions data from an aluminum retailer. We find that customers in our dataset can be best represented by two latent states - a "vigilant" state characterized by heightened price sensitivity and a cautious approach to ordering, and a more "relaxed" state. The seller's pricing decisions can transition customers between these two states. An optimal dynamic and targeted pricing strategy based on our model suggests a 52% improvement in profitability compared to the status quo. Furthermore, a counterfactual analysis which examines the optimal policy under fluctuating commodity prices reveals that the seller should pass much of the costs to customers when commodity prices increase, but hoard most of the profit when commodity prices (seller's costs) decrease.

1. Introduction

Despite the major role of the business-to-business (B2B) sector in the U.S. and world economy, marketing modelers have paid scant attention to B2B issues. While the B2B sector commands more than 50% market share of all commerce in the U.S. (U.S. Department of Commerce 2007), only a small fraction (approximately 3.4%) of the articles in the top four marketing journals deal with B2B contexts (LaPlaca and Katrichis 2009). Among the different marketing decisions, pricing in B2B environments is particularly under-researched (Reid and Plank 2004). In this paper, we address this imbalance by developing an integrated framework for modeling the multiple impacts of pricing in B2B contexts. We then illustrate how our framework can aid sellers in implementing first-degree and inter-temporal price discrimination for long-run profitability.

Pricing decisions in B2B contexts are different from those within business-to-consumer (B2C) environments on multiple facets. First, B2B settings are often characterized by highly variable costs of goods and order sizes, the prevalence of product and service customization, and the reliance on personal selling to cement transactions. These factors create a fertile environment for implementing first degree and inter-temporal price discrimination. Sellers in many B2B situations can easily vary prices across buyers and can even change prices between subsequent purchases of the same buyer. This flexibility generates significant opportunities for sellers to pursue long-term profitability. In contrast, B2C retailers rely on a restricted set of targeted price discrimination mechanisms (such as coupons), and targeted pricing is difficult, in general, because of logistical and ethical concerns.

Second, the B2B environment involves evolving long-term relationships between buyers and sellers (Morgan and Hunt 1994). The development of trust in such relationships can impact price sensitivities over time. Pricing decisions, in turn, can play a vital role in developing and sustaining such relationships (Kalwani and Narayandas 1995). Third, transactions in B2B markets exhibit greater complexity as the business customer typically makes several inter-related decisions. Specifically, B2B buyers not only choose what, when, and how much to buy, but also decide on how to buy. In B2B settings, buyers often choose whether to ask for a price quote (offering the seller the opportunity to provide a price quote) or whether to order directly from the seller without asking for a price. Such requests for a price quote allow sellers to observe demand and price sensitivity even when a sale is not made (i.e., when a bid is made and the buyer rejects the bid). Such data are rarely observed in B2C settings (Khan et al. 2009). Similarly, the buyer's decision to order directly from the seller (without observing the price) indicates the strength of relationship. Our modeling approach accommodates this unique aspect of the buying process.

Finally, decision makers (buyers and sellers) in B2B settings are often assumed to behave rationally (Reid and Plank 2004). Thus, while behavioral pricing theory suggests that internally constructed reference prices play an important role in customer purchase decisions, and that loss aversion should be considered in modeling reference price effects (Kalyanaram and Winer 1995), it is not clear whether such behavioral effects are operant in the B2B domain.

In this paper, we develop a modeling framework that incorporates these distinguishing aspects of the B2B environment. Our framework models customer decisions on each potential purchasing occasion in an integrated fashion using different model components. It uses hierarchical Bayesian copulas (Nelsen 2006; Pitt et al. 2006; Trivedi and Zimmer 2007) to flexibly model the four decisions (purchase timing, purchase amount, quote request, and quote acceptance) jointly. It accounts for heterogeneity in customer preferences and behaviors to facilitate targeting, incorporates asymmetric reference price effects and other behavioral effects, handles price endogeneity using a Bayesian analog of control functions (Park and Gupta 2009; Petrin and Train 2010), and accommodates purchase dynamics and the short and long-term effects of pricing using a multivariate non-homogeneous hidden Markov model (Montoya et al. 2010; Netzer et al. 2008; Schweidel et al. 2011).

We apply our framework using longitudinal transaction data from an aluminum retailer selling to industrial buyers. We identify two latent buying behavior states: a vigilant state characterized by high customer price sensitivity and a cautious approach towards ordering, and a relaxed state that is characterized by more direct orders and lower price sensitivity. We also find strong evidence for asymmetric reference price effects including loss aversion and gain seeking. Consistent with the hedonic adaptation theory (Frederick and Lowenstein 1999), we find that buyers not only weight price losses more than gains, but also take longer to adapt to losses than to gains. To the best of our knowledge, this is the first empirical evidence for the hedonic adaptation theory using secondary data. We also find that the proposed model exhibits superior out-of-sample predictive ability relative to several benchmark models.

We illustrate how our model can be used for computing optimal individual targeted prices that maximize long-term customer value. The optimal pricing policy balances between short- and long-term perspectives. In the short-term, the price is determined by the tradeoff between converting orders to sales and the desire to increase margins. In the long-term, the seller balances the wish to reduce prices to retain customers in the relaxed state and the wish to keep prices high to avoid reducing internal reference prices. The optimization indicates that the optimal dynamic targeted pricing policy can increase the seller's profitability by as much as 52% compared to the status-quo. We also use a counterfactual analysis to examine the firm's optimal pricing policy in the presence of a volatile aluminum commodity market. Volatility in the aluminum commodity market alters the cost structure for the seller while also changing the external reference point for buyers. Consistent with the dual entitlement principle, our simulation results indicate that when the commodity prices increase, the seller should pass to the buyers much of the cost increase. However, it can "hoard" some of the benefits of a cost decrease when the commodity market prices decrease.

In summary, our research pushes forward the pricing literature in several directions. On the methodological front, it provides a state-of-the-art and unique hierarchical Bayesian framework that weaves together a multivariate non-homogenous HMM, copulas, heterogeneity, and control functions to effectively capture relevant aspects of the B2B settings. More importantly, on the substantive front, it yields insights about how the short- and long-term effects of behavioral demand parameters such as loss aversion, reference price, and latent buying behavior states shape customer demands in what is traditionally considered to be “rational” purchasing activity. Furthermore, our findings provide strong evidence for the potential to employ value-based pricing policies, even in a traditional B2B industry characterized by cost-plus pricing practices.

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 metal retailer. Section 4 outlines the modeling framework. Section 5 illustrates the application of our modeling framework to the data. Section 6 describes the dynamic targeted pricing optimization based on the estimated model and Section 7 concludes by discussing practical implications, theoretical contributions, and future directions.

2. Targeted Pricing Decisions in B2B Settings

Our research lies at the intersection of multiple research streams involving pricing in B2B markets, targeted pricing, pricing dynamics and reference prices. In this section, we briefly review these literatures.

2.1 Pricing in B2B Markets

The majority of the research on B2B pricing is conceptual and survey based (Johnston and Lewin 1996), and scant attention is given to quantitative pricing models. The overall neglect could stem from conflicting views about the role and importance of pricing relative to other attributes in B2B contexts (see Lehmann and O’Shaughnessy 1974; Wilson 1994 and Hinterhuber 2004). In this paper, we empirically investigate the multifaceted impact of pricing using data from a B2B seller.

2.2 Dynamic and Individually Targeted Pricing

Targeting and customization are emerging topics in academic research and are immensely relevant to the world of practice. The empirical literature on targeting, however, has primarily focused on non-price marketing actions such as catalog mailing (Gönül and Shi 1998; Simester et al. 2006), coupons (Rossi et al. 1996; Shaffer and Zhang 1995), digital marketing campaigns

(Ansari and Mela 2003), B2B communication contacts (Venkatesan and Kumar 2004) and pharmaceutical detailing and sampling (Dong et al. 2009; Montoya et al. 2010) as these marketing actions are viewed as being naturally customizable and targetable at the individual level. In contrast, empirical research on *individually* targeted pricing has been relatively sparse, possibly due to the logistical, ethical, and legal issues concerning price discrimination in traditional (B2C) settings. Exceptions include Zhang and Krishnamurthi (2004) and Lewis (2005) who study pricing and promotion targeted at the segments level. Khan et al. (2009) demonstrate the importance of individually targeting consumers and the value of both inter-temporal and cross-sectional targeting for a mix of promotional activities. The authors note that one of the limitations of their study is the inability to observe purchase intents that did not result in a purchase. They also highlight the logistical problem of individually targeted pricing in the B2C brick and mortar context.

Dynamics can be incorporated in pricing models using different mechanisms. Winer (1986) uses reference prices to capture dynamics. Greenleaf (1995) and Kopalle et al. (1996) study the implications of reference prices for the firm's dynamic pricing policy, Lewis (2005) incorporates customers' forward looking expectations about future prices, and Chan and Seetharaman (2004) model the relationship between state-dependence and pricing decisions. We differ from these studies on several aspects. First, we use a multivariate non-homogenous HMM to allow for dynamic pricing and to capture the enduring impact of reference prices. Second, we leverage these dynamics to target prices both at the *individual level*, as well as *temporally,* over repeated transactions for the same customer. Finally, most pricing models investigate the effect of the firm’s pricing on brand choices or on a single purchase decision. In contrast, we investigate the impact of the seller's pricing policy and the resulting reference prices on a sequence of inter-related customer decisions that are typical of B2B environments.

2.3 Reference Prices in Customer Buying Behavior

The notion that consumers use reference prices in assessing the attractiveness of offers is well established within marketing (Hardie et al. 1993; Kalyanaram and Winer 1995; Kalwani et al.

1990; Krishnamurthi et al. 1992; Winer 1986). The literature distinguishes between “internal” and “external” reference prices (Mayhew and Winer 1992). External reference prices (e.g., MSRP, and prices of other brands) are generally observable and common to all customers, whereas internal reference prices are assumed to be individual-specific and are often constructed based on the customer’s observed prices on previous purchase occasions. Briesch et al. (1997) compares several external and internal reference price mechanisms and concludes that internal reference prices have greater empirical support.

A rich body of experimental and empirical work demonstrates the behavioral underpinnings of reference price effects (Kalawani et al. 1990; Wedel and Leeflang 1998). Erdem et al. (2010), in contrast, proposes a “rational” explanation for reference price effects based on quality signaling and price expectations. In behavioral pricing, the observed price in relation to a reference price is often encoded as either a loss or a gain and can thus have an asymmetric impact on brand choice (Kalwani et al. 1990; Putler 1992), purchase timing (Bell and Bucklin 1999), and purchase quantity (Krishnamurthi et al. 1992).

Despite the voluminous literature on reference prices, these have found little application in B2B pricing models based on the presumed "rationality" of B2B decision makers (Kalayanaram and Winer 1995). Internal reference prices, however, can play an important role in B2B purchasing because transactions are formally recorded by most buyers and relationships are typically long-term in nature. In addition, we know little about the impact of ignoring reference prices on targeting policy and performance as they have not been studied in conjunction with targeted pricing. We further the reference price literature along several dimensions by studying the asymmetric impacts of internal reference prices in multi-decisional contexts which permit first-degree and inter-temporal price discrimination. We also investigate the possible long-term effects of reference prices in a B2B context. Next, we describe our dataset and the business context in which the seller operates.

3. Data

Our data come from an East Coast local aluminum retailer that supplies to industrial clients such as machine shops and fabricators operating in its geographical trading area. This dataset is typical of what is found in B2B selling in commodity markets. The data contains customer-level information on purchase events over a 21 month period 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 may request a price quote (usually via the phone or fax). For example, a typical direct order may be received by the morning via a fax saying: "Send me four aluminum sheets A inch by B inch and thickness of C inch by tomorrow afternoon". Direct orders are generally fulfilled immediately and the customer is charged a price determined by the seller. Alternatively, if the customer requests a quote (i.e., an "indirect order") the firm bids for the customer's business, and can only "win" the business if the customer accepts the quoted price.[1] Thus, in our setting, purchase events include not only completed transactions but also lost transactions involving quotes that were not accepted, thus allowing for a better understanding of customer price sensitivity.

The company has a large number of 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 company does not maintain a price list and determines the price to charge or quote on a case by case basis. Because of the variation in order quantities and the large number of SKUs, a common measure for prices in this industry is a "price per pound" measure that incorporates the cutting costs and complexity of the order. As is typical of most customer relationship management (CRM) datasets in B2B settings, our dataset does not include

[1] Discussions with the management and sales personnel in the company that provided the data revealed that negotiations, beyond the request for the price quote and the firm initial bid, are not common. Analysis of the price and bid quantity relative to price and quantity on the order invoice confirmed that both price and quantity rarely changed from the original bid to the final order. This provides empirical evidence for minimal negotiation beyond the bid process.

information about the competition. However, unfulfilled indirect orders provide an indirect signal for a purchase that went to the competition.

Our sample contains 1,859 customers for whom we observe at least 7 purchase events (quotes or orders) in the data period (see Tables 1 and 2 for summary statistics of the data). On average, a customer 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 Shipley and Jobber (2001) who argue that in B2B settings, customers' needs are often urgent and account for a small proportion of the buyer's total expenses, leading to relatively low price sensitivity. Of the indirect orders, 47% were won by the seller.

An average purchase event involves 457 lbs 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) customers are heterogeneous in terms of their metal needs and transactions with the firm; 3) customers exhibit different propensities to order directly without asking for a price quote, suggesting variation in their attitudes and latent relationships with the firm; 4) about half of the orders are direct (which result in a sale regardless of the price charged). This suggests that the firm might be tempted to charge "any" price on such orders. However, such pricing behavior can have negative long-term consequences. We now describe our modeling framework that is designed to account for the above aspects of our dataset.

4. Model

We model a sequence of purchase events for each customer. A *purchase event* is characterized by four interrelated customer 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 request a quote, hence allowing the seller to bid for business, and 4) whether to accept the quote if a bid was made by the seller. We can write the vector of observed customer behaviors for customer 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 inter-purchase-event time), and b_{ij} and w_{ij} are the binary quote request and quote acceptance decisions, respectively. The seller observes the buying behavior and pricing history for each customer and the marketing environment before deciding on the unit price p_{ij} to charge customer *i* at purchase event *j*.

To model customer dynamics over repeated purchase events we allow the customer to transition between different latent buying behavior states that differentially impact the four customer decisions. The seller's past pricing decisions may affect the customer's transition between states. For example, a customer who was 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, coupled with a higher price sensitivity.

We capture such dynamics using a multivariate non-homogeneous HMM. In the HMM, the joint probability of a sequence of interrelated decisions up to purchase event *j*, for customer *i*, $\left\{\mathbf{Y_{i1}} = \mathbf{y_{i1}}, ..., \mathbf{Y_{ij}} = \mathbf{y_{ij}}\right\}$, is a function of three main components: (1) the initial hidden state membership probabilities $\boldsymbol{\pi_i}$ (2) a matrix of transition probabilities among the buying-behavior states $\mathbf{\Omega_{i,j\text{-}1\to j}}$, and (3) a multivariate likelihood of the interrelated customer decisions conditional on the customer's buying-behavior state $L_{ij|s} = f_{is}(q_{ij}, t_{ij}, b_{ij}, w_{ij})$. We describe our formulation of each of these components next.

4.1 Initial State Distribution

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

4.2 The Markov Chain Transition Matrix

We model the transitions between states as a Markov process. Each element of the transition matrix ($\mathbf{\Omega_{i,j\text{-}1\to j}}$) can be defined as $\omega_{ijss'} = P(S_{ij} = s' \mid S_{ij-1} = s)$, the conditional probability that customer *i* moves from state *s* at purchase event *j*-1 to state *s'* at purchase event *j*, and where

$0 \le \omega_{ijss'} \le 1 \;\; \forall s, s'$, and $\sum_{s'} \omega_{ijss'} = 1$. As the transition probabilities are influenced by the seller's pricing decisions at the previous purchase event *j*-1, we define:

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

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

4.3 The State Dependent Multivariate Interrelated Decisions

Conditional on being in state *s* at purchase event *j*, the customer makes the four interrelated decisions. These decisions are unconditionally interrelated as they all depend on the customer's latent state. To allow these to be conditionally dependent, we use a copula approach (Danaher and Smith 2011; Trivedi and Zimmer 2007). Copulas enable us to model each decision flexibly using appropriate marginal distributions while at the same time allowing for interdependence.

Given that customer *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 as:[2]

$$L_{ij|s} = f_{is}(q_{ij}, t_{ij}, b_{ij}, w_{ij}) = f_{is}(q_{ij}, t_{ij}) \Pr_{is}(b_{ij}, w_{ij} \mid q_{ij}, t_{ij}). \qquad (2)$$

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

$$\Pr_{is}(b_{ij}, w_{ij} \mid q_{ij}, t_{ij}) = \Pr_{is}(b_{ij} = 0 \mid q_{ij}, t_{ij})^{1-\delta_{ij}^{b}} \left[\Pr_{is}(w_{ij} \mid b_{ij} = 1, q_{ij}, t_{ij}) \Pr_{is}(b_{ij} = 1 \mid q_{ij}, t_{ij}) \right]^{\delta_{ij}^{b}}, \quad (3)$$

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

In modeling the time between purchase events, t_{ij}, the last observation for each customer, t_{ij}^{*}, is censored because of the fixed time horizon of the dataset. Let $S\left(t_{ij}^{*}\right)$ be the survival

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

function for the censored observation, and let δ_{ij}^{c} be a censoring indicator, which equals 1 if observation *j* for customer *i* is censored, and 0, otherwise. Accordingly, accounting for censoring and inserting Equation (3) into Equation (2) we can re-write Equation (2) as follows

$$L_{ij|s} = f_{is}(q_{ij}, t_{ij}, b_{ij}, w_{ij}) =$$

$$= S_{is}(t_{ij}^{*})^{\delta_{ij}^{c}} \left[f_{is}(q_{ij}, t_{ij}) \Pr_{is}(b_{ij} = 0 \mid q_{ij}, t_{ij})^{1-\delta_{ij}^{b}} \left[\Pr_{is}(w_{ij} \mid b_{ij} = 1, q_{ij}, t_{ij}) \Pr_{is}(b_{ij} = 1 \mid q_{ij}, t_{ij}) \right]^{\delta_{ij}^{b}} \right]^{1-\delta_{ij}^{c}} . \tag{4}$$

Next, we describe the distributional assumptions for each of the four decisions and provide details about the copula approach.

4.3.1 Modeling Quantity and Time between Events. We use a bivariate copula to model the joint density of quantity and time between events, $f(q_{ij}, t_{ij})$. The copula weaves together the univariate marginal distributions into a joint distribution. That is, the joint CDF of q_{ij} and t_{ij} is

$$F_{is}(q_{ij}, t_{ij}) = C(F_{isq}(q_{ij}), F_{ist}(t_{ij})) , \tag{5}$$

where, *C*( ) is a copula function, and F_{isq} and F_{ist} are the c.d.f's for the quantity and inter-purchase-event time variables, respectively. The joint density can then be written as

$$f_{is}(q_{ij}, t_{ij}) = f_{isq}(q_{ij}) f_{ist}(t_{ij}) H_{iqt}[F_{isq}(q_{ij}), F_{ist}(t_{ij})] , \tag{6}$$

where, f_{isq} and f_{ist} are the univariate marginal densities, and $H_{iqt}[.,.]$ is the double derivative of the copula function *C*( ) with respect to the two marginal c.d.f's. We condition on the state throughout to highlight that the distributions are state dependent.

While many different marginal distributions can be used, here, we assume that the purchase-event times follow a log-logistic distribution (Lancaster 1990, pp. 44; Kumar et al. 2008) because of its flexibility in accommodating both monotonic and non-monotonic hazards. The p.d.f and c.d.f of the log-logistic are given by

$$f_{is}(t_{ij}) = \frac{e^{\mathbf{x_{ij}}'\boldsymbol{\beta}_{\mathbf{tsi}} + \xi_{ij}^{t}} \alpha_s t_{ij}^{\alpha_s - 1}}{(1 + e^{\mathbf{x_{ij}}'\boldsymbol{\beta}_{\mathbf{tsi}} + \xi_{ij}^{t}} t_{ij}^{\alpha_s})^2} , \quad F_{is}(t_{ij}) = \frac{e^{\mathbf{x_{ij}}\boldsymbol{\beta}_{\mathbf{tsi}} + \xi_{ij}^{t}} t_{ij}^{\alpha_s}}{1 + e^{\mathbf{x_{ij}}\boldsymbol{\beta}_{\mathbf{tsi}} + \xi_{ij}^{t}} t_{ij}^{\alpha_s}} , \tag{7}$$

where, $\alpha_s > 0$, is a shape parameter, $\boldsymbol{\beta}_{\mathbf{tsi}}$ is vector of coefficients for customer-level, purchase-event specific covariates such as prices or reference prices, and ξ_{ij}^{t} represents an unobserved shock associated with the inter-purchase-event time. The random shock ξ_{ij}^{t} is assumed to be correlated with the unobserved shock of the pricing equation to account for possible endogeneity (see section 4.4).

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

$$f_{is}(q_{ij}) = \frac{\phi(\frac{\log q_{ij} - \mathbf{x_{ij}}'\boldsymbol{\beta}_{\mathbf{qsi}} - \xi_{ij}^{q}}{\sigma})}{\sigma q_{ij}}, \quad F_{is}(q_{ij}) = \Phi(\frac{\log(q_{ij}) - \mathbf{x_{ij}}'\boldsymbol{\beta}_{\mathbf{qsi}} - \xi_{ij}^{q}}{\sigma}), \tag{8}$$

where, $\boldsymbol{\beta}_{\mathbf{qsi}}$ is a vector of coefficients for a set of customer-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 that is 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.

There are many families of copulas which differ in the nature of the dependence they represent. We use the Frank copula (Frank 1979; Trivedi and Zimmer 2007) to model the interdependence between quantity and inter-purchase-event times because it covers the entire domain between the Frechet-Hoeffding bounds and thus allows for both positive and negative interdependence. The Frank copula for quantity and inter-purchase-event time can be written as

$$C[F_{isq}(q_{ij}), F_{ist}(t_{ij})] = -\frac{1}{\theta_{qt}} \ln\left(1 + \frac{(e^{-\theta_{qt}F_{isq}(q_{ij})} - 1)(e^{-\theta_{qt}F_{ist}(t_{ij})} - 1)}{e^{-\theta_{qt}} - 1}\right), \tag{9}$$

where, the parameter θ_{qt} captures the interdependence between the quantity and inter-purchase-event time decisions. For $\theta_{qt} > 0$ the interdependence between quantity and inter-purchase-event time is positive, and for $\theta_{qt} < 0$ the interdependence is negative.

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

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

Similarly, conditional on a price quote, customer *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:

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

The joint distribution of b_{ij} and w_{ij} is relevant when a bid is made (i.e., b_{ij} = 1). For such quote orders, we can distinguish between those that were accepted (the seller "wins" the bid; $w_{ij} = 1$) and those that were rejected (the seller "loses" the bid; $w_{ij} = 0$). For accepted orders, we have

$$\Pr_{is}(w_{ij} = 1, b_{ij} = 1 \mid t_{ij}, q_{ij}) = 1 - \Pr_{is}(b_{ij}^* < 0) - \Pr_{is}(w_{ij}^* < 0) + \Pr_{is}(b_{ij}^* < 0, w_{ij}^* < 0)\,. \tag{10}$$

Similarly, for rejected bids, we have

$$\Pr_{is}(w_{ij} = 0, b_{ij} = 1 \mid t_{ij}, q_{ij}) = \Pr_{is}(w_{ij}^* < 0) - \Pr_{is}(b_{ij}^* < 0, w_{ij}^* < 0)\,. \tag{11}$$

We model the joint probability $\Pr_s(b_{ij}^* < 0, w_{ij}^* < 0)$ in Equations (10) and (11) using the Frank copula such that,

$$\Pr_{is}(b_{ij}^* < 0, w_{ij}^* < 0) = C(F_{isb^*}(0), F_{isw^*}(0)) = -\frac{1}{\theta_{bw}} \ln\left(1 + \frac{(e^{-\theta_{bw} F_{isb^*}(0)} - 1)(e^{-\theta_{bw} F_{isw^*}(0)} - 1)}{e^{-\theta_{bw}} - 1}\right), \tag{12}$$

where, the parameter θ_{bw} captures the interdependence between the quote request behavior and quote acceptance decision. For the marginal distributions, we assume that each of the latent variables, b_{ij}^* and w_{ij}^* are distributed logistic. Thus,

$$\Pr_{is}(b_{ij}^* < 0) = \frac{1}{1 + e^{\mathbf{x_{ij}\beta_{bsi}} + \xi_{ij}^b}} \quad \text{and} \quad \Pr_{is}(w_{ij}^* < 0) = \frac{1}{1 + e^{\mathbf{x_{ij}\beta_{wsi}} + \xi_{ij}^w}}\,. \tag{13}$$

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

Inserting equations (6)-(13) into equation (4) we get the full likelihood of observing the four interrelated customer decisions conditional on the customer's state.

4.4 The HMM Likelihood Function

Given the Markovian structure of the model, the likelihood of observing a set of joint customer decisions at purchase event J is dependent on all decisions up to that event. The likelihood of observing the customer's decisions over J purchase events ($\mathbf{Y_{i1}}, \mathbf{Y_{i2}}, ..., \mathbf{Y_{iJ}}$) can be succinctly written as (McDonald and Zucchini 1997)

$$L_{iJ} = P(\mathbf{Y_{i1}} = \mathbf{y_{i1}}, ..., \mathbf{Y_{iJ}} = \mathbf{y_{iJ}}) = \boldsymbol{\pi}_\mathbf{i} \mathbf{M_{i1}} \boldsymbol{\Omega}_{\mathbf{i},1\to\mathbf{2}} \mathbf{M_{i2}} ... \boldsymbol{\Omega}_{\mathbf{i},\mathbf{J}\text{-}\mathbf{1}\to\mathbf{T}} \mathbf{M_{iJ}} \mathbf{1}', \tag{14}$$

where, $\boldsymbol{\pi}_\mathbf{i}$ is the initial state distribution described in Section 4.1, $\boldsymbol{\Omega}_{\mathbf{i},\mathbf{j}\to\mathbf{j+1}}$ is the transition matrix described in section 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.

To ensure identification of the states, we restrict the probability of quote request to be non-decreasing in the buying behavior states. We impose the restriction $\tilde{\beta}_{b01i} \leq \tilde{\beta}_{b02i} \leq ... \leq \tilde{\beta}_{b0Si}$ by setting $\tilde{\beta}_{b0si} = \beta_{b01i} + \sum_{s'=2}^{s} \exp(\beta_{b0s'i})$; s=2,…,$S$. 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 $\mathbf{x_{ij}}$. To avoid underflow, we scale the likelihood function in Equation (14) following the approach suggested by MacDonald and Zucchini (1997, p. 79).

4.5 Recovering the State Membership Distribution

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

$$P(S_{ij} = s \mid \mathbf{Y_{i1}}, \mathbf{Y_{i2}}, .., \mathbf{Y_{ij}}) = \boldsymbol{\pi}_\mathbf{i} \mathbf{M_{i1}} \boldsymbol{\Omega}_{\mathbf{i},1\to\mathbf{2}} \mathbf{M_{i2}} ... \boldsymbol{\Omega}_{\mathbf{i},\mathbf{j}\text{-}\mathbf{1}\to\mathbf{t}\bullet\mathbf{s}} L_{ij|s} / L_{ij}, \tag{15}$$

where, $\boldsymbol{\Omega}_{\mathbf{i,j\text{-}1\to j.s}}$ is the s^{th} column of the transition matrix $\boldsymbol{\Omega}_{\mathbf{i,j\text{-}1\to j}}$, and, L_{ij} is the likelihood of the observed sequence of joint decisions up to purchase event j from Equation (14).

4.6 The Control Function Approach to Price Endogeneity

It is possible that the seller's pricing decisions are based on unobserved factors that also impact the buyers' decisions. For example, the seller can target each customer differentially based on private information it has that is not observed by the researcher. In such case, price will be correlated with the unobserved components (the $\xi's$) of the four distributions. Ignoring this endogeneity can result in misleading inferences about the price sensitivities of customers (Villas-Boas and Winer 1999). We use a Bayesian analog of the control function approach to account for price endogeneity (Park and Gupta 2009; Rossi et al. 2005). Specifically, 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 impacting the four decisions. Formally, we have

$$p_{ij} = \gamma_1 + \gamma_2 z_{ij} + \mu_{ij}, \tag{16}$$

where, μ_{ij} are unobserved factors influencing the pricing decision. We assume that μ_{ij} is distributed jointly bivariate normal with each of $\xi_{ij}^l, l \in \{t, q, b, w\}$ in Equations (7), (8), and (13) to account for endogeneity. For each of the four decisions, the bivariate normal distribution can be written as

$$f(\mu_{ij}, \xi_{ij}^l) \sim MVN\left(\begin{pmatrix}0\\0\end{pmatrix}, \begin{pmatrix}\sigma_p^2 & \rho_{pl}\\ \rho_{pl} & \sigma_l^2\end{pmatrix}\right) \quad l \in \{t, q, b, w\}, \tag{17}$$

where, σ_p^2 is the variance of μ_{ij}, σ_l^2 is the variance for the random shock ξ_{ij}^l, and $\rho_{p,l}$ is the covariance between μ_{ij} and ξ_{ij}^l. We use the wholesale price as an instrument for price (z_{ij} in Equation (16)). The wholesale price is the seller's wholesale cost, that is, the cost that the seller pays to the mills for the metal. This is observed by the seller but not by the buyers. Conversations with the management team indicate that the salespersons observe the wholesale cost on their computer screen and rely heavily on it when setting the price. Wholesale price has been commonly used as an instrument for price (e.g., Chintagunta 2002).

5. Model Estimation and Results

In the previous section, we focused on the overall modeling framework. We now describe how this is instantiated in our application. We also report and discuss the parameter estimates from our model.

5.1 Description of Variables

Asymmetric Reference Price Effects **-** We define the reference price for customer *i,* at purchase event *j,* as a quantity weighted average of the customer's past observed prices (in \$/lb)[3], i.e.,

$$reference\_price_{ij} = \left(\sum_{k=1}^{j-1} quantity_{ik} \times price\_lb_{ik} \right) / \left(\sum_{k=1}^{j-1} quantity_{ik} \right),$$

where, $price\_lb_{ij}$ is the price per pound observed by customer *i* in purchase event *j.* The quantity weighting reflects that customers attend closely to larger orders relative to smaller ones. Consistent with the rich literature on behavioral pricing, we incorporate asymmetric reference price effects (Putler 1992) using "gain" and "loss" variables:

$$gain_{ij} = \begin{cases} reference\_price_{ij} - price\_lb_{ij}, & \text{If } price\_lb_{ij} < reference\_price_{ij}, \\ 0, & \text{Otherwise.} \end{cases} \quad \text{and}$$

$$loss_{ij} = \begin{cases} price\_lb_{ij} - reference\_price_{ij}, & \text{If } price\_lb_{ij} > reference\_price_{ij}, \\ 0, & \text{Otherwise.} \end{cases}$$

Commodity market **-** We use the daily aluminum spot prices on the London Metal Exchange (LME) to control for the fluctuations in commodity market. The LME also serves as an external reference price. We define,

lme_{ij} = the aluminum spot price (dollars per metric ton) on the LME at purchase event *j*, for customer *i*, and

[3] We tested several alternative specifications of the reference price variable including time-weighted reference prices, and the external reference price from the commodity market. The quantity weighted reference price resulted in the best model fit. Incorporating time decay to the reference price formulation did not result in significant improvement in model's fit.

$lme\_volatility_{ij}$ = the volatility of the aluminum spot price on the LME, calculated as the standard deviation of the LME daily returns over the seven trading days prior to purchase event *j* for customer *i*.

We now describe how these variables are included in each of the customer 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 (7) 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 (8) includes the covariates: $gain_{ij-1}$, and $loss_{ij-1}$, lme_{ij}, and $lme\_volatility_{ij}$.

3. *Quote Request*: We expect that customers would 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, the internal reference price effects experienced on the previous purchase event could impacts the decision to request a quote. Thus, $\mathbf{x_{ij}}$ in Equation (13) includes: t_{ij}, $quantity_{ij}$, $gain_{ij-1}$, $loss_{ij-1}$, lme_{ij}, and $lme\_volatility_{ij}$

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 customer experiences a price gain. 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 (13) includes: t_{ij}, $quantity_{ij}$, $gain_{ij}$, $loss_{ij}$, $gain_{ij} \times quantity_{ij}$, $loss_{ij} \times quantity_{ij}$, lme_{ij}, and $lme\_volatility_{ij}$.

5.1.2 The Non-Homogenous Transition Matrix - The "gain" or "loss" experienced on the previous purchase event can affect the buyer's evaluation (or re-evaluation) of the relationship with the seller and trigger a transition to a different buying behavior state thereby affecting purchases in the long-run. Thus, $\mathbf{x_{ij-1}}$ in Equation (1) includes: $gain_{ij-1}$, and $loss_{ij-1}$.

5.2 Heterogeneity Specification

We allow the coefficients for the reference price effects and the parameters that govern the HMM dynamics to vary across customers. This facilitates targeting, properly accounts for reference price effects and is also crucial for empirically distinguishing dynamics from cross-customer heterogeneity (Heckman 1981). Specifically, we allow all the transition matrix parameters, as well as the four intercepts and the coefficients of the reference price related variables, (i.e., $gain_{ij}$, $gain_{ij-1}$, $loss_{ij}$, $loss_{ij-1}$, $quantity_{ij} \times gain_{ij}$, $quantity_{ij} \times loss_{ij}$, lme_{ij}, $lme\_volatility_{ij}$) to vary across customers. For parsimony, all other parameters in the model are assumed to be invariant across customers (see column 3 in Table 4 for a full list of the random- and fixed-effect parameters). We also restrict the Frank copula dependence parameters (θ_{qt} and θ_{wb}) and the scale parameter of the lognormal distribution (σ) to be the same across the latent states.

5.3 Estimation Procedure

We use a hierarchical Bayesian approach based on Markov Chain Monte Carlo (MCMC) methods for inference. In particular, we follow Netzer et al. (2008) and use the HMM likelihood function in Equation (14) in conjunction with the pricing equation to create a joint likelihood function. The inherent complexity of the HMM often leads to significant autocorrelation among the draws. We therefore use the adaptive Metropolis procedure in Atchadé (2006) which significantly improves 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. 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 draws from the full conditional distributions.

5.4 Model Fit and Predictive Ability

We compare the fit and predictive ability of the proposed model (**Full** model) to that of four benchmark models which vary with respect to a) the extent of interdependence among the four decisions, b) the extent of heterogeneity, c) the degree of dynamics, and d) whether they account for price endogeneity. All our benchmark models rely on the same marginal distributions for the four decisions as the proposed model. For all models, we use a Bayesian approach for inference. The benchmark models are:

1) Nohet: This model ignores the heterogeneity in the Full model and thus differs from it in that all parameters are assumed to be invariant across customers. A comparison of this model with the Full model therefore allows us to assess the importance of modeling heterogeneity.

2**) Indep:** This is a single state heterogeneous model in which the distributions for the four decisions are independent. Thus, this model ignores the interdependence that is created from the copulas and the HMM. A comparison with the Full model can indicate the value of modeling the decisions jointly and of capturing relationship dynamics.

3) NoDyn: In this model we ignore the two sources of dynamics present in the Full model: 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. Comparing this "static" model to the proposed model allows us to assess the value of capturing dynamics.

4) NoEndo: This model is identical to the Full model except that prices are considered exogenous. Thus, this model does not have the Bayesian control function component. A comparison of this model to the Full model can highlight the extent of price endogeneity as well as the perils of ignoring it**.**

We compare the fit and predictive ability of the five models using the Log Marginal Density (LMD) and the Deviance Information Criterion (DIC) on the calibration sample, and the validation log-likelihood on the validation sample. We also assess the component-specific fit and prediction using the root mean squared error (RMSE) and the mean absolute deviation (MAD) between the predicted and observed values of the four outcome variables and the hit-rates for the binary quote-request and conditional quote acceptance decisions within the calibration and validation samples.

Table 3 presents the model comparison statistics for the five models. We see that the **Full** model outperforms the benchmark models on both the component-specific and the overall measures for both samples. The relatively poor performance of the no-heterogeneity model indicates that the extent of buyer heterogeneity in our data is substantial and suggests an opportunity for individually targeted pricing. Similarly, accounting for dynamics improves the representation and prediction of buying behavior, suggesting that inter-temporal targeting may be advantageous. In contrast, accounting for price endogeneity results in only a marginal improvement. This is consistent with the use of a "cost-plus" pricing strategy by the seller. A regression of price on wholesale cost yields an R-squared of 0.84. We now discuss the parameter estimates of the proposed model.

5.5 Parameter Estimates

In this section we discuss the parameter estimates of the **Full** model with two latent states. Models with a larger number of HMM states resulted in worse performance (higher DIC).

5.5.1 The HMM States. Table 4 contains the posterior means, standard deviations and the 95% posterior intervals for the **Full** model. Recall that all covariates are mean-centered, thus, the intercept of each equation captures the average response tendency. A comparison of the parameters across the two states indicates that customers in State 1 are more likely to request a price quote, but are less likely to accept the quote relative to customers in State 2. Also, customers in State 2 are more sensitive to reference price effects. These customers exhibit stronger loss aversion in the inter-purchase-event time, quote request and quote acceptance decisions and stronger gain seeking in the quantity decision. Customers who are in State 2 at a given purchase event also tend to be more responsive to the commodity market (LME), and pay closer attention to the characteristics of the order itself (quantity and inter-purchase-event time) in making their decisions.

Overall, this multidimensional view of the two states (see a summary of the two states in Table 5) implies that State 2 characterizes a more cautious approach towards ordering, whereas, customers 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*. Customer trust, which is an important facet in B2B relationships (Morgan and Hunt 1994), is the most likely underlying mechanism that governs the existence of these two states. Using the perspective of regulatory-focus theory (Higgins 1997), we can characterize customers in the vigilant state as more "prevention focused," i.e., focusing on avoiding price losses, and those in the relaxed state as being more "promotion-focused", i.e., concentrating more on gains and price savings.

5.5.2 Customer Dynamics. Customers can transition between the two states over time, thus creating long-term dynamics. The parameter estimates in Table 6 and their transition matrix representation in Table 7 together 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 price). One can see that the states are relatively sticky. The likelihood of remaining in the vigilant state is 92%. However, the seller can use its pricing policy to affect customer 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 customers perceive that they are treated well, they are more likely to transition to or remain in a state of more favorable relationship with the firm. In contrast, the right matrix in Table 7 shows that a 10% price increase ("loss") in the previous purchase event increases the likelihood that the customer will transition to the vigilant state by almost 7% and also increases the likelihood of staying in the vigilant state by 3.5%, which in turn, translates to further increased sensitivity to losses. Thus, a price increase may have a long-term effect by transitioning the customer 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 a separate transition matrix for each buyer. We leverage this heterogeneity in developing a targeted pricing policy in Section 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. We now look at the extent of interdependence among the decisions.

5.5.3 Interdependent Decisions. The parameters of the Frank copula in Table 6 indicate significant and substantial positive dependence between the quantity and inter-purchase-event time decisions, and negative dependence between the quote request and quote acceptance decisions. The latter negative correlation indicates selectivity (Heckman 1981) and implies that factors that increase the likelihood of quote requests also result in lower quote acceptance rates, a kind of a "double jeopardy" for the firm. The positive correlation between the inter-purchase-event time and quantity decisions can be partly explained by the presence of inventories. Finally, the shape parameters for the log-logistic distribution of the inter-purchase-event times in Table 6 indicate that the baseline hazard for a purchase first increases and then decreases with time, irrespective of the latent state.

5.5.4 Investigating Price Endogeneity. Table 8 compares the price sensitivity estimates for the **Full** model with those from the **NoEndo** model to assess the benefits of accounting for price endogeneity. The bottom portion of Table 8 shows the correlations between the unobserved error in the price equation and each of the random shocks (ξ's) for the four decisions. The results show mildly negative correlations for the quantity and quote acceptance decisions and a mildly positive correlation for the inter-purchase-event time and quote request decisions. These correlations are all in the expected direction and can imply that 1) despite the predominant practice of cost-plus pricing, the seller does occasionally target price-insensitive (sensitive) customers by charging them higher (lower) prices, and that 2) unobserved shocks could influence both the seller's pricing decisions and the buyers' behavior. Hence, a failure to properly account for price endogeneity would result in overestimation of the effects of price gains, and underestimation of the impact of price losses (see Table 8). However, it should be noted that overall the differences in the price sensitivity estimates between the two models are relatively small. Moreover, the model fit statistics for the **Full** model and Model 4 in Table 3 also suggest that the gains from modeling price endogeneity are modest in our application.

5.6 Disentangling the Effects of Pricing

Assessing the marginal and the integrative impact of price is not straight forward from the reference price coefficients in Tables 4 and 6 as 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 impact. Following the one-time price shock, we set prices to the reference price level for the remaining 19 purchase events. The short-term elasticity captures the immediate impact (i.e., on purchase event 1), whereas, the long-term elasticity is computed using the effect over the next 19 purchase events. These short- and long-term elasticities are reported in Table 9. In addition, Figures 1 and 2 illustrate the asymmetric percentage increases and decreases in each affected variable over the simulated 20 periods following the one-time 10% price increase or decrease.

Several insights can be derived from Table 9 and the two figures. 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%-50% of the total price elasticities. The average short-term quantity elasticities are consistent with those reported in Jedidi et al. (1999), and fall within the range of the price elasticities reported in the meta-analysis by Tellis (1988) and Bijmolt et al. (2005). 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 and loss aversion (Kahneman and Tversky 1979), we find that price increases (losses) have a stronger impact than price decreases (gains) for all decisions, except quantity. The quantity decision, in contrast, exhibits gain seeking. Thus, in the context of the B2B retailer studied here, behavioral reference pricing effects are present and significant. Fifth, Figures 1 and 2 show that the negative effects of a price hike (loss) persist longer than the positive effects of price drop (gain). This result is consistent with the *hedonic adaptation theory* which states that individuals adapt faster to improvements than to deteriorations (Frederick and Lowenstein 1999). To the best of our knowledge, this is the first investigation of the long-term effects of asymmetric reference price effects, and the first empirical demonstration that the effect of price losses is not only stronger in magnitude than the effect of price gains, but also lasts longer.

Overall, these results imply that in B2B contexts, studies that consider only the short-term effects of pricing can significantly and substantially underestimate the overall impact of pricing. Furthermore, studies that do not examine the asymmetry in reference price effects caused by gains and losses can miss important insights into how B2B customers behave, both in the short and long run. From a managerial perspective, B2B sellers in general need to be cautious in increasing prices because of the immediate as well as the persistent negative impact of loss aversion and the protracted adaptation. In the next section, we investigate such issues and discuss how to leverage these behavioral insights to dynamically target individual customers using the model's parameter estimates.

6. Targeted and Dynamic Price Policy Optimization

The separate and varied influences of price on the four customer decisions can be integrated using the metric of Long-term Customer Value (LCV; Gupta and Lehmann 2002) which captures the overall long-term impact of pricing. The LCV for customer i is the sum of the discounted stream of future profits. Specifically, customer value over the next J_i purchase events can be written as:

$$LCV_i = \sum_{j=1}^{J_i} \frac{(\text{Price}_{ij} - \text{Cost}_{ij}) q_{ij} \left[\left(1 - \delta_{ij}^b\right) + \delta_{ij}^b \delta_{ij}^w \right]}{(1+r)^{\tau_{ij}}} \tag{18}$$

where, δ_{ij}^b is a binary indicator that equals 1 when a quote is requested on purchase event j for customer i, and is zero for a direct order, δ_{ij}^w is another binary indicator that equals 1 when the quote is accepted and is zero otherwise, r represents the discount rate, and τ_{ij} is the cumulative time until the j^{th} purchase event. The seller's objective is to design a targeted dynamic pricing policy for each customer to maximize the LCV.

We conduct the optimization over a nine-month horizon. We therefore split the data into a calibration sample covering the duration from January to December 2007 and a holdout period ranging from January 2008 to September 2008.[4] We then use the parameter estimates from the calibration dataset to conduct the price optimization over the holdout period of 9 months. The optimization is performed for a representative sample of 300 customers who experienced between 6 and 16 purchase events over the calibration period and an average of ten purchase events in the nine-month planning horizon.

We discretize the continuous pricing decision on any purchase event to a set of five customer-specific price points that form the quintiles of the range of prices that the customer experienced in the calibration period. We choose to stay within each customer's historical range of experienced prices to avoid a drastic change in the price regime while accounting for the price variance experienced by each customer. 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 purchase events of the customer in the nine-month planning horizon. The optimization process is initialized for each customer by setting the state membership probabilities and the reference price to their values at the end of the calibration period. The forward simulation proceeds by generating a sequence of purchase events, characterized by a new set of quantity, inter-event time, quote request decisions and the associated reference price and latent state-membership

[4] For the purpose of price policy simulation, we re-estimate the proposed full 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, reported in Table 4.

probabilities. At each simulated purchase event, profits are computed by weighting the HMM latent state-specific profits by the state membership probabilities. For each customer we simulated 200 random sequences of purchase events over all possible price paths. Given the 5 price points at each purchase event, there are $5^{J_{ik}}$ possible price paths for customer i, where J_{ik} is the number of purchase events for customer i in the k^{th} random sequence. We then use the 200 Monte Carlo sequences to calculate the LCV by computing for each price path the average net present profits over the nine months, assuming a discount rate of 12% (a weekly discount rate r of 0.22%).[5] This yields an optimal dynamic price policy for each customer. Full details of the price simulation and optimization are available in the Appendix.

6.1 Price Policy Simulation Results

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

(1) Individually targeted static pricing policy – In this policy a single price is determined for each buyer for the entire nine-month planning horizon. This 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 that the buyer belong to on that occasion.[6] This policy accounts for dynamics arising from the latent transitions but enables limited targeting as only two possible prices are used.

(3) Segment targeted static policy – Analogous to the segment targeted dynamic policy, this policy includes only two possible prices. However, unlike in the dynamic segment policy, in this policy the buyer receives the same price for the entire planning horizon. This price is

[5] We tested the improvement in precision that can be gained from increasing the number of draws. We find diminishing marginal gains in precision with number of draws. We choose 200 draws as a good compromise between precision and computational time, as it offers 8% improvement in profits over 100 draws, but only underperforms 300 draws by 1%.

[6] In the segment policies the latent state membership at each purchase event is determined by the state with the highest membership probability.

based on the buyer's state membership in the first period of the planning horizon. Thus, this policy accounts for limited heterogeneity and ignores the possibility of transitions between the HMM states.

(4) 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.

(5) 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, policies (1) to (4) all use LCV maximization as the criterion when deciding which prices to charge, whereas, the myopic dynamic policy considers only the short-term effect of pricing in each period.

(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 stem from individual-level targeting, dynamic pricing, and from adopting a long-term perspective. Table 10 shows how the seven price policies perform. The proposed policy yields the highest LCV per customer of $4,795 over nine months. It outperforms the single static price policy by as much as 73%. The individually targeted static policy, which leverages heterogeneity, generates a 40% improvement in the LCV when compared to the aggregate static policy ($3,855 vs. $2,759). An additional 24% improvement results from dynamic targeting ($4,795 vs. $3,855). This result is consistent with the findings of Khan et al. (2009), which highlights the potential gains from inter-temporal targeting. Similarly, by leveraging the dynamic impact of pricing and by adopting a long-term perspective, the proposed policy improves profits by 17% over the myopic policy.

The firm's current policy is better than the static policies and yields profits similar to that of the segment dynamic policy. However, the current policy is worse than the individually targeted and dynamic policies. Finally, it is important to note that the proposed policy suggests a 52% improvement in profitability over the status-quo. Taking into account the entire customer base of the seller, this translates into a potential profit improvement of approximately $4 million annually.

Figure 3 shows the dynamic impact of different pricing policies by plotting the average monthly profits over the nine-month planning horizon. Several interesting insights can be gleaned from the figure. First, all the individual-level policies (static, dynamic or myopic) outperform the current policy, highlighting the rewards from individual-level targeted pricing. Second, the proposed policy generally dominates the static policy, showing the impact of temporal variation in pricing. Third, the proposed policy carefully balances the interplay between several forces that govern customers' buying behaviors, namely, 1) charging lower prices to increase quote acceptance; 2) charging lower prices to keep customers in the relaxed state or to transition them towards it 3) charging higher prices to increase margins, and 4) charging higher prices to keep customers' 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 10). This strategy leads to higher profits than the proposed policy in the first few months. However, charging lower prices results in lower reference prices, making it more likely that future prices are perceived as losses by customers. This subsequently increases the likelihood of customer migration to the vigilant state. The lowered reference prices create 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, thus demonstrating the importance of a long-term perspective when setting prices. Thus, Figure 3 suggests that in the world of B2B, where long-term relationship are crucial, myopia in price setting can be a slippery slope.

6.2 Pricing to Customers in the Vigilant and Relaxed States

Recall that Table 5 highlights the differences in customer behavior in the vigilant and relaxed states. We now examine how these parameter differences translate into differences in the targeted pricing policy and the LCV for these states. To do so, we assign the customer to the state with the highest posterior probability (see Equation 15). The resulting distributions of the state-specific LCVs are shown in Figure 4. Customers who are in the vigilant state have an average LCV of $3,551 relative to an average LCV of $6,137 for those on relaxed state. Figure 4 illustrates that the stark contrast in profitability across the two states results in a bimodal pattern of the overall LCV distribution.

To better understand how the proposed pricing strategy leverages the state membership to drive profitability, we plot in Figure 5 the monthly LCV values in conjunction with the proportion of customers who are in relaxed state in a given month. The figure shows that the pricing policy effects a transition from the vigilant state to the relaxed state over the planning horizon. This shift is most pronounced in the first few months. Once customers transition to the relatively sticky relaxed state after the first few months, they are likely to stay there and contribute much higher profits and LCV in the remaining months.

Figure 6 shows how the optimal pricing policy differs across the two states. The proposed price policy recommends a much higher price for customers who are in the relaxed state versus those that are in the vigilant state ($3.65/lb vs. $3.22/lb). The current policy that is used by the seller, however, does not price differentiate between customers in the two latent states. Charging a higher price for customers 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 customers 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 contrast, charging lower prices for customers who are in the vigilant state can help transition them to the relaxed state.

In summary, the superior profitability of our individual dynamic targeted price policy, relative to the current policy stems from its ability to leverage the 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) tradeoff between the short and long-run impacts of pricing arising from the different sources of dynamics in customer behavior.

6.3 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 look at how exogeneous fluctuations in the raw aluminum prices on the LME would impact optimal pricing decisions.

The aluminum prices on the LME fluctuated between US$2,393 to US$3,318 per tonne over the duration of our data. Much of this fluctuation occurred from August 2007 until January 2008 after an initial period of price stability in the first half of 2007. A change in the commodity prices can lead to at least two opposing impacts on the profitability of the seller. On the one hand, alumnium prices influence the seller's cost of replensihment[7]. On the other hand, buyers use the price on the LME as an external reference price. Thus, rising alumnimum prices, for example, increase the seller's future costs, but at same time serve as a high reference price for buyers, influencing them to buy larger quantities, order directly, or accept price quotes more readily. 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 empirically by computing our individual dynamic price policy 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. We then compare the resulting prices and profits to those at the original LME levels.[8]

Figure 7 shows the distribution of the optimal prices under each LME regime. For a 20% increase in the LME prices, the mean optimal price is $3.70 – a 9% price increase relative to the optimal price under the original LME levels ($3.39). However, when the LME drops by 20%, the optimal policy suggests a mean price of $3.30, which is a meager 2.7% price reduction. Figure 8 shows additional implications for the customers in the relaxed and vigilant states. We see that when the LME price increases by 20%, it is optimal for the seller to pass through this increase in

---

[7] The seller in our empirical application keeps up to 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 replinshment.
[8] Our use of a 20% shock in LME falls within the range of fluctuations observed during the data period. The maximum daily, monthly and 3-monthly LME fluctuations in our data were 5.6%, 22%, and 31%, respectively.

replacement cost by increasing the unit prices by 12% for customers who are in the vigilant state. In comparison, a modest increase of 4.6% is optimal for customers in the relaxed state. This result is consistent with the higher sensitivity to changes in the LME prices for the customers who are in the vigilant state, relative to the customers who are in the relaxed state (see Table 4). Figure 8 also shows that when LME prices drop by 20%, it is optimal for the firm to "hoard" most of the cost saving and drop prices by only 2.5%-2.8% for customers in both states. The rationale here is that lowering the price results in a corresponding lowering of the internal reference price, and this can have long terms 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* (Akerloff 1979, Kahneman et al. 1986, Okun 1981, Urbany et al. 1989). The dual entitlement principle states that 1) customers use reference cues when making transactions; 2) firms are "justified", in the eyes of customers, to increase prices when costs increase, in order to protect firms' normal profits; and 3) firms do not need to lower prices when costs drop, because customers perceptions are mainly driven by their past reference prices. While we do not model fairness directly, the external and internal asymmetric reference prices capture similar effects. To the best of our knowledge, this is the first paper to empirically demonstrate the dual-entitlement principle and measure its extent in a B2B transaction setting.

How does the above pricing strategy translate into long-term profitability? Figure 9 shows the LCV under each LME regime and broken down by each state. We find that when the LME increases, the firm is "entitled" to keep the normal profits it had prior to the price increase (the LCV drops only slightly from $4,795 to $4,701). However, the optimal price response to a drop in LME prices increases the LCV by approximately 7%, from $4,795 to $5,103.

In summary, to profitably manage fluctuations in economic conditions, sellers should pass on the most of the cost increases, especially to buyers who are in the vigilant state. In contrast, cost decreases present a good opportunity for sellers to enjoy a period of heightened profitability by keeping prices the same, at least in the short-run.

7. General Discussion

Understanding and managing the impact of pricing on buyer behavior is critical for firms' profitability in the long-run. The ability to customize prices for each customer at each transaction allows firms to extract additional profits from the customer base. To ensure long-term profitability from each customer, firms need to model the multifaceted nature of the customer buying process taking into account preference heterogeneity and response dynamics. Furthermore, firms need to appropriately factor in both the short-term and long-term impact of pricing when setting prices to maximize long-term customer value.

In this paper, we present an integrative framework to model the customer-level and long-term impact of pricing decisions in B2B settings. Our framework jointly models the interrelated customer decisions via copulas, captures the response dynamics and the long-term effect of pricing decisions via a non-homogeneous HMM, and accounts for potential price endogeneity in a Bayesian control function framework. To the best of our knowledge this is the first paper to incorporate copulas as well as price endogeneity within a HMM.

We identify several substantive insights in the under-researched area of B2B pricing. First, we uncover how industrial customers may transition over time between a relaxed and a vigilant state of buying behavior. Second, we find that pricing in B2B contexts significantly influences customer decisions not only in the short-run but also in the long-term. Thus, B2B pricing models that do not account for the long-term effects may severely underestimate the impact of pricing. Third, despite the view that B2B customers are rational, we find significant asymmetric reference effects. Fourth, we find that the impact of a price "loss" is not only stronger than that of a price "gain", but that it also takes much longer for customers to adapt to losses. Therefore, sellers should be cautious of the potential long-term impact of either lowering the reference prices of buyers or of charging a higher price on any particular order.

Using a series of price policy simulations, we demonstrate that the proposed dynamic targeted price policy can offer up to 52% improvement in long-term customer profitability over the firm's current pricing policy. Further, we show that the LCV of a customer in the relaxed state is almost twice as high as that of customers in the vigilant state. The proposed policy balances two forces 1) lowering prices to win business and to transition or keep customers in the relaxed state, and 2) increasing prices to maximize margins and thereby avoid a decrease in the internal reference price.

We use a what-if simulation to investigate what the optimal pricing policy suggests in a volatile economic environment with cost fluctuations. We find that the seller should pass on a part of the increased costs to customers, but in contrast, 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. These policy implications are consistent with the dual-entitlement principle.

More generally, this research offers B2B sellers a comprehensive modeling framework to manage their customer base via 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 using a value-based pricing policy, taking advantage of the flexibility to implement first-degree inter-temporal price discrimination.

We highlight several limitations and directions for future research. First, we assume that customers are not forward looking with respect to the seller's pricing decisions. One could extend our framework to incorporate customer's expectation about future price changes (e.g., Lewis 2005). Second, as is typical in CRM contexts, our dataset does not include competitive information. While quote requests and unfulfilled quotes provide indirect evidence for competition, if competitive data are available directly, one can extend our framework to this setting. Third, while we focus here on B2B pricing, our methodological framework is also appropriate for those B2C settings in which the customer buying process is composed of several interrelated decisions, and where the firm has the opportunity to price discriminate to varying degrees. Finally, in this paper we take an initial step toward studying the under-explored terrain of B2B pricing using a specific empirical application of a metals retailer selling to industrial buyers. Future researchers can extend our framework to other B2B environments to investigate further the generalizability of our conclusions.

References

Akerlof, G. A. 1979. The Case against Conservative Macroeconomics: An Inaugural Lecture. *Economica*, 46(8), 219-237.

Anderson, J.C., D.C. Jain, P. Chintagunta. 1993. Customer Value Assessment in Business Markets: A State of the Practice Study. *Journal of Business-to-Business Marketing,* 1(1), 3-28.

Ansari, A., C. F. Mela. 2003. E-customization. *Journal of Marketing Research, 40*(2), 131-145.

Atchade, Y. 2006. An Adaptive Version for the Metropolis Adjusted Langevin Algorithm with a Truncated Drift. *Methodology and Computing in Applied Probability,* 8, 235-254.

Bell, D., R. Bucklin. 1999. The Role of Internal Reference Price in the Category Purchase Decision. *Journal of Consumer Research,* 26(2), 128–43.

Bijmolt, T. H.A., H. J. van Heerde, R. G. M. Pieters. 2005. New Empirical Generalizations on the Determinants of Price Elasticity. *Journal of Marketing Research*, 42(2), 141-156.

Briesch, R., L. Krishnamurthi, T. Mazumdar, S. Raj. 1997. A Comparative Analysis of Reference Price Models. *Journal of Consumer Research*, 24(2), 202-214.

Chan, T., P.B. Seetharaman. 2004. Estimating Dynamic Pricing Decisions in Oligopolistic Markets: an Empirical Approach Using Micro- and Macro-level Data. *Working Paper.* Washington University.

Chintagunta, P. (2002), Investigating Category Pricing Behavior at a Retail Chain. *Journal of Marketing Research*, 39 (2), 141-154.

Danaher, P., M. Smith. 2011. Modeling Multivariate Distributions Using Copulas: Applications in Marketing. *Marketing Science,*30(1), 4-21.

Dong, X., P. Manchanda, P. Chintagunta. 2009. Quantifying the Benefits of Individual Level Targeting in the Presence of Firm Strategic Behavior. *Journal of Marketing Research,* 46(4), 207-221.

Frank, M. J. (1979) On the Simultaneous Associativity of F(x, y) and x + y - F(x, y). *Aequationes Mathematicae,* 19, 194-226.

Frederick, S., G. Loewenstein. 1999. *Hedonic adaptation.* in D. Kahneman, E. Diener, and N. Schwartz [eds]. *Scientific Perspectives on Enjoyment, Suffering, and Well-Being*. New York. Russell Sage Foundation.

Gonul, F., M. Shi. 1998. Optimal Mailing of Catalogs: A New Methodology Using Estimable Structural Dynamic Programming Models. *Management Science,* 44(9), 1249-1262.

Greenleaf, E. 1995. The Impact of Reference Price Effects on the Profitability of Price Promotions. *Marketing Science,*14(1), 82–104.

Gupta, S., D. Hanssens, B. Hardie, W. Kahn, V. Kumar, N. Lin, N. Ravishanker, S. Sriram. 2006. Modeling Customer Lifetime Value. *Journal of Service Research,* 9(2) 139-155.

Gupta, S., D. R. Lehmann. 2002. Customers As Assets. *Journal of Interactive Marketing*, 17(1). 9-24.

Heckman, J.J. 1981. "Statistical Models for Discrete Panel Data." Pp. 114-78. Structural Analysis of Discrete Panel Data with Econometric Applications. C.F Manski and D. McFadden, editors. MIT Press.

Hamilton, J.D. 1989. A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica,* 57(2), 357-384.

Hardie, B.G. S., E. J. Johnson, P. S. Fader. 1993. Modeling Loss Aversion and Reference Dependence Effects on Brand Choice. *Marketing Science*, 12(3), 378-394.

Higgins, E. T. 1997. Beyond Pleasure and Pain. *American Psychologist,* 52, 1280-1300.

Hinterhuber, A. 2004. Towards Value-based Pricing - An Integrative Framework for Decision Making. *Industrial Marketing Management,* 33(8), 765-778.

Jedidi, K., C. Mela, S. Gupta. 1999. Managing Advertising and Promotion for Long-Run Profitability. *Marketing Science,* 18(1), 1-22.

Johnston, W.J., J.E. Lewin. 1996. Organizational Buying Behavior: Toward an Integrative Framework. *Journal of Business Research*, 35(1), 1–15.

Kahneman, D., J. L. Knetsch, and R. Thaler. 1986. Fairness as a Constraint on Profit Seeking: Entitlements in the Market. *American Economic Review*, 76 (September), 728-741.

Kahneman, D., A. Tversky. 1979. Prospect Theory: An Analysis of Decision Under Risk. *Econometrica,* 47(2), 263-291.

Kalwani, M.U., N. Narayandas. 1995. Long-Term Manufacturer-Supplier Relationships: Do They Pay Off for Supplier Firms? *Journal of Marketing,* 59(1), 1–16.

Kalwani, M.U., H.J. Rinne, Y. Sugita, C.K. Yim. 1990. A Price Expectations Model of Customer Brand Choice. *Journal of Marketing Research*, 27(8), 251–262.

Kalyanaram, G., R. Winer. 1995. Reference Price and Asymmetric Price Response Effects: Empirical Generalizations and Future Research. *Marketing Science,* 14(3), 212-221.

Khan, R., M. Lewis, V. Singh. 2009. Dynamic Customer Management and the Value of One-to-One Marketing. *Marketing Science*, 28(6), 1063-1079.

Kopalle, P. K., A. G. Rao, J. L. Assunção. 1996. Asymmetric Reference Price Effects and Dynamic Pricing Policies. *Marketing Science*, 15(1), 60-85.

Krishnamurthi, L., T. Mazumdar, and S. P. Raj. 1992. Asymmetric Response to Price in Consumer Choice and Purchase Quantity Decisions. *Journal of Consumer Research,* 19(12), 387-400.

Kumar, V., R. Venkatesan, W. Reinartz. 2008. Performance Implications of Adopting a Customer-focused Sales Campaign. *Journal of Marketing*, 72(5), 50–68.

LaPlaca P., J. Katrichis. 2009. Relative Presence of Business-to-Business Research in the Marketing Literature. *Journal of Business-to-Business Marketing*, 16(1), 1-22.

Lehmann, D.R., J. O'Shaughnessy. 1974. Difference in Attribute Importance for Different Industrial Products. *Journal of Marketing,* 38(2), 36–42.

Lewis, M. 2005. A Dynamic Programming Approach to Customer Relationship Pricing. *Management Science,* 51(6), 986-94.

McDonald, I., W. Zucchini. 1997. *Hidden Markov and Other Models for Discrete Valued Time Series.* Chapman and Hall, London.

Montoya, R., O. Netzer, K. Jedidi. 2010. Dynamic Allocation of Pharmaceutical Detailing and Sampling for Long-Term Profitability. *Marketing Science,* 29(5), 909-924.

Morgan, R. M., S.D. Hunt. 1994. The Commitment-Trust Theory of Relationship Marketing. *Journal of Marketing.* 58(3), 20–38.

Nelsen, R. B. 2006. *An Introduction to Copulas (2nd ed.).* Springer-Verlag, New York.

Netzer, O., J. Lattin, V. Srinivasan. 2008. A Hidden Markov Model of Customer Relationship Dynamics. *Marketing Science,* 27(2), 185-204.

Okun, A.M. 1981. *Price and Quantities: A Macroeconomic Analysis*. Washington, D.C.: The Brookings Institute.

Park, S., S. Gupta 2009. A Simulated Maximum Likelihood Estimator for the Random Coefficient Logit Model Using Aggregate Data. *Journal of Marketing Research 46 (4), 531-542.*

Petrin, A., K. Train 2010. A Control Function Approach to Endogeneity in Consumer Choice Models. *Journal of Marketing Research 46 (4), 531-542.*

Pitt, M., D. Chan, R. Kohn. 2006. Efficient Bayesian Inference for Gaussian Copula Regression Models. *Biometrika,* 93(3), 537-554.

Putler, D. S. 1992. Incorporating Reference Price Effects into a Theory of Consumer Choice. *Marketing Science,* 11(3), 287–309.

Reid, D.A., R.E. Plank. 2004. *Fundamentals of Business Marketing Research.* The Haworth Press. New York.

Rossi, P. E.,G. M. Allenby, R. McCulloch. 2005. *Bayesian Statistics and Marketing*, John

Wiley & Sons.

Rossi, P. E., R. McCulloch, G. M. Allenby. 1996. The Value of Purchase History Data in Target Marketing. *Marketing Science,* 15(4), 321-340.

Schweidel, D., E. Bradlow, P. Fader. 2011. Portfolio Dynamics for Customers of a Multi-Service Provider. *Management Science.* Forthcoming.

Shaffer, G., Z. J. Zhang. 1995. Competitive Coupon Targeting. *Marketing Science*, 14(4), 395–416.

Shipley, D., D. Jobber. 2001. Integrative Pricing via the Pricing Wheel. *Industrial Marketing Management,* 30(3), 301–314.

Simester, D., P. Sun, J. Tsitsiklis. 2006. Dynamic Catalog Mailing Policies. *Management Science,* 52(5), 683-696.

Tellis, G. 1988. The Price Elasticity of Selective Demand: A Meta-Analysis of Econometric Models of Sales. *Journal of Marketing Research,* 25(4), 331-341.

Trivedi, P. K., D. M. Zimmer. 2007. *Copula Modeling: An Introduction for Practitioners. Foundations and Trends in Econometrics.* Now Publishers.

Urbany, J. E., T. J. Madden, and P. R. Dickson. 1989. All's Not Fair in Pricing: An Initial Look at the Dual Entitlement Principle. *Marketing Letters*, 1 (1), 17-25

Venkatesan, R., V. Kumar. 2004. A Customer Lifetime Value Framework for Customer Selection and Resource Allocation Strategy. *Journal of Marketing,* 68 (4), 106-125.

Villas-Boas, M. J., R. S. Winer. 1999, Endogeneity in Brand Choice Models. *Management Science*, 45 (October), 1324-1338.

Wilson, E.J. 1994. The Relative Importance of Supplier Selection Criteria: A Review and Update. *Journal of Supply Chain Management,* 30(3), 34-41.

Winer, R. S. 1986. A Reference Price Model of Demand for Frequently-Purchased Products. *Journal of Consumer Research*, 13 (September), 250-256.

Zhang, J., L. Krishnamurthi. 2004. Customizing Promotions in Online Stores. *Marketing Science,* 23(4), 561–578.

Appendix – Price Simulation Procedure

We initialize the simulation by calculating for each customer the state membership probability and the reference price at end of the calibration period. The simulation procedure continues as follows:

1. Draw for each customer for the first purchase event four independent uniform variates (u_1, v_2, u_3, v_4) and use the same random draws for all price points at the purchase event.
2. Using the Frank copula (Nelsen 2006) convert v_2 to u_2

   $$u_2 = -\frac{1}{\theta_{qt}} \ln\left(1 + \frac{v_2(1 - e^{-\theta_{qt}})}{v_2(e^{-\theta_{qt}u_1} - 1) - e^{-\theta_{qt}u_1}}\right),$$

   where, θ_{qt} is the dependence parameter for the Frank copula for the quantity and inter-purchase event time decisions.
3. Generate state specific quantity q_{sij} such that $q_{sij} = F_q^{-1}(u_1)$ with mean $\mathbf{x_{ij}\beta_{qsi}}$ and parameter σ.
4. Update the quantity weighted reference price.
5. Conditioned on q_{sij}, generate t_{sij} such $t_{sij} = F_t^{-1}(u_2)$ with mean $\mathbf{x_{ij}\beta_{qsi}}$ and parameter α_s.[9]
6. Convert v_4 to u_4 using the Frank copula as in point 2 above using the copula parameter θ_{bw}.
7. Calculate the latent utility for quote request $b_{sij}^* = \mathbf{x_{ij}\beta_{bsi}} + e_{bij}$, where

   $$e_{bij} = \ln\left(\frac{u_3}{1 - u_3}\right)$$

8. Calculate the latent utility for quote request, where $w_{sij}^* = \mathbf{x_{ij}\beta}_{\mathbf{wsi}} + e_{wij}$, where

   $$e_{wij} = \ln\left(\frac{u_4}{1 - u_4}\right)$$

   The observed decisions b_{sij} and w_{sij} are governed by the underlying latent utilities as described in section 4.3.2.

[9] Because quantity and inter-purchase event time can have extreme values, which can determine other decisions and reference price going forward, we truncate extreme draws of q_{sij} and t_{sij}. Specifically, we resample q_{sij} or t_{sij} if the current draw is less than half of the minimum, or more than double of the maximum values experienced by the customer during the calibration period.

9. Given the four state-specific behaviors, calculate state-conditional profits. The unconditional profit for the purchase event is calculated as the weighted average of the state-conditional profits, weighted by state-membership probability for the purchase event.
10. Update the state membership for each customer using Equation (15).
11. For each price point in the current purchase event, repeat steps 1-10 to generate the profits in the next purchase event for each new possible price point.
12. Repeat step 11 until the cumulative inter-purchase-event time just exceeds the planning horizon (9 months). This step would result in 5^{J_i} possible price paths for the customer, where 5 is the number of price points, and J_i is the number of purchase events for customer i in the planning horizon.
13. Repeat step 1-12, 200 times to generate 200 Monte Carlo profit draws for each possible price path for each customer.
14. Calculate the average LCV of each price path across the 200 draws. Choose the price path with the highest average LCV as the optimal price path for that customer.
15. Repeat steps 1-14 for each customer.

For the purpose of profit calculations we make several assumptions based on discussion with the seller: (1) we use a discount rate of 12% (a weekly discount rate r of 0.22%); (2) for cost we use actual average wholesale price recorded by the seller for each date in the planning horizon; (3) we add a \$50 fixed cost to each order to capture shipping and administrative costs (4) we use the actual LME commodity prices on each day of the planning horizon.

Table 1: Overall Statistics

Number of customers	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 Customer

	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$)	1,236	1,471	292	808	2,458
Quantity (lbs)	457	553	92	288	968
Inter-purchase-event time (weeks)	6.41	4.31	1.80	5.23	12.45
Unit price (US$)	3.24	0.76	2.44	3.08	4.29

Table 3: Model Selection and Predictive Validity

		Full model		**Benchmark 1 NoHet**		**Benchmark 2 Indep**		**Benchmark 3 NoDyn**		**Benchmark 4 NoEndo**	
	Heterogeneity	Yes		No		Yes		Yes		Yes	
	Interdependence	Copula+HMM		Copula+HMM		No		Copula		Copula+HMM	
	Dynamics	HMM + Reference Price		HMM + Reference Price		Reference Price		No Dynamics		HMM + Reference Price	
	Control Function	Yes		Yes		Yes		Yes		No	
	LMD	**-62,531***		-90,223		-68,959		-72,173		-62,956	
	DIC	**129,996**		180,899		141,478		148,218		130,675	
	Complexity pD	4,779		75		3,432		3,253		4,763	
	Validation LL	**-15,961**		-20,947		-16,455		-16,789		-16,009	
		In-sample	out-sample	In-sample	out-sample	In-sample	out-sample	In-sample	out-sample	In-sample	out-sample
Hit Rate	Quote request	**0.75**	**0.65**	0.63	0.60	0.69	0.64	0.65	0.63	**0.75**	0.65
	Quote acceptance	**0.72**	0.59	0.59	0.57	0.68	**0.61**	0.67	0.60	**0.72**	0.59
RMSE	Quantity	**0.95**	**1.03**	1.73	1.74	1.04	1.07	1.09	1.08	**0.95**	1.03
	Inter-purchase time	**16.22**	**20.03**	25.32	28.65	21.74	23.18	21.18	22.28	16.25	20.07
	Quote request	**0.53**	**0.59**	0.62	0.91	0.61	0.63	0.60	0.68	**0.53**	**0.59**
	Quote acceptance	**0.55**	**0.67**	0.77	1.01	0.67	0.68	0.73	0.72	**0.55**	**0.67**
MAD	Quantity	**0.42**	**0.42**	0.72	0.70	0.42	0.47	0.46	0.50	0.42	0.42
	Inter-purchase time	**12.15**	**14.85**	18.49	18.85	14.98	16.00	14.88	15.67	12.17	14.88
	Quote request	**0.31**	**0.37**	0.45	0.55	0.32	0.39	0.36	0.43	**0.31**	**0.37**
	Quote acceptance	0.41	0.49	0.63	0.76	**0.40**	**0.48**	**0.40**	0.49	0.41	0.49

* Numbers in bold represent the best fit/predictive ability from among the five models.

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

4a – Quantity Decision

	Parameter	Aggregation Level***	Mean	Std Dev.	2.50%	97.50%
Quantity						
State1	Intercept	RE	**1.823**	0.039	-1.897	-1.751
	Gain(j-1)	RE	0.004	0.021	-0.037	0.043
	Loss(j-1)	RE	**-0.035**	0.013	-0.061	-0.010
	LME(j)	RE	0.029	0.129	-0.220	0.268
	LME_Volatility(j)	RE	-0.782	0.486	-1.720	0.122
State2	Intercept	RE	**-1.523**	0.033	-1.587	-1.461
	Gain(j-1)	RE	**0.133**	0.019	0.096	0.169
	Loss(j-1)	RE	**-0.096**	0.014	-0.123	-0.070
	LME(j)	RE	**2.813**	0.539	1.774	3.815
	LME_Volatility(j-1)	RE	**-3.350**	1.054	-5.383	-1.390

4b – Inter-Purchase-Event Time Decision

	Parameter	Aggregation Level	Mean	Std Dev.	2.50%	97.50%
Inter-purchase-event time						
State1	Intercept	RE	**0.977**	0.029	0.919	1.035
	Quantity(j-1)	FE	-0.009	0.029	-0.067	0.048
	Gain(j-1)	RE	**-0.062**	0.014	-0.091	-0.034
	Loss(j-1)	RE	**0.046**	0.014	0.019	0.074
State2	Intercept	RE	**0.854**	0.031	0.793	0.915
	Quantity(j-1)	FE	0.031	0.021	-0.010	0.072
	Gain(j-1)	RE	-0.016	0.016	-0.047	0.016
	Loss(j-1)	RE	**0.037**	0.010	0.016	0.057

4c – Quote Request Decision

	Parameter	Aggregation Level	Mean	Std Dev.	2.50%	97.50%
Quote request vs. direct order behavior (quote request=1)						
State1	Intercept	RE	**-1.125**	0.034	-1.190	-1.062
	Quantity(j)	FE	**0.418**	0.097	0.230	0.599
	Interpurchase time(j)	FE	**0.030**	0.004	0.023	0.036
	Gain(j-1)	RE	**-0.081**	0.026	-0.132	-0.033
	Loss(j-1)	RE	**0.096**	0.024	0.050	0.140
	LME(j)	RE	-0.981	0.878	-2.674	0.652
	LME_Volatility(j)	RE	1.180	1.056	-0.858	3.144
State2	Intercept	FE	**1.037**	0.080	0.881	1.186
	Quantity(j)	FE	**2.090**	0.575	0.981	3.160
	Interpurchase time(j)	FE	0.004	0.007	-0.009	0.017
	Gain(j-1)	RE	**-0.312**	0.015	-0.340	-0.285
	Loss(j-1)	RE	**0.768**	0.017	0.735	0.799
	LME(j)	RE	**-1.880**	0.589	-3.017	-0.784
	LME_Volatility(j)	RE	**3.349**	1.293	0.854	5.753

4d – Quote acceptance Decision

	Parameter	Aggregation Level	Mean	Std Dev.	2.50%	97.50%
Quote acceptance vs. rejection behavior (accept=1)						
State1	Intercept	RE	**0.214**	0.015	0.186	0.242
	Quantity(j)	FE	**-0.220**	0.032	-0.281	-0.162
	Interpurchase time(j)	FE	0.002	0.017	-0.031	0.033
	Gain(j)	RE	**0.211**	0.015	0.183	0.239
	Loss(j)	RE	**-0.180**	0.016	-0.212	-0.150
	Quantity(j) x Gain(j)	RE	**0.087**	0.018	0.051	0.121
	Quantity(j) x Loss(j)	RE	**-0.047**	0.016	-0.079	-0.017
	LME(j)	RE	0.212	0.500	-0.752	1.141
	LME_Volatility(j)	RE	-0.362	1.056	-2.399	1.602
State2	Intercept	RE	**0.096**	0.049	0.002	0.187
	Quantity(j)	FE	**-0.530**	0.046	-0.618	-0.444
	Interpurchase time(j)	FE	**-0.032**	0.008	-0.046	-0.017
	Gain(j)	RE	**0.133**	0.015	0.104	0.161
	Loss(j)	RE	**-0.211**	0.020	-0.250	-0.173
	Quantity(j) x Gain(j)	RE	**0.285**	0.015	0.256	0.312
	Quantity(j) x Loss(j)	RE	**-1.518**	0.018	-1.554	-1.484
	LME(j)	RE	**1.055**	0.460	0.164	1.884
	LME_Volatility(j)	RE	**-2.398**	1.123	-4.751	-0.079

*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).
*** RE represent random-effect parameter and FE represent fixed-effect (common) parameter.

Table 5: Description of the Two HMM States

	"Relaxed" State	"Vigilant" State
Quote request probability	23%	86%
Quote accept probability	65%	52%
Average quantity ordered	432 lb	502 lb
Inter-purchase time	5.5 weeks	8.1 weeks
Average price elasticity	1.3	3.4
Average loss aversion ratio	0.92	3.11
Average sensitivity to LME	0.8	6.7

Table 6: HMM and Distributional Parameter Estimates*,**

	Parameter	Aggregation Level***	Mean	Std Dev.	2.50%	97.50%
Transition matrix						
State1	Intercept	RE	**2.078**	0.026	2.027	2.127
	Gain(j-1)	RE	**0.052**	0.018	0.017	0.086
	Loss(j-1)	RE	**-0.105**	0.018	-0.140	-0.072
State2	intercept,	RE	**2.514**	0.043	2.431	2.594
	Gain(j-1)	RE	**-0.505**	0.018	-0.539	-0.471
	Loss(j-1)	RE	**0.712**	0.019	0.675	0.748
Initial state membership probability						
	Vigilant state (π)	RE	**0.697**	0.030	0.639	0.752
Distributional Parameters						
Std dev. for the quantity model, log scale (σ)		FE	**0.126**	0.052	0.027	0.222
State 1 Shape parameter for inter-purchase event time, log scale (α_1)		FE	**0.116**	0.019	0.079	0.152
State 2 Shape parameter for inter-purchase event time, log scale (α_2)		FE	**0.079**	0.012	0.055	0.102
Frank copula parameter for inter-purchase even and quantity (θ_{qt})		FE	**0.606**	0.178	0.264	0.937
Frank copula parameter for quote request and quote acceptance (θ_{bw})		FE	**-10.023**	0.563	-11.110	-8.975

*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).
*** RE represent random-effect parameter and FE represent fixed-effect (common) parameter

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)	0.895	0.105	0.864	0.136	0.795	0.205
Vigilant (j)	0.096	0.904	0.078	0.922	0.043	0.957

Table 8: Reference Price Coefficients – Full Model vs. NoEndo Model*

		Full Model		NoEndo	
	Parameter	Coefficient	Std Dev.	Coefficient	Std Dev.
Quantity					
Relaxed	Gain(j-1)	0.004	0.021	0.007	0.016
	Loss(j-1)	**-0.035**	0.013	**-0.028**	0.011
Vigilant	Gain(j-1)	**0.133**	0.019	**0.150**	0.013
	Loss(j-1)	**-0.096**	0.014	**-0.072**	0.010
Inter-purchase-event time					
Relaxed	Gain(j-1)	**-0.062**	0.014	**-0.079**	0.013
	Loss(j-1)	**0.046**	0.014	**0.047**	0.013
Vigilant	Gain(j-1)	-0.016	0.016	-0.025	0.014
	Loss(j-1)	**0.037**	0.010	**0.032**	0.008
Quote request vs. direct order behavior (quote request=1)					
Relaxed	Gain(j-1)	**-0.081**	0.026	**-0.122**	0.025
	Loss(j-1)	**0.096**	0.024	**0.118**	0.019
Vigilant	Gain(j-1)	**-0.312**	0.015	**-0.390**	0.018
	Loss(j-1)	**0.768**	0.017	**0.689**	0.018
Quote acceptance vs. rejection behavior (accept=1)					
Relaxed	Gain(j)	**0.211**	0.015	**0.253**	0.014
	Loss(j)	**-0.180**	0.016	**-0.182**	0.017
Vigilant	Gain(j)	**0.133**	0.015	**0.145**	0.017
	Loss(j)	**-0.211**	0.020	**-0.224**	0.020
Inferred Correlation Structure w/ Price					
Quantity	Relaxed	-0.011	0.012		
	Vigilant	-0.021	0.013		
Inter-purchase-event time	Relaxed	0.017	0.016		
	Vigilant	0.025	0.016		
Quote request	Relaxed	0.134	0.072		
	Vigilant	**0.211**	0.065		
Quote acceptance	Relaxed	-0.108	0.059		
	Vigilant	**-0.144**	0.061		
Log Marginal Density (LMD)		-62,531			-62,956

* Estimates in bold indicate a significant effect (that is, 95% posterior interval exclude 0).

Table 9: 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
Inter-purchase-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

Table 10: Policy Performance Comparison Over Nine Months

Policy	**Current Pricing**	**Individual Dynamic Pricing**	**Individual Static Pricing**	**Segment Dynamic Pricing**	**Segment Static Pricing**	**Aggregate Static Pricing**	**Individual Dynamic Pricing (Myopic)**
Average LCV Per Customer ($)	$3,158	$4,795	$3,855	$3,318	$2,759	$2,457	$4,105
Mean Price ($/lb)	$3.24	$3.39	$3.27	$3.43	$3.29 ($3.3; $2.9)*	$3.40	$3.27
Median Price ($/lb)	$2.87	$3.29	$3.10	$3.22	N/A	N/A	$3.10

* Optimal prices in the relaxed and vigilant states, respectively.

Figure 1: Duration of the Asymmetric Price Effects on the Four Decision Components

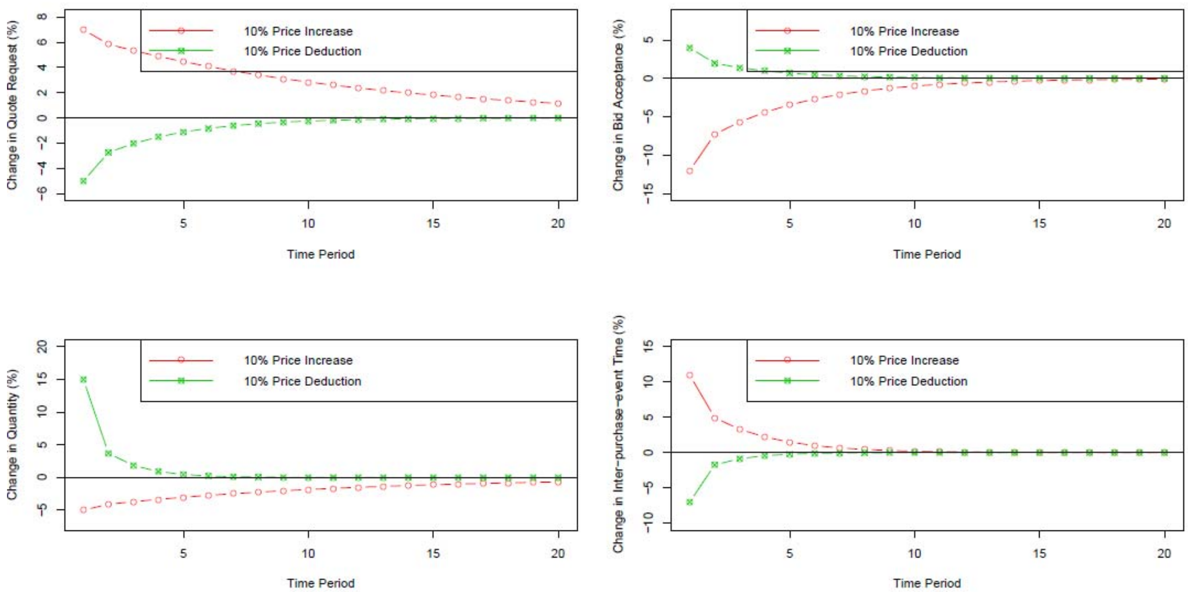

Figure 2: Duration of the Asymmetric Price Effects on Vigilant State Membership

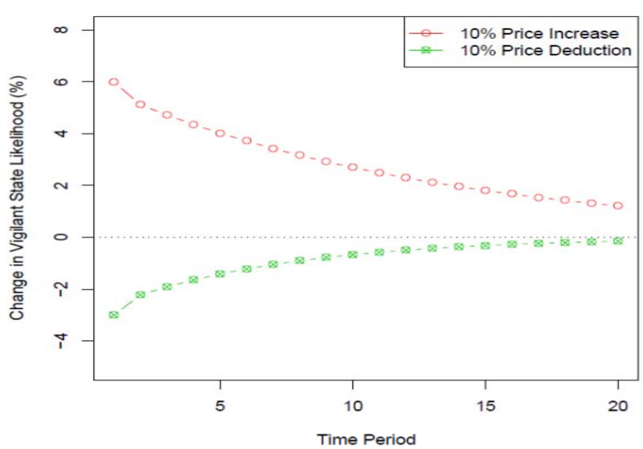

Figure 3: Policy performance Comparison Over Nine Months

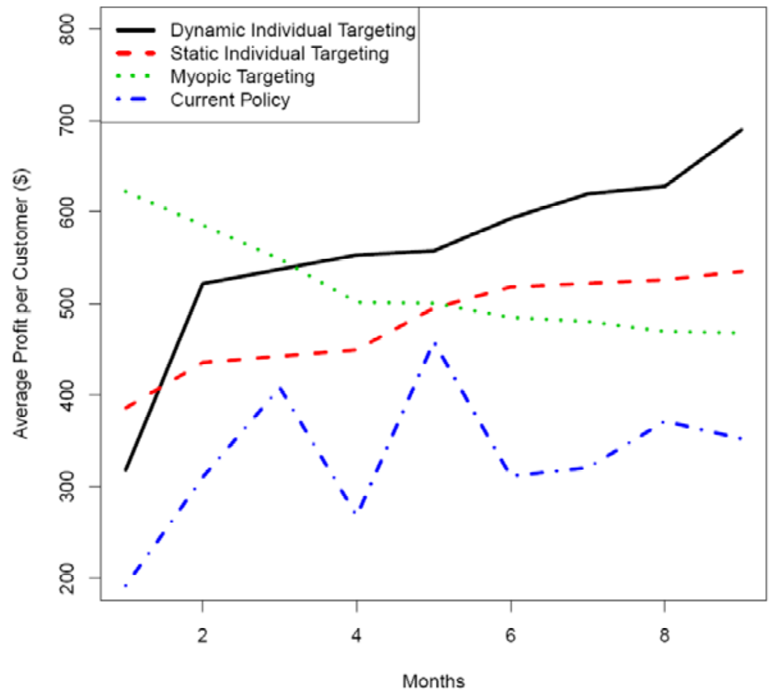

Figure 4: Long-term Customer Value (Nine months)

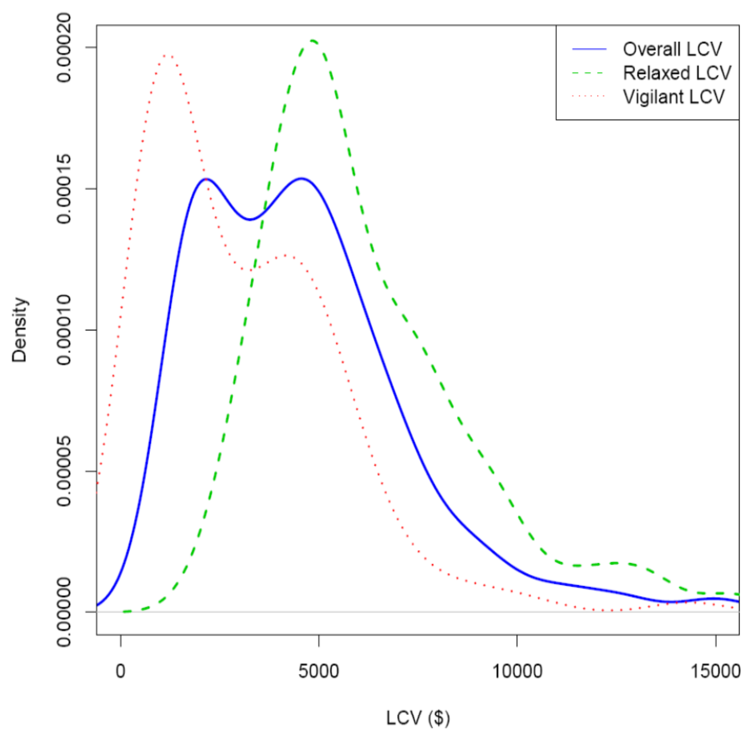

Figure 5: Relaxed State Membership and Profitability over Time

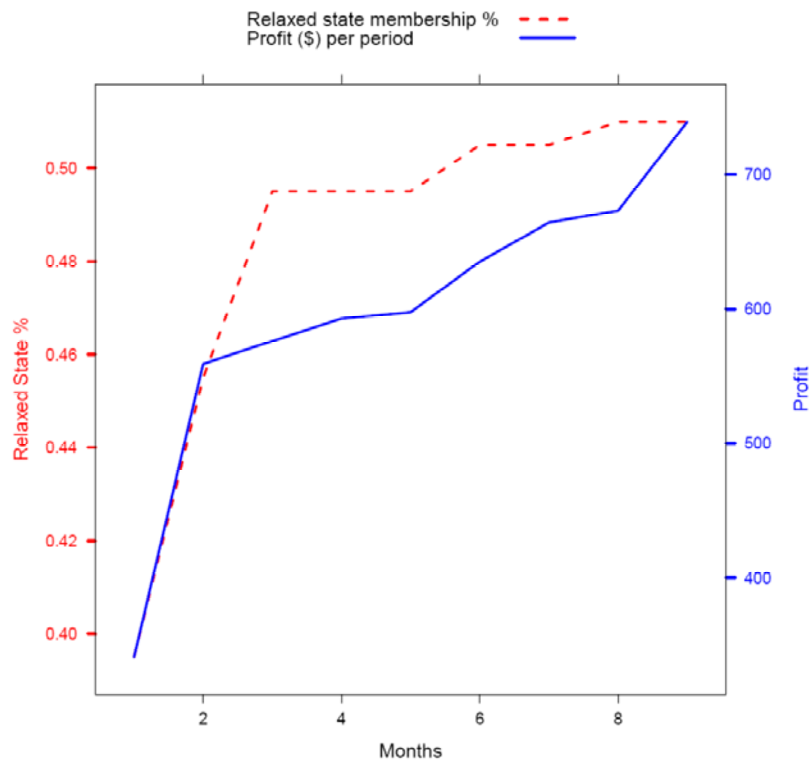

Figure 6: Optimal and Current Unit Price per Pound by State

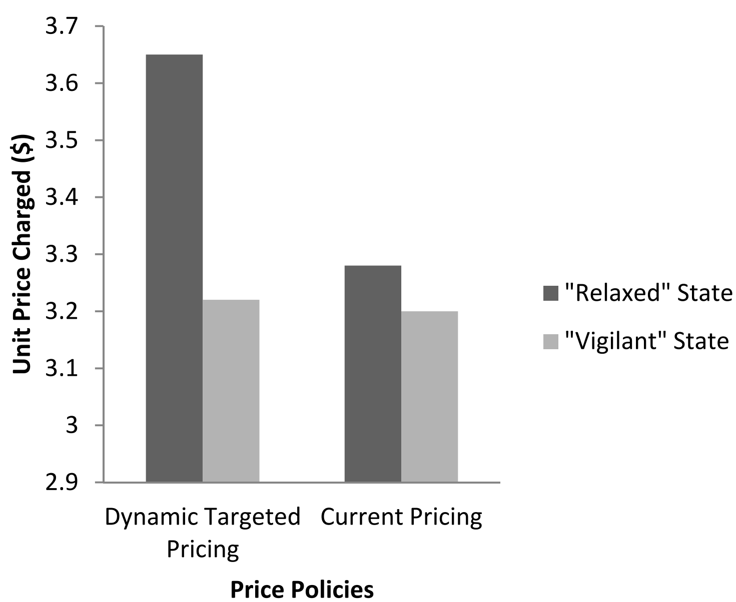

Figure 7: Price Distributions at Different LME Levels

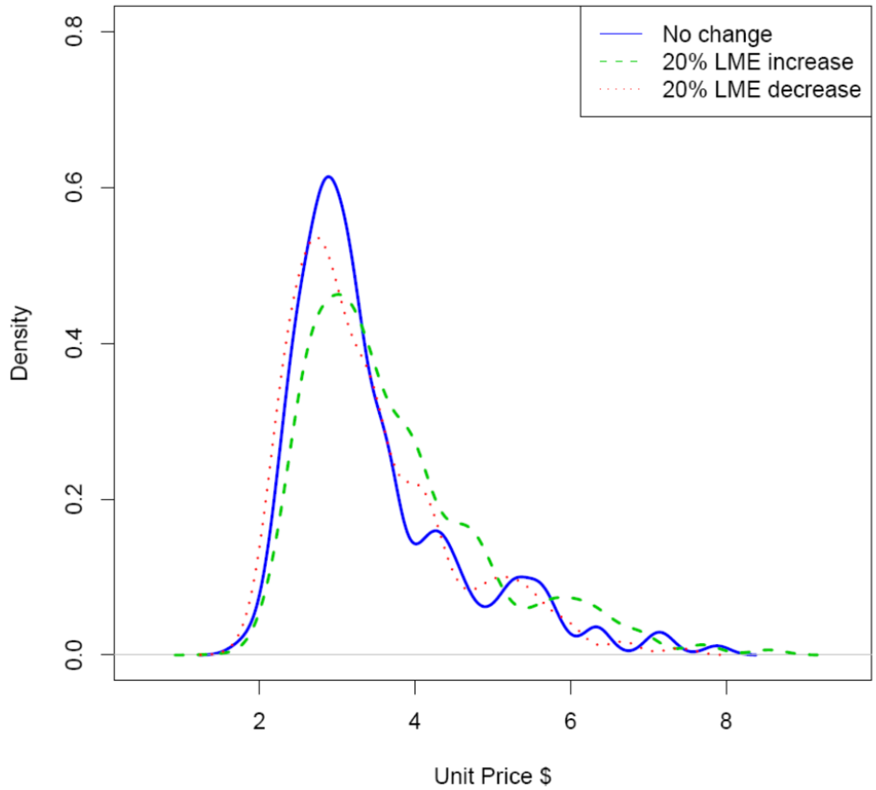

Figure 8 – Optimal and Current Unit Price per Pound by State, under Different LME Regimes

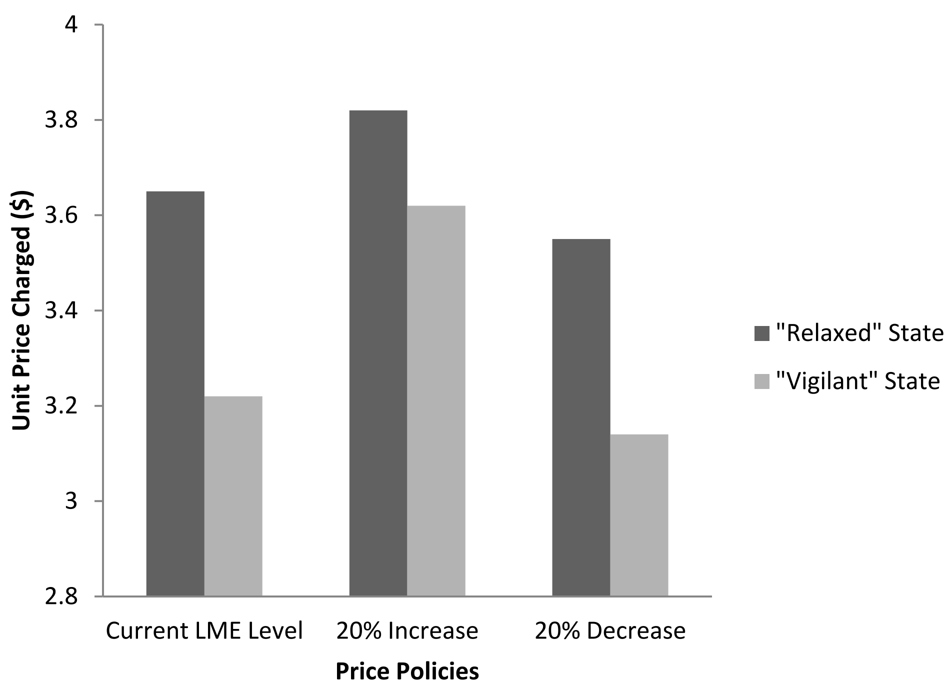

Figure 9 – LCV per State, under Different LME Regimes