

Dynamic Targeted Pricing in B2B Settings

Zaozao Zhang

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
under the Executive Committee
of the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2011

© 2011
Zaozao Zhang
All Rights Reserved

ABSTRACT

Dynamic Targeted Pricing in B2B Settings

Zaozao Zhang

This research models the impact of firm pricing decisions on different facets of the customer purchasing process in business-to-business (B2B) contexts and develops an integrated framework for inter-temporal targeted pricing to optimize long-term profitability for the firm.

Pricing decisions in B2B settings are inherently different from those within the business-to-consumer (B2C) environment, commonly studied in marketing. First, B2B pricing often offers considerable flexibility in implementing first degree and inter-temporal price discrimination, i.e., sellers in B2B contexts can easily vary prices across customers and can even change prices between subsequent purchases by the same customer. While this flexibility affords significant opportunities for the firm, it also requires great care in ensuring long-term profitability. Second, transactions in the B2B environment are often more complex than those in B2C settings. Specifically, the business customer typically makes several interrelated decisions (e.g., when and how much to buy, whether to ask for a quote and whether to accept the seller's bid), which need to be modeled jointly.

The proposed model considers these inter-related decisions in an integrated fashion. In addition, the model accounts for heterogeneity in customers' preferences and behaviors, asymmetric reference price effects, and purchase dynamics, while taking into account the short- and long-term implications of the pricing policy. To model the complexity of inter-

related joint customer decisions, we use hierarchical Bayesian copulas, which weave together different marginal distributions to form joint distributions. To account for dynamics in purchase behavior and to model the possible long-term impact of experienced prices on the different components of the customer's decision, we use a non-homogenous hidden Markov model with multivariate interrelated state-dependent behaviors. In addition, we rely on the behavioral pricing literature in modeling the effect of price, using asymmetric reference price effects.

We calibrated the model using longitudinal transaction data from a metals retailer. The results reveal several substantive insights about the short- and long-term impact of the firm's pricing decisions on each of the inter-related components of the customer's purchasing behavior. Specifically, we find positive interdependence between the quantity and purchase timing decisions and strong negative interdependence between the decision to request a quote and the decision to accept it. Capturing the long-term and asymmetric impact of reference prices, we find that losses not only have larger negative effects relative to gains on customers' buying behavior, but customers also take longer to adapt to losses than they do to gains. Furthermore, the firm's pricing decisions could have a long-term impact on its customers by shifting their preferences between a "vigilant" state - characterized by a cautious approach towards ordering and heightened price sensitivity, and a more "relaxed" state. These dynamics imply that the B2B seller needs to carefully consider both the short- and the long-term consequences of its pricing policy when setting prices for each order. Additionally, the proposed model exhibits superior predictive performance relative to several benchmark models, and in a price policy simulation results in a 52% improvement in profitability compared to the company's current practice. Through pricing simulations

performed are made when pricing in volatile economic environments. Other policy simulations are conducted to examine how the B2B firm should price in the recent economic environment with volatile commodity prices. We find when commodity prices increase, the firm should pass the costs to the customers, when the prices decrease, the firm should “hoard” the profit.

Table of Contents

Acknowledgements	v
Dedication	vii
1. Introduction.....	1
2. Dynamic Individually Targeted Pricing in B2B Settings.....	5
2.1 Pricing in Business-to-Business Markets.....	5
2.2 Dynamic and Individually Targeted Pricing.....	6
2.3 Reference Prices in Customer Buying Behavior.....	7
3. Modeling Framework.....	9
3.1 Initial State Distribution.....	10
3.2 The Markov Chain Transition Matrix.....	10
3.3 The State Dependent Multivariate Interrelated Decisions.....	10
3.3.1 Modeling Quantity and Time between Events.....	12
3.3.2 Modeling Customer Quote Request and Acceptance Decisions.....	14
3.4 The Control Function Approach to Price Endogeneity.....	15
3.5 The HMM Likelihood Function.....	16
3.6 Recovering the State Membership Distribution.....	17

4. Empirical Application.....	17
4.1 Data.....	17
4.2 Variables Description.....	20
4.2.1 Asymmetric Reference Price Effects.....	20
4.2.2 Additional Variables.....	21
4.3 Model Specification.....	22
4.3.1 The State Dependent Decisions.....	22
4.3.2 The Non-Homogenous Transition Matrix.....	24
4.3.3 Price-Generating Mechanism Specification.....	24
4.3.4 Heterogeneity Specification.....	24
4.4 Model Estimation.....	25
4.5 Results and Marketing Implications.....	25
4.5.1. Benchmark Models.....	26
4.5.2. Model Fit and Predictive Ability.....	27
4.5.3. The HMM Parameter Estimates.....	29
4.5.4. Disentangling the Effects of Pricing.....	37
4.5.5 Heterogeneity across Customers.....	40
5. Price Policy Simulation Procedure.....	42
6. Price Policy Simulation Results	43
6.1. Comparison of Prices Charged Under Different Price Policies.....	47
6.2. How do We Value Vigilant and Relaxed Customers Differently?.....	48

6.3 How Much to Charge, and to Whom?	50
6.4 Pricing in a Volatile Economic Environment – the Role of External Reference Price.....	52
7. General Discussion.....	57
References.....	60
 Appendices	
Appendix A. Priors and Full Conditionals Distributions for Hierarchical Bayesian Estimation.....	63
Appendix B – Price Simulation Procedure	64
 List of Tables	
Table 1: Overall Statistics	19
Table 2: Descriptive Statistics Per Customer.....	19
Table 3: Model Selection and Predictive Validity	28
Table 4: State Dependence Estimates for the Four Decisions	32
Table 5: HMM and Distributional Parameter Estimates.....	34
Table 6: Posterior Mean of Transition Matrix Across Customers.....	34
Table 7: Price Parameters - Bayesian Control Function Approach	36
Table 8: Price Elasticities of the Decision Components and the Vigilant State Membership.....	38

Table 9: Policy Performance Comparison Over 9 Months.....47

List of Figures

Figure 1: Duration of the Asymmetric Price effects on Quote Request, Bid Acceptance, Quantity, and Inter-purchase-event Time.....38

Figure 2: Duration of the Asymmetric Price Effects on Vigilant State Membership.....39

Figure 3 Heterogeneity in Price Elasticities Across Customers.....41

Figure 4: Policy performance Comparison Over 9 Months.....46

Figure 5: Comparison of Price Distributions.....47

Figure 6: Long-term Customer Value (9 months).....49

Figure 7: Impact of Relaxed State Membership on Profitability.....50

Figure 8: Optimal and Current Unit Price per Pound by State.....52

Figure 9: Price Distributions at Different LME Levels.....54

Figure 10 – Optimal and Current Unit Price per Pound by State,
under Different LME Regimes.....55

Figure 11 – LCV per State, under Different LME Regimes.....56

Acknowledgement

When all is said and done, and the only step left is writing this acknowledgement section, I can't help but to experience a flashback of the past five years. It was an exciting journey of not only academic discovery, but also one of self-discovery and personal growth. It was a journey with obstacles, winding roads, long stretches of frustration and the occasional despair, but also one studded with moments of “aha” and the subsequent euphoria. Most importantly, it was a journey where I encountered a group of extraordinary people that have given me advice, guidance, and support to complete my study. I will always be professionally and personally indebted to them.

First and foremost, I would like to express my most sincere gratitude to my thesis advisors Asim Ansari and Oded Netzer. I consider myself extremely lucky to be their student. Their energy, enthusiasm and the level of precision that epitomizes true scientists continue to move and inspire me. From day one, their supports have been tremendous and unwavering. Sitting across the table from them in long discussions, I learned how to think analytically, creatively and practically. Getting an email reply from Oded at 3 a.m., or watching Asim coming to the office seven days a week and writing his own C code, exposed me to a life of dedication, discipline and academic purism. Listening to their jokes over the years taught me that above anything else, one needs to harbor an optimistic outlook in life, and, as Asim so elegantly puts it, “have fun”. For these invaluable life lessons, I am deeply grateful.

Second, I thank the remaining members of my dissertation committee, namely Brett Gordon, Raghu Iyengar, and Scott Neslin for their time and constructive comments on the thesis.

Third, I thank the entire faculty at Marketing Division of Columbia Business School, for a world-class, intellectually satisfying and well-rounded education that was crucial for the completion of this thesis. The courses that I took from Don Lehmann, Ran Kivetz, Michel Pham, Kamel Jedidi, Rajeev Kohli, Oded Koenigsberg and my advisors Oded and Asim, have benefited me long after the classes were over. In particular, I want to thank Kamel and Don for bringing

this 22 year old kid (who initially did not know the difference between “CB” and “Quant” and thought the premier marketing journal was the *HBR*) into the doctoral program and encouraging my academic development.

Fourth, I thank Gilad Fishman for giving me timely access to the data used in this dissertation. Without this rich dataset, I would not be able to answer many of the questions that were personally curious, academically interesting, and managerially relevant.

Fifth, I thank past and current students in the program for their camaraderie, constructive discussions, and friendship. Eric, Isaac, Andrew, Ricardo, Hannah, and Yang – the memorable experiences that I shared with you guys through the process were as diverse as they were unifying. In particular, the late night discussions with Yang Li, on topics of mutual interests and concerns, have been insightful and inspirational. I have no doubt that you will be a great scholar one day.

Sixth, I have been fortunate to undertake my Ph.D. in the connected and constantly throbbing city of New York, and I was even more fortunate to have made some great friends with shared interests and life perspectives. Tony, Carl, Paul, Anthony, Thomas, Eddie, you offered the welcomed diversion when I was frustrated. You made a Ph.D. student’s life more diverse and more interesting.

Last and an ocean away from least, I want to thank my mom for all of her hard work in raising me, in providing me with wisdom and education about life. The sacrifice and efforts that she has made to get me to where I am today are immense. Without her, none of this is possible. I owe her my deepest gratitude.

All in all, it was an interesting five years, and I thoroughly enjoyed the process and the amazing people that I met along the way.

Dedication

For my mother, and all of those who have been there for me.

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 and business-to-consumers (B2C) markets are approximately equal in size (U.S. Department of Commerce 2007), it is estimated that 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, the study of pricing in B2B environments is particularly under researched (Reid and Plank 2004). While Kalayanaram and Winer (1995), in their review of the literature on pricing and reference prices, called for further research to investigate the impact of pricing and reference prices on industrial buyers and urged researchers to go beyond contexts involving frequently purchased packaged goods, in the decade that has passed since this call to action, little has been done to address this gap.

Pricing decisions in B2B contexts differ inherently from those within business-to-consumer (B2C) environments on multiple facets. First, B2B environments, which are often characterized by highly variable costs of goods, variable order sizes and personal selling oriented transactions, offer considerable flexibility in implementing first degree and inter-temporal price discrimination. That is, sellers in B2B situations can easily vary prices across customers and can even change prices between subsequent purchases of the same customer. While this flexibility generates significant opportunities for the firm, it also requires great care in ensuring long-term profitability.

Second, transactions in B2B markets are more complex than those in B2C situations. The business customer typically makes several inter-related decisions (e.g., when and how much to buy, whether to ask for a price quote and whether to accept the seller's bid) thus requiring an

integrated approach for modeling the impact of pricing on these different facets of purchasing behavior.

Third, the B2B environment is often characterized by long-term relationships between buyers and sellers (Morgan and Hunt 1994). Pricing decisions can play a vital role in developing and sustaining such relationships (Kalwani and Narayandas 1995). Accordingly, pricing models in the B2B context need to account for purchasing dynamics and the possible long-term effect of pricing decisions, and should be capable of representing the evolution of the relationships over time (Netzer et al. 2008). For example, the development of trust between the buyer and the seller over time may shift the buyer from a state of higher price sensitivity to a state involving lower price sensitivity, thus creating dynamic price effects.

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 customers' internally constructed reference prices play an important role in their 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.

The objective of this paper is to develop an integrated framework for modeling the multiple impacts of pricing in B2B transactions, and for aiding firms in implementing first-degree and inter-temporal price discrimination for long-run profitability. Our modeling framework jointly considers the multiple inter-related decisions that customers make on each potential purchasing occasion in an integrated fashion. It accounts for heterogeneity in customer preferences and behaviors to facilitate targeting, incorporates asymmetric reference price effects to understand the role of behavioral effects, and models purchase dynamics and the possibly long-term effect of pricing to facilitate optimal dynamic targeting.

In our modeling framework, we incorporate the marginal effects of reference prices on each one of the customer decisions and then weave together these interrelated decisions using a hierarchical Bayesian copula approach (Nelsen 2006; Pitt et al. 2006; Trivedi and Zimmer 2007) under a control function framework (Gupta and Park 2009; Petrin and Train 2010). Copulas

allow for the construction of joint distributions from univariate marginal distributions. Specifically, we use the Frank copula (Frank 1979; Trivedi and Zimmer 2007), which accommodates both positive and negative dependence, to combine the four customer decisions: purchase timing, purchase amount, quote request, and bid acceptance distributions in a flexible fashion, while simultaneously accounting for appropriate conditionality and selectivity. The Bayesian control function framework allows us to fully account for potential endogeneity in prices, and to avoid biased estimates on customer's price sensitivities. We then represent dynamics in purchasing behavior using a multivariate non-homogeneous hidden Markov model (Montoya et al. 2009; Netzer et al. 2008; Schweidel et al 2009). Thus, prices are allowed to affect in the short run the four inter-related behaviors, the actual impact being conditional on the latent behavioral state of the customer. In addition, prices can also influence the transition between latent states and thereby have a long-term impact.

Thus, our model is unique in its ability to represent the short- and long-term impact of pricing on multiple facets of customer purchasing behavior and provides the firm greater fine-tuning capability in eking out more profits from each customer an over time. Methodologically, to the best of our knowledge, this is one of the first papers to apply copulas in marketing (see Danaher and Smith, forthcoming for an exception), and the first paper to integrate copulas, control functions with a hidden Markov model (HMM) in a comprehensive Bayesian framework.

We calibrate the model using longitudinal transaction data from a retailer of metals (aluminum and steel) selling to industrial customers. The results of our empirical application demonstrate that the asymmetric reference price has varying short- and long-term effects on each component of the industrial customer's decisions. For example, a price increase has the strongest immediate negative impact on the customer's decisions on when to buy and whether to accept or reject the seller's bid. The strongest long-term impact of a price increase, on the other hand, is on the customer's decision whether to request a price quote (as opposed to order directly without asking for a quote first). The decisions with respect to when to purchase, whether to ask for a quote and whether to accept the quote, all exhibit loss aversion, whereas the decision of

how much to buy exhibits gain seeking. Furthermore, we identify two latent buying behavior states: a vigilant state characterized by enhanced price sensitivity and a cautious approach towards ordering, and a relaxed state characterized by more direct orders and lower price sensitivity. We also find that past prices can have a long-term effect on the customer buying behavior, and price changes (particularly price increases) can take the customer long time to adapt to. These dynamics imply that the B2B seller needs to carefully consider both the short- and the long-term consequences of the pricing policy when setting prices for each order.

The application of the model for individual targeted pricing is illustrated using a “what-if” analysis of alternative pricing policies. We show that if the company adopts our dynamic targeted pricing policy, it would drastically increase the profitability. Additionally, using a validation sample, we demonstrate that the proposed model improves the ability to predict all of the dimensions of the buying process, relative to several benchmark models.

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 framework for dynamically modeling multiple facets of customer purchasing in B2B settings. On the substantive front, it provides an individually targeted dynamic pricing framework for B2B settings, which 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 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 presents the proposed modeling approach for capturing dynamic pricing effects on the industrial buyer’s demand. Section 4 illustrates the proposed model using B2B transaction data from an industrial metals retailer. Section 5 concludes by discussing practical implications, theoretical contributions, and future directions.

2. Dynamic Individually Targeted Pricing in B2B Settings

Our research lies at the intersection of two literatures. These include the relatively thin body of work involving choice modeling in B2B settings and the abundant literature on pricing which has focused, primarily, on the B2C context. In this section, we briefly review these twin literatures and highlight the challenges involved in adapting the received wisdom on price targeting, dynamic pricing and behavioral modeling to the B2B context. We also focus on the opportunities that exist in this area and showcase how we contribute to the above literatures.

2.1 Pricing in Business-to-Business Markets

Although pricing is one of the most important aspects of B2B marketing, it is paradoxically also one of the most neglected areas of study (Reid and Plank 2004). The majority of the research on B2B pricing is based on conceptual models and surveys (Johnston and Lewin 1996). The scant attention given to pricing models in B2B research could stem from the view advanced by some researchers that price is of little importance relative to other attributes in B2B settings, and that business customers are less sensitive to price when compared to consumers in B2C markets (Avila et al 1993; Lehmann and O'Shaughnessy 1974 and 1982; Wilson 1994). However, others have argued that prices do in fact play a major role in many B2B situations (e.g., Hinterhuber 2004). In this paper, we demonstrate that because of long-term buyer-seller relationships, business customers can transition over time between states of low and high price sensitivity, thus possibly reconciling these two opposing views.

Decision makers (buyers and sellers) in B2B settings are often assumed to behave rationally (Reid and Plank 2004). Accordingly, the B2B pricing literature sparingly draws upon the voluminous behavioral pricing research on reference pricing and loss aversion (see Briech et al. 1997 and Kalyanaram and Winer, 1995 for a review). One of our objectives in this research is to examine whether agents in B2B settings are sensitive to behavioral pricing effects.

2.2 Dynamic and Individually Targeted Pricing

Customization is an emerging topic in academic research (Gupta et al. 2006) and in the world of practice (Peppers and Rogers 1994). The empirical literature on targeting has primarily focused on non-price marketing actions such as catalog mailing (Bitran and Mondschein 1996; Gönül and Shi 1998; Gönül and Ter Hofsted 2006; Simester et al. 2006), coupons (Gönül and Srinivasan 1996; Rossi et al. 1996), e-mail marketing campaigns (Ansari and Mela 2003), and pharmaceutical detailing and sampling (Dong et al. 2009; Montoya et al. 2009), as it is natural to think of such marketing actions as 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.

Several studies, however, have investigated targeted pricing at a more aggregate level (i.e., at the level of customer segments). Taking an analytical, competitive strategy, approach, Chen (1997), Klemperer (1995), and Shaffer and Zhang (1995 and 2000) derived qualitative conclusions about which segments of customers to target to enhance profitability. In the empirical domain, Zhang and Krishnamurthi (2004) showed how promotions can be targeted to segments of online customers and Lewis (2005) proposed targeting of newspaper subscription rates to different segments of customers.

In addition to targeting, dynamics has been incorporated in pricing models in myriad ways. Reference price models based on the customer's history of purchases (e.g., Hardie et al. 1993) capture some degree of dynamics in price sensitivity and can generate implications for the firm's dynamic pricing policy (Greenleaf 1995; Kopalle et al. 1996). In addition, changing customer preferences may also generate dynamics in price sensitivity. For example, Lewis (2005) modeled dynamics in the form of customers' forward looking expectations about future prices. In a related vein, several studies have investigated the relationship between state-dependence and pricing decisions (Berry and Pakes 2000; Chan and Seetharaman 2004; Pakes and McGuire 2001).

Our research differs from the above studies in several aspects. First, because our context involves longitudinal transaction from a B2B retailer, our objective is to customize prices both at the *individual level*, as well as *temporally* over repeated transactions for the same customer. Second, most pricing models investigate the effect of the firm's pricing on brand choices or a single purchase decision. For example, Lewis (2005) investigated the effect of pricing on customers' decisions to renew newspaper subscriptions. On the other hand, we investigate the impact of the firm's pricing policy on a complex sequence of inter-related customer decisions that are typical of B2B environments. Finally, we allow for dynamic pricing and for the enduring impact of reference price effect using a state-of-the-art multivariate non-homogenous HMM.

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; Lattin and Bucklin 1989; Winer 1986). However, reference price effects have not been incorporated in the empirical research on targeted pricing. Ignoring the effect of reference prices could potentially bias the proposed targeting policy. Furthermore, despite the popularity of reference prices, they have been used almost exclusively in B2C contexts (often using scanner panel data). In this research we explore the role of reference prices in B2B settings. B2B purchasing is often considered rational because of the relative expertise and sophistication of the B2B buyer, and due to the typically large size of transactions. Accordingly, the existence of behavioral pricing effects in this context is not obvious.

The literature distinguishes between "internal" and "external" reference prices. External reference prices stem from forces exogenous to customers such as the manufacturer's suggested retail price (MSRP), regular prices (Burton et al. 1993; Urbany et al. 1988), or the prices of other products in the consideration set (Hardie et al. 1993), and are generally assumed to be observable and common to all customers. On the other hand, internal reference prices are assumed to be individual-specific, and are often constructed based on the customer's observed prices on

previous purchase occasions (Hardie et al. 1993; Kalwani et al. 1990; Krishnamurthi et al. 1992; Lattin and Bucklin 1989; Winer 1986). Operationally, internal reference prices are often calculated as the price observed in the previous purchase occasion, or a weighted average of past observed price. Briesch et al. (1997) compared several reference price mechanisms, including several forms of internal and external reference prices, and concluded that internal reference prices best fit the data. Given the relatively long buyer-seller relationships in the B2B environment, and because past transactions are often formally recorded by buyers, one may argue that internal reference prices could play an important role in B2B purchasing as well.

Reference prices are an application of the more general notion of reference point in prospect theory (Tversky and Kahneman 1979). According to prospect theory, gains and losses relative to the reference point have asymmetric effects, such that losses loom larger than gains. Thus, an observed price that is lower than the reference price is seen as a “gain”, whereas, a higher observed price is perceived as a “loss”. Several studies have demonstrated such asymmetric reference price effects in the context of brand choice (Chang et al. 1999; Hardie et al. 1993; Kalwani et al. 1990; Putler 1992), purchase timing (Bell and Bucklin 1999), and purchase quantity (Krishnamurthi et al 1992). In our modeling framework, we incorporate internal reference price effects based on the buyer’s past experienced prices and allow for asymmetric gain and loss effects. Furthermore, we allow reference prices to affect the customer’s transitions between several buying behavior states, capturing the, possibly, long-term effect of reference prices. Thus, our research furthers the reference price literature along several dimensions: 1) we examine reference price effects in a B2B setting, which traditionally assumed normative customers, 2) the complexity of the purchasing decision in B2B settings offers an opportunity to investigate the effect of reference prices not only on the decision to purchase but also their simultaneous effect on the timing of the purchase, the quantity purchased and the nature of the ordering process, 3) we account for reference price effects in a context which permits both first degree as well as inter-temporal price discrimination, 4) we investigate the

possibly long-term effect of reference prices. In the next section we formulate our modeling framework aimed at achieving these objectives.

3. Modeling Framework

The model described in this section is a customer-level model of buying behavior comprised of four interrelated customer decisions: 1) when to buy, 2) how much to buy, 3) whether to ask for a price quote (hence allowing the seller to bid for business) or order directly without asking for a quote first, and 4) whether to accept the quote if a bid is made by the seller. We consider a set of customers, each of whom is involved in repeated interactions with the seller. The seller observes the buying behavior and pricing history for each customer as well as the marketing environment at every time period.

We model a sequence of purchase events for each customer. A *purchase event* is described by either a direct order from the customer which always result in a purchase, or a request for a quote followed by the acceptance or rejection of the quote. Each purchase event is also described by the timing of the event and a quantity that is desired. On receiving a quote request or a direct order, the seller decides the unit price to charge for the customer.

To allow for customer dynamics over repeated purchase events we allow the customer to transition between different latent buying behavior states that differentially impact the customer decisions. For example, a state with a high propensity to request quotes rather than order directly when coupled with high price sensitivity may characterize a customer who is in an evaluative mode and is testing the seller. Over time, the customer may develop trust in the relationship with the seller and transition to a more relaxed or trusting state characterized by a higher propensity to order directly, without a quote request, and a diminished sensitivity to prices. The seller's past pricing decisions may affect the customer's transition between successive states. For example, pricing that is favorable to the customer may facilitate a transition from an evaluative state to a more relaxed state.

To model such behavior, we define a multivariate non-homogeneous hidden Markov model. In the HMM, the joint probability of a sequence of interrelated decisions up to purchase event j for customer i , $\{Y_{i1} = y_{i1}, \dots, Y_{ij} = y_{ij}\}$ is a function of three main components: (1) the initial hidden states membership probabilities (π_i), (2) a matrix of transition probabilities among the buying-behavior states ($\Omega_{ij-1 \rightarrow j}$), and (3) a multivariate likelihood of the interrelated customer decisions conditional on the customer's buying-behavior state. We describe our formulation of each of these components next.

3.1 Initial State Distribution - Let s denote a buying-behavior state ($s = 1, 2, \dots, S$). Let π_{is} be the probability that customer i is initially 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.

3.2 The Markov Chain Transition Matrix - We model the transitions between states as a Markov process. Each element of the transition matrix ($\Omega_{i,j-1 \rightarrow j}$) can be defined as

$\omega_{ijss'} = P(S_{ij} = s' | S_{ij-1} = s)$, the conditional probability that customer i moves from state s at purchase event $j-1$ to state s' at purchase event j , and where $0 \leq \omega_{ijss'} \leq 1 \quad \forall s, s'$, 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:

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

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

3.3 The State Dependent Multivariate Interrelated Decisions - Conditional on being in state s at purchase event j , the customer decides when to order, how much to order, whether to order directly or submit a request for price quote, and whether to accept the quote (firm's bid) if a quote was requested. These four decisions are unconditionally interrelated as they all depend on

the customer's latent state. To allow these decisions to also be conditionally interrelated, we model dependence using a copula (Danaher and Smith 2009; Trivedi and Zimmer 2007). The use of copula allows us to model each decision flexibly using appropriate marginal distributions. At the same time, copulas capture interdependence to yield a rich multivariate specification.

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 quote or order (i.e., the inter-purchase-event time), and b_{ij} and w_{ij} are the binary quote request and quote acceptance decisions, respectively.

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:¹

$$L_{ij|s} = f_s(q_{ij}, t_{ij}, b_{ij}, w_{ij}) = f_s(q_{ij}, t_{ij}) \Pr_s(b_{ij}, w_{ij} | q_{ij}, t_{ij}) \quad (2)$$

In the above, we assume that the joint decisions of when to order and how much to order arise 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. For example, it is possible that customers are more likely to request a price quote on larger orders. Naturally, the decision to accept or reject a quote occurs only on occasions when the customer decides to request a quote rather than order directly from the seller ($b_{ij} = 1$). Let δ_{ij}^b be a quote request indicator, which equals 1, if purchase event j for customer i is a quote request and 0, if it is a direct order. Accordingly, we specify the joint probability of b_{ij} and w_{ij} using a selectivity approach:

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

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 $S(t_{ij}^*)$ be the survival function for the censored observation, and let δ_{ij}^c be a censoring indicator, which equals 1 if

¹ To avoid clutter, we describe first the model in the general distribution form, omitting the customer-level subscript i , and then outline the particular distributions and parameterization used.

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:

$$\begin{aligned} L_{ij|s} &= f_s(q_{ij}, t_{ij}, b_{ij}, w_{ij}) = \\ &= S_s(t_{ij}^*)^{\delta_{ij}^c} \left[f_s(q_{ij}, t_{ij}) \Pr_s(b_{ij} = 0 | q_{ij}, t_{ij})^{1-\delta_{ij}^b} \left[\Pr_s(w_{ij} | b_{ij} = 1, q_{ij}, t_{ij}) \Pr_s(b_{ij} = 1 | q_{ij}, t_{ij}) \right]^{\delta_{ij}^b} \right]^{1-\delta_{ij}^c}. \end{aligned} \quad (4)$$

Next, we define the distributional assumptions of each of the four customer decisions and the copula approach to interrelate these distributions.

3.3.1 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

$$F_s(q_{ij}, t_{ij}) = C(F_{sq}(q_{ij}), F_{st}(t_{ij})), \quad (5)$$

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

$$f_s(q_{ij}, t_{ij}) = f_{sq}(q_{ij}) f_{st}(t_{ij}) H_{qt}[F_{sq}(q_{ij}), F_{st}(t_{ij})], \quad (6)$$

where, $f_{sq}(q_{ij})$ and $f_{st}(t_{ij})$ are the univariate marginal densities, and $H_{qt}[\cdot, \cdot]$ 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.

Many different distributions are possible candidates for modeling the inter-purchase-event time and quantity variables. In the current application, we assume that the times between purchase events follow a log-logistic distribution (Lancaster 1990, pp 44). The p.d.f of the log-logistic is given by

$$f_s(t_{ij}) = \frac{e^{x_{ij} \beta_{tsi} + \xi_{ij}^t} \alpha_s t_{ij}^{\alpha_s - 1}}{(1 + e^{x_{ij} \beta_{tsi} + \xi_{ij}^t} t_{ij}^{\alpha_s})^2}, \quad (7)$$

and the c.d.f is written as

$$F_s(t_{ij}) = \frac{e^{x_{ij} \beta_{tsi} + \xi_{ij}^t} t_{ij}^{\alpha_s}}{1 + e^{x_{ij} \beta_{tsi} + \xi_{ij}^t} t_{ij}^{\alpha_s}}, \quad (8)$$

where, $\alpha_s > 0$, is a shape parameter, β_{tsi} is vector of coefficients for customer-level, purchase-event specific covariates such as prices and 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 endogeneity (see section @@). We use the log-logistic distribution as it is flexible in accommodating both monotonic and non-monotonic hazard functions. For $\alpha_s \leq 1$ the hazard decreases with t_{ij} ; for $\alpha_s > 1$ the hazard function first increases and then decreases with t_{ij} , reaching a maximum at $t = [(\alpha_s - 1) / e^{\mathbf{x}_{ij}'\beta_{tsi} + \xi_{ij}^t}]^{1/\alpha_s}$.

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

$$f_s(q_{ij}) = \frac{\phi\left(\frac{\log q_{ij} - \mathbf{x}_{ij}'\beta_{qsi} - \xi_{ij}^q}{\sigma}\right)}{\sigma q_{ij}}, \quad (9)$$

where β_{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 discussed subsequently, σ is the scale parameter, and ϕ represents the p.d.f of a standard normal distribution. The corresponding c.d.f is given by

$$F_s(q_{ij}) = \Phi\left(\frac{\log(q_{ij}) - \mathbf{x}_{ij}'\beta_{qsi} - \xi_{ij}^q}{\sigma}\right), \quad (10)$$

where, Φ represents the c.d.f of the standard normal distribution.

Given these marginals, we use the Frank copula (Frank 1979; Trivedi and Zimmer 2007) to model the interdependence between quantity and inter-purchase-event times. There are many families of copulas which differ in the nature of the dependence they represent. We chose the Frank copula as 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

$$C[F_{sq}(q_{ij}), F_{st}(t_{ij})] = -\frac{1}{\theta_{qt}} \ln \left(1 + \frac{(e^{-\theta_{qt} F_{sq}(q_{ij})} - 1)(e^{-\theta_{qt} F_{st}(t_{ij})} - 1)}{e^{-\theta_{qt}} - 1} \right), \quad (11)$$

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. The resulting joint density of the two variables is derived in Appendix A.

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

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

Similarly, conditional on a price quote let customer's i 's binary decision to accept or reject the quote on purchase event j be given by w_{ij} . We assume that this is driven by the latent utility w_{ij}^* such that:

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

The joint distribution of b_{ij} and w_{ij} can be specified as the product of the marginal distribution for the quote request (i.e., b_{ij}) and the conditional distribution of quote acceptance or rejection given a quote. Thus, for 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 ($w_{ij} = 1$) we have (see Appendix A for the full derivation)

$$\begin{aligned} \Pr_s(w_{ij} = 1 | b_{ij} = 1, t_{ij}, q_{ij}) \times \Pr_s(b_{ij} = 1 | t_{ij}, q_{ij}) = \\ \Pr_s(w_{ij} = 1, b_{ij} = 1 | t_{ij}, q_{ij}) = 1 - \Pr_s(b_{ij}^* < 0) - \Pr_s(w_{ij}^* < 0) + \Pr_s(b_{ij}^* < 0, w_{ij}^* < 0) \end{aligned}, \quad (12)$$

Similarly, for rejected orders ($w_{ij} = 0$) we have

$$\begin{aligned} & \Pr_s(w_{ij} = 0 | b_{ij} = 1, t_{ij}, q_{ij}) \Pr_s(b_{ij} = 1 | t_{ij}, q_{ij}) = \\ & \Pr_s(w_{ij} = 0, b_{ij} = 1 | t_{ij}, q_{ij}) = \Pr_s(w_{ij}^* < 0) - \Pr_s(b_{ij}^* < 0, w_{ij}^* < 0) \end{aligned} \quad (13)$$

As for the quantity and inter-purchase-event time copula, we model the joint probability $\Pr_s(b_{ij}^* < 0, w_{ij}^* < 0)$ using the Frank copula. Thus,

$$\Pr_s(b_{ij}^* < 0, w_{ij}^* < 0) = C(F_{sb^*}(0), F_{sw^*}(0)) = -\frac{1}{\theta_{bw}} \ln \left(1 + \frac{(e^{-\theta_{bw} F_{sb^*}(0)} - 1)(e^{-\theta_{bw} F_{sw^*}(0)} - 1)}{e^{-\theta_{bw}} - 1} \right), \quad (14)$$

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

$$\Pr_s(b_{ij}^* < 0) = \frac{1}{1 + e^{\mathbf{x}_{ij} \boldsymbol{\beta}_{bsi} + \xi_{ij}^b}} \quad \text{and} \quad \Pr_s(w_{ij}^* < 0) = \frac{1}{1 + e^{\mathbf{x}_{ij} \boldsymbol{\beta}_{wsi} + \xi_{ij}^w}}. \quad (15)$$

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

Inserting equations (6)-(15) into equation (4) we get the full likelihood of observing the four interrelated customer decisions conditional on the customers state. To complete the HMM specification, we vectorize the state-specific full conditional multivariate behaviors

$\mathbf{m}'_{ij} = [L_{ij|1}, L_{ij|2}, \dots, L_{ij|S}]$. 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 $\tilde{\beta}_{b01i} \leq \tilde{\beta}_{b02i} \leq \dots \leq \tilde{\beta}_{b0Si}$ by setting $\tilde{\beta}_{b0si} = \beta_{b01i} + \sum_{s'=2}^s \exp(\beta_{b0s'i})$; $s=2, \dots, S$. As both the intercepts and the response parameters are state-specific, we impose this restriction at the mean of the vector of covariates, i.e., by mean-entering \mathbf{x}_{ij} .

3.4 The Control Function Approach to Price Endogeneity

It is possible that the firm pricing decisions are based on unobserved factors that also impact the customer decisions. For example, it is possible that seller targets each customer differentially

based on private information it has that is not observed to the researcher. In such case, price will be correlated with the unobserved components (the ξ 's) of the four distributions. Ignoring this endogeneity can lead to misleading inferences about the price sensitivities of consumers (Winer and Villas Boas 1999). We use a Bayesian control function approach to account for price endogeneity (Gupta and Park 2009; Petrin and Train 2010). In the control function approach, we express price as a function of observed instrument variable z_{ij} that is correlated with price, but is uncorrelated with the unobserved factors impacting four decisions. Formally, we have

$$p_{ij} = \gamma_1 + \gamma_2 z_{ij} + \mu_{ij}, \quad (16)$$

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), (9), (10), and (15) to account for endogeneity. For each one of the four decisions the bivariate normal distribution can be written as

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

where σ_p^2 is the variance of μ_{ij} , σ_l^2 is the variance for the random shocks ξ_{ij}^l , and $\rho_{p,l}$ is the covariance terms capturing the correlation between μ_{ij} and ξ_{ij}^l . We use wholesale price as an instrument for price. The wholesale price is observed by the retailer but not to customers. In particular, in our applications, discussions with the management indicate that the salesperson observes the wholesale price on their computer screen when determining the price. Wholesale price has been commonly used as instrument for price (cite). After incorporating the price endogeneity, the likelihood in Equation (2) can be amended to obtain the state-specific likelihood for all endogenous variables.

$$L_{ij|s} = f_s(q_{ij}, t_{ij}, b_{ij}, w_{ij}, p_{ij}) = f_s(q_{ij}, t_{ij} | p_{ij}) \Pr_s(b_{ij}, w_{ij} | q_{ij}, t_{ij}, p_{ij}) f(p_{ij}). \quad (18)$$

3.5 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 purchase event J ($\mathbf{Y}_{i1}, \mathbf{Y}_{i2}, \dots, \mathbf{Y}_{iJ}$).

The likelihood of observing the customer's decisions over J purchase events can be succinctly written as (McDonald and Zucchini 1997):

$$L_{iJ} = P(\mathbf{Y}_{i1} = \mathbf{y}_{i1}, \dots, \mathbf{Y}_{iJ} = \mathbf{y}_{iJ}) = \boldsymbol{\pi}_i \mathbf{M}_{i1} \boldsymbol{\Omega}_{i,1 \rightarrow 2} \mathbf{M}_{i2} \dots \boldsymbol{\Omega}_{i,J-1 \rightarrow T} \mathbf{M}_{iJ} \mathbf{1}', \quad (19)$$

where $\boldsymbol{\pi}_i$ is the initial state distribution described in section 3.1, $\boldsymbol{\Omega}_{i,j \rightarrow j+1}$ is the transition matrix described in section 3.2, \mathbf{M}_{ij} is a $S \times S$ diagonal matrix with the elements $L_{ij|s}$ from Equation (18) on the diagonal, and $\mathbf{1}'$ is a $S \times 1$ vector of ones. To avoid underflow, we scale the likelihood function in Equation (19) using the approach suggested by MacDonal and Zucchini (1997, p. 79).

3.6 Recovering the State Membership Distribution

We use the filtering approach (Hamilton 1989), to determine the customers' membership probability in any state at any given purchase event. The filtering probability that customer i is in state s at purchase event j conditioned on the customer's history of choices is given by:

$$P(S_{ij} = s | \mathbf{Y}_{i1}, \mathbf{Y}_{i2}, \dots, \mathbf{Y}_{ij}) = \boldsymbol{\pi}_i \mathbf{M}_{i1} \boldsymbol{\Omega}_{i,1 \rightarrow 2} \mathbf{M}_{i2} \dots \boldsymbol{\Omega}_{i,j-1 \rightarrow t,s} L_{ij|s} / L_{ij}, \quad (20)$$

where, $\mathbf{Q}_{i,j-1 \rightarrow j,s}$ is the s^{th} column of the transition matrix $\mathbf{Q}_{i,j-1 \rightarrow j}$, and, L_{ij} is the likelihood of the observed sequence of joint decisions up to purchase event j from Equation (19).

4. Empirical Application

In this section, we apply our model to two years of transactional data from a B2B retailer of metals. We aim to investigate the short- and long-term impact of the retailer's pricing decisions on the buying decisions of customers. We begin by describing the characteristics of the data. We then estimate the multi-decision HMM model and assess the impact and duration of the pricing effects. Finally, we use the estimated parameters to conduct "what-if" simulations to investigate how various managerial metrics are impacted by alternative pricing policies.

4.1 Data

The dataset comes from an east coast local B2B metals retailer that specializes in supplying aluminum to industrial clients such as machine shops and fabricators that operate in its

geographical trading area. The dataset consists of customer-level transactions over a 21 month period from January 2007 to September 2008. As the purchasing process in B2B settings can be complex, it is noteworthy to first describe the typical flow of a purchase opportunity. After a need for a certain quantity arises at a certain point in time, the customer places an order with the retailer. An order can originate either with a request for a price quote (usually via the phone or via fax) or by a direct order without asking for a price quote (hereafter “direct order”). For example, a typical direct order may be received by fax in the morning 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 (hereafter “indirect order”) the firm bids for the customer’s business, and can only “win” the business if the customer accepts the quoted price.² Our dataset is unique as it includes not only completed transactions but also lost transactions involving bids that were lost, thus allowing us a better grasp of customer price sensitivity. The company has a very 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. These two facts preclude the maintenance of a price list, therefore, as typical in this industry, the company determines the price to charge or bid on a case by case basis.³

We use the term “purchase event” to define an indirect or a direct order. For every purchase event we observe: the date that the customer initiates the order; the quantity requested; whether the customer ordered directly or requested a quote; if a quote, whether the customer purchased or not; and the price the firm quoted and/or charged. To account for the variation in order quantities, and because of the large number of SKUs, we normalize prices to price per pound.

² 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.

³ As is typical of most customer relationship management (CRM) datasets in B2B settings, our dataset does not include 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 (bids or orders) in the data period. Tables 1 and 2 contain the summary statistics of the data. The data contains 33,925 purchase events (bids and/or orders) totaling 16 million lbs (approximately 7,250 metric tons) of aluminum. Of these, 53% (18,038 purchase events) were direct orders in 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, customer 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% (7,501 orders) were won by the company.

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

On average, a customer in our sample engages in 23.6 purchase events during the span of 21 months. An average purchase event involves 457 lbs of aluminum with an average price of 3.24 \$/lb.

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 2 show that 1) direct orders tend to be smaller, possibly suggesting 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) heterogeneity in the average price paid is relatively low 4) customers exhibit different degrees of caution when it comes to ordering behavior, as some are more prone to order directly without asking for a price quote than others, suggesting that customers may vary in their attitudes and latent relationships with the firm; 5) the fact that about half of the orders are direct, which result in a sale regardless of the price charged, suggests that the firm might be able to charge “any” price on such orders. As we demonstrate later, our dynamic model allows us to estimate the possibly negative long-term consequences of such “over-charging” pricing behavior.

Applying the modeling framework described in Section 3 to our dataset, we wish to answer the following questions: 1) How does pricing affect each dimension of the buying process? 2) What are the short- and long-term effects of observed price? 3) How can the firm use pricing to change the behaviors of its customers? 4) What are the implications of the firm’s pricing policy for long-term profitability?

4.2 Variables Description

In this section we describe the variables that potentially impact the different components of the customer decision.

4.2.1 Asymmetric Reference Price Effects

First, we specify how we account for reference price effects. We assume that customers pay closer attention to the prices of larger orders relative to smaller ones. We therefore define the reference price as a quantity weighted average of the customer’s past observed prices (in \$/lb). The reference price is constructed internally for each customer and is updated at every time period to include the most recently observed price. The reference price for customer i , at purchase event j , can be written as:⁴

⁴ We compared the reference price in Equation (21) against several alternative reference prices such as simple average of past purchases, the price in the last purchase event and an external reference price from the LME spot price. The above reference price is both theoretically sound and provides the best fit for the data.

$$reference_price_{ij} = \frac{\sum_{k=1}^{k=j-1} quantity_{ik} \times price_lb_{ik}}{\sum_{k=1}^{k=j-1} quantity_{ik}} \quad (21)$$

We define asymmetric reference price effects (Putler 1992) via “gain” and “loss” variables computed from the difference between the current price and the reference price. Specifically, we define:

1. $gain_{ij}$: The difference between the price per pound in the current purchase event and the customer’s internally constructed reference price as calculated in Equation (21), for customer i at purchase event j , when the current price is lower than the reference price (i.e., the current price is perceived as gain).

$$gain_{ij} = \begin{cases} reference_price_{ij} - price_lb_{ij} & , \text{ If } price_lb_{ij} < reference_price_{ij}, \\ 0 & , \text{ Otherwise.} \end{cases}$$

Similarly we define:

2. $loss_{ij}$: The difference between the price per pound in the current purchase event and the customer’s internally constructed reference price for customer i at purchase event j , when the current price is higher than the reference price (i.e., the current price is perceived as loss).

$$loss_{ij} = \begin{cases} price_lb_{ij} - reference_price_{ij} & , \text{ If } price_lb_{ij} > reference_price_{ij}, \\ 0 & , \text{ Otherwise.} \end{cases}$$

4.2.2 Additional Variables

1. lme_j : To control for fluctuations and trends in aluminum prices, we overlay the transactions dataset with the daily aluminum spot prices on the London Metal Exchange (LME). This variable denotes the level of aluminum spot price (dollars per metric ton) on LME.

2. $lme_volatility_j$: 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 .
3. $quantity_{ij}$: The quantity (pounds of aluminum) that the customer i requested on the purchase event j .
4. t_{ij} : The time in weeks between customer i 's purchase event $j-1$, and purchase event j .

4.3 Model Specification

We now describe how the above mentioned variables are incorporated in the general model described in section 3.

4.3.1 The State Dependent Decisions

1. *Quantity Decision*: We let the lognormal distribution of the quantity ordered be a function of past internal reference price effects and the external price level and volatility in metal prices. We expect that price loss [gain] in the previous purchase event will lead to lower [higher] quantity ordered in the current purchase event. Accordingly, the vector \mathbf{x}_{ij} in Equation (9) includes the covariates: $gain_{ij-1}$, $loss_{ij-1}$, lme_j , and $lme_volatility_j$.
2. *Purchase-event Timing Decision*: We let the log-logistic distribution of inter-purchase-event time be a function of the previous quantity ordered and past internal reference price effects. We expect that price loss [gain] in the previous purchase event would lead to longer [shorter] inter-purchase-event time, and because of the role of inventory, the larger the quantity ordered on the previous purchase-event, the longer the inter-purchase-event time will be. Accordingly, the vector \mathbf{x}_{ij} in Equation (7) includes the covariates: $gain_{ij-1}$, $loss_{ij-1}$, and $quantity_{ij-1}$.
3. *Quote Request Decision*: We assume that the probability of requesting a quote (rather than ordering directly), is conditioned on the time since the last purchase event (t_{ij}) and the quantity ordered ($quantity_{ij}$). We also allow the market conditions to affect the level of caution that a customer exercises in ordering. We expect that customers would avoid ordering directly when the quantity purchased is large, when it has been a long time since

the last purchase event, or when the market conditions are volatile. When the LME metal price is high, however, we expect customers to be less cautious in his ordering behavior because, firstly, a high commodity price can indicate a benign economic environment, and secondly, a high LME price might be presented to customers as a high external reference price that will make the seller's current price more attractive. Furthermore, the decision to request a quote also depends on the internal reference price effects that are experienced directly on the previous purchase event, and as we will describe later, also indirectly via the customer's latent state. Thus, the vector \mathbf{x}_{ij} in Equation (15) includes the covariates: t_{ij} , $quantity_{ij}$, $gain_{ij-1}$, $loss_{ij-1}$, lme_j , and $lme_volatility_j$

4. *Quote Acceptance Decision*: Once a quote has been requested the customer's probability of accepting the quote may depend on the price that is bid, the quantity ordered, and the time elapsed since the last purchase event. Furthermore, we test for interaction effect between the quantity ordered and the reference price effects. We predict that that likelihood of accepting a bid will be higher when the quantity ordered is small, purchases are frequent, and the customer experiences either a gain (or an absence of a loss) in reference price. We also expect the gain and loss effects to be magnified for larger orders. Furthermore, the decision might be marginally affected by the market condition. Accordingly, the vector \mathbf{x}_{ij} in Equation (15) includes the covariates: t_{ij} , $quantity_{ij}$, $gain_{ij}$, $loss_{ij}$, $gain_{ij} \times quantity_{ij}$, $loss_{ij} \times quantity_{ij}$, lme_j , and $lme_volatility_j$.

The above set of four decisions characterizes customer behavior given the customer's latent state. Thus, a jointly constructed set of equations for the four decisions are defined for each state of the HMM following Equation (2). We model the interdependency between the quantity ordered and the inter-purchase-event time decisions via the Frank copula with dependency parameter θ_{qt} (see in Equation (6)). Similarly, as bid acceptance can only take place when a quote is requested, bid acceptance behavior and quote request behaviors are modeled jointly using a copula framework for selectivity. These binary decisions follow the standard logit specifications in equation (15), and are jointly modeled via a Frank copula with parameter θ_{bw} (see equation (14)).

To capture the dynamics in buying behavior we specify a two-state HMM. To ensure identification of the states, we define the intercepts in the quote request equation (Equation (15)) to be increasing in state's cardinality, such that (at the mean of the covariates) the likelihood of requesting a quote (rather than ordering directly) is higher in State 2 relative to State 1. Accordingly, all covariates are mean-centered in the model. Our two-state HMM specification is consistent with a "relaxed" state characterized by a trustful attitude towards the seller, and a more "vigilant" state, in which the customer is more cautious about ordering from the seller.⁵

4.3.2 The Non-Homogenous Transition Matrix

We include the asymmetric reference price effects, namely $gain_{ij-1}$ and $loss_{ij-1}$, in the transition probabilities specified in Equation (1). This specification suggests that the "joy of gain" and the "pain of loss", experienced on the previous purchase event, can affect the customer's evaluation (or re-evaluation) of the relationship with the firm, and trigger a transition to a different buying behavior state. Such a specification allows us to capture the possibly long-term effect of reference prices. Accordingly, the vector \mathbf{x}'_{ij} in Equation (1) includes the covariates: $gain_{ij-1}$, and $loss_{ij-1}$.

4.3.3 Price-Generating Mechanism Specification

We have data on the seller's wholesale cost, that is, the cost that the seller pays to the mills for a particular wholesale lot of metal. Conversations with the management team reveal that each customer's order is retrieved from a wholesale lot, and that the salesperson would focus on the wholesale cost when quoting the price to the customer. Thus, we include the wholesale cost as z_{ij} in the price-generating equation in Equation (16).

4.3.4 Heterogeneity Specification

⁵ For model parsimony and ease of interpretation we do not extend our model to a three or more HMM states.

To distinguish dynamics from cross-customer heterogeneity (Heckman 1981), we specify the model parameters that govern the HMM dynamics and those that are related to reference price effects to be customer-specific. Specifically, we allow the intercepts of the four decisions equations 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_j , $lme_volatility_j$.) to vary randomly across customers. Additionally, all the parameters that influence the initial state membership and the customer dynamics in the transition matrix are assumed to be random across customers. For parsimony, all other parameters in the model are assumed to be constant across customers (see column 3 in Table 4 for a full list of the random- and fixed-effect parameters).⁶

4.4 Model Estimation

In this section we briefly discuss the estimation procedure. We use a hierarchical Bayesian approach for inference and estimate the parameters using Markov Chain Monte Carlo (MCMC) methods. Given that many of the full conditional distributions in our model do not have a closed form, we use a combination of random-walk Metropolis-Hasting steps and Gibbs sampling steps to draw parameter values. We follow Netzer et al (2008) and use the HMM likelihood function in Equation (19) directly in formulating our MCMC methods. The inherent complexity of the HMM often leads to significant autocorrelation between successive draws of the MCMC chain. We therefore use the adaptive Metropolis procedure in Atchadé (2006), which significantly improves mixing and convergence. See Appendix for a detailed description of the MCMC estimation procedure, along with the prior and posterior distributions.

We use the first 18 months of data to estimate the model, and the last three months for validation purposes. We ran the MCMC methods for a million iterations. The first 700,000 iterations were used as a “burn-in” period, and the remaining 300,000 iterations were used for parameter inference.

⁶ For parsimony reasons, we do not let the Frank copula correlation parameter (θ_{qt} and θ_{wb}) and the scale parameter of the lognormal distribution (σ) to vary by states.

4.5 Results and Marketing Implications

In this section, we first describe the benchmark models we used to assess the relative model fit and predictions of our proposed model, and the parameter estimates. We then use the estimated parameters as input into a set of simulation to disentangle the short- and long-term effects of the pricing decisions on the four decision variables for the customer. We also derive implication for customer lifetime value and long-term profitability.

4.5.1. Benchmark Models

We compare the fit and predictive ability of our proposed two-state non-homogenous HMM model in which the multivariate outcomes are correlated via copulas, with that of four benchmark models. The benchmark models vary with respect to the degree of heterogeneity, the degree of dynamics (i.e. via HMM specification to capture latent behavior dynamics, and via the reference price specification to capture dependence on past prices), the degree in which price endogeneity is accounted for, and the interdependence between the components of the customer's buying process that they accommodate. We use the first 18 months in our dataset to calibrate the model and the last three months for validation. As for our proposed model, we assume that in all the benchmark models, quantities follow a lognormal distribution, inter-purchase-event times follow a log-logistic distribution, and the quote request and bid acceptance decisions both follow logistic distributions. For all models, we use a Bayesian approach for inference.

Benchmark Model 1 (no heterogeneity). This model is identical to the proposed model in its structure except that all its parameters are considered to be invariant across customers. A comparison of this model with the proposed model allows us to assess the importance of capturing heterogeneity in the context of our empirical application.

Benchmark Model 2 (independent decisions, no HMM). There are two sources for interdependence among the four decision components of the customer. First, the HMM model

structure creates dependence among the components. Second, conditioned on the state, we allow interdependence among the components using copulas. Thus, to create a benchmark model that assumes independent decisions, we estimate a single state heterogeneous model in which the quantity, inter-purchase-event time, quote request and bid acceptance distributions are all estimated independently. This benchmark model allows us to examine the value of joint decisions modeling and of capturing relationship dynamics.

Benchmark model 3 (no dynamics). This model has a single state (i.e., no HMM) and uses actual prices instead of reference prices to remove dependence on past prices. This allows us to assess the value of capturing dynamics.

Benchmark model 4 (not counting for price endogeneity). This model is identical to the proposed model in its structure except for the Bayesian control function component. This will allow us to first assess the value of accounting for price endogeneity, and secondly, examine the extent of the bias of sensitivity if price endogeneity is not properly accounted for.

4.5.2. Model Fit and Predictive Ability

We use log marginal density (LMD) and the Deviance Information Criterion (DIC) to assess model fit. Both measures are calculated from the output of the Metropolis Hastings sampler. In addition, we use the validation log-likelihood to assess overall predictive ability. Apart from these overall measures of model performance, we also compute component-specific measures. Specifically, we calculate the root mean squared error (RMSE) and the mean absolute deviation (MAD) between the predicted and observed values of the timing, quantity, quote request, and bid acceptance decisions and the hit-rates for the binary quote request and conditional bid acceptance decisions in both the calibration and validation samples.

Table 3 presents the fit and predictive ability measures for the five models. The results suggest that the proposed model dominates the benchmark models on the overall measures and outperforms these other models on each of the components, in both the calibration and validation

periods. The relatively poor fit and poor predictive ability of the no-heterogeneity model suggests that the degree of customer heterogeneity in our data is substantial. Similarly, by modeling purchasing decisions jointly and by incorporating dynamics via HMM and via the reference price specification we improve the representation and prediction of buying behavior. Furthermore, accounting for price endogeneity only marginally improves the model fit and prediction, which is not surprising as this business is traditionally based on a “cost-plus” pricing practice, and an regression of price on wholesale cost yields a R-squared of 0.84. We now discuss the parameter estimates of the proposed model.

Table 3: Model Selection and Predictive Validity

	Proposed model		Benchmark 1		Benchmark 2		Benchmark 3		Benchmark 4	
	Heterogeneity		Heterogeneity		Heterogeneity		Heterogeneity		Interdependence	
	Interdependence		Interdependence		Interdependence		Interdependence		Heterogeneity	
	Dynamics (HMM + Reference Price)		Dynamics (HMM + Reference Price)		Dynamics (Reference Price Only)		Dynamics (Reference Price Only)		Dynamics (HMM + Reference Price)	
	Control Function		Control Function		Control Function		Control Function		Control Function	
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
Bid 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
Bid 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
Bid acceptance	0.41	0.49	0.63	0.76	0.40	0.48	0.40	0.49	0.41	0.49

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

4.5.3. The HMM Parameter Estimates

Tables 4 and 5, contain the posterior means, standard deviations, and 95% posterior intervals for the parameters of the proposed model. Table 4 reports the state-specific parameters for each component of the buying process, whereas, Table 5 presents the HMM transition matrix and initial state distribution parameters, the dependence parameters for the two copulas, and the other parameters for the distributions of the decision variables.

Interpreting the HMM States. We focus on Table 4 to characterize the two states of the HMM. As all covariates are mean-centered the intercept of each equation captures the average tendency to behave in a particular manner. Hence, comparing the parameters for the quote request and bid acceptance decisions across the two states, we see that purchase events in State 2 have a higher chance of involving indirect orders and are associated with a lower probability of bid acceptance compared to events in State 1. This suggests that State 2 represents a more cautious behavior towards ordering. Further supporting this characterization is the fact that, relative to events when the customer is in State 1, events when the customer is in State 2 exhibit a higher sensitivity to reference price effects, stronger loss aversion in the inter-purchase-event time and bid acceptance decisions, and stronger gain seeking in the quantity decision. Customers who are in State 2 at a given event also tend to be more responsive to the general economic environment (e.g. the level and the volatility in the LME), and pay closer attention to the characteristics of the order itself (quantity and inter-purchase-event time) in making their decisions. We therefore call State 2 the “*vigilant*” state.

In contrast, customers in State 1 tend to order directly without asking for a price quote, are more likely to accept the seller’s bid and are generally less price sensitive than customers in State 2. Thus, these customers are more relaxed in their relationship with the seller. Accordingly, we label State 1 as the “*relaxed*” state. Furthermore, customers in the relaxed state, instead of exhibiting loss aversion, exhibit the diametrically opposing behavior of gain “seeking”

in the bid acceptance decision. That is, for customers in the relaxed state the positive impact of gains is stronger than the negative effect of losses, suggesting that customers when in the relaxed state are focus more on gains than on losses. This finding is consistent with regulatory-focus theory (Higgins 1997). Customers in the vigilant state are more “prevention focused,” focusing on avoiding losses, and engage in a more vigilant and scrutinizing relationship with the firm. On the other hand, customers in the relaxed state are more “promotion-focused”, concentrate more on gains and price savings, and adopt a more favorable buying relationship with the seller.

Trust has been suggested as an important facet in B2B buyer-seller relationships (Morgan and Hunt 1994). Trust may serve as an underlying mechanism for the relaxed and vigilant states found in our empirical application. Specifically, one may argue that customers in the relaxed state tend to exhibit buying patterns that are consistent with a more trustful attitude towards the seller relative to customers in the vigilant state.

Reference Price Effects. The results show that B2B customers respond to both internal and external reference prices. Internally, price gains would make the current price more attractive and would hence lead to more favorable ordering behaviors, and price losses would conversely lead to more cautious behaviors. Externally, the results suggest that a high level of metal price on the LME could be interpreted as a high reference point, making the seller’s current price more attractive and thus leading to more favorable ordering behaviors. A high volatility of the LME captures uncertainty in the current economic environment, resulting in a more cautious approach to ordering. Once again, both internal and external reference price effects are stronger for customers currently in the vigilant state.

Customer Dynamics. According to the HMM, customers can transition over time between the two states, thus exhibiting dynamic behavior. The parameter estimates in Table 5 and their transition matrix representation in Table 6, depict these dynamics. As can be seen in the middle matrix in Table 6, the states are relatively sticky. The vigilant state is particularly sticky. Assuming the customer is priced at its reference price, the likelihood that a customer in the vigilant state remains in this state in the next period is 91%. However, the seller can use its

pricing policy to affect customer transitions between the two states. The transition matrix parameter estimates in Table 5, and the left matrix in Table 6, show that a price “gain” in the previous purchase event increases the likelihood of a shift from the vigilant to the relaxed state by 2% and increases the likelihood of the customer remaining in the relaxed state by 3%. This suggests that when customers perceive that they are treated well, they may transition to or remain in a state of more favorable relationship with the firm. On the other hand, the right matrix in Table 6 shows that a price “loss” in the previous purchase event increases the likelihood that the customer will transition to the vigilant state by 7% and increases the likelihood of remaining in the vigilant state by 3.5%, which in turn, translates to further increased sensitivity to losses. As shown in Tables 5 and 6, the average loss aversion in the effect of reference prices on the transition between the states ranges from 1.5 to 2.5. That is, the increase in likelihood of moving to or staying in a vigilant state due to a price loss is one and a half to two times the corresponding decrease in the likelihood of moving to or staying in a vigilant state because of an identical price gain.

Table 4: State Dependence Estimates for the Four Decisions*

	Parameter	Aggregation Level**	Mean	Std Dev.	2.50%	97.50%
<i>Quantity</i>						
State1	Intercept	RE	1.823***	0.039	-1.897	-1.751
	Gain(t-1)	RE	0.004	0.021	-0.037	0.043
	Loss(t-1)	RE	-0.035	0.013	-0.061	-0.010
	LME(t)	RE	0.029	0.129	-0.220	0.268
	LME_Volatility(t)	RE	-0.782	0.486	-1.720	0.122
State2	Intercept	RE	-1.523	0.033	-1.587	-1.461
	Gain(t-1)	RE	0.133	0.019	0.096	0.169
	Loss(t-1)	RE	-0.096	0.014	-0.123	-0.070
	LME(t)	RE	2.813	0.539	1.774	3.815
	LME_Volatility(t-1)	RE	-3.350	1.054	-5.383	-1.390
<i>Inter-purchase-event time</i>						
State1	Intercept	RE	0.977	0.029	0.919	1.035
	Quantity(t-1)	FE	-0.009	0.029	-0.067	0.048
	Gain(t-1)	RE	-0.062	0.014	-0.091	-0.034
	Loss(t-1)	RE	0.046	0.014	0.019	0.074
State2	Intercept	RE	0.854	0.031	0.793	0.915
	Quantity(t-1)	FE	0.031	0.021	-0.010	0.072
	Gain(t-1)	RE	-0.016	0.016	-0.047	0.016
	Loss(t-1)	RE	0.037	0.010	0.016	0.057

<i>Quote request vs. direct order behavior (I= quote request)</i>						
State1	Intercept	RE	-1.125	0.034	-1.190	-1.062
	Quantity(t)	FE	0.418	0.097	0.230	0.599
	Interpurchase time(t)	FE	0.030	0.004	0.023	0.036
	Gain(t-1)	RE	-0.081	0.026	-0.132	-0.033
	Loss(t-1)	RE	0.096	0.024	0.050	0.140
	LME(t)	RE	-0.981	0.878	-2.674	0.652
	LME_Volatility(t)	RE	1.180	1.056	-0.858	3.144
State2	Intercept	FE	1.037	0.080	0.881	1.186
	Quantity(t)	FE	2.090	0.575	0.981	3.160
	Interpurchase time(t)	FE	0.004	0.007	-0.009	0.017
	Gain(t-1)	RE	-0.312	0.015	-0.340	-0.285
	Loss(t-1)	RE	0.768	0.017	0.735	0.799
	LME(t)	RE	-1.880	0.589	-3.017	-0.784
	LME_Volatility(t)	RE	3.349	1.293	0.854	5.753
<i>Bid acceptance vs. rejection behavior (I= accept)</i>						
State1	Intercept	RE	0.214	0.015	0.186	0.242
	Quantity(t)	FE	-0.220	0.032	-0.281	-0.162
	Interpurchase time(t)	FE	0.002	0.017	-0.031	0.033
	Gain(t)	RE	0.211	0.015	0.183	0.239
	Loss(t)	RE	-0.180	0.016	-0.212	-0.150
	Quantity(t) x Gain(t)	RE	0.087	0.018	0.051	0.121
	Quantity(t) x Loss(t)	RE	-0.047	0.016	-0.079	-0.017
	LME(t)	RE	0.212	0.500	-0.752	1.141
	LME_Volatility(t)	RE	-0.362	1.056	-2.399	1.602
State2	Intercept	RE	0.096	0.049	0.002	0.187
	Quantity(t)	FE	-0.530	0.046	-0.618	-0.444
	Interpurchase time(t)	FE	-0.032	0.008	-0.046	-0.017
	Gain(t)	RE	0.133	0.015	0.104	0.161
	Loss(t)	RE	-0.211	0.020	-0.250	-0.173
	Quantity(t) x Gain(t)	RE	0.285	0.015	0.256	0.312
	Quantity(t) x Loss(t)	RE	-1.518	0.018	-1.554	-1.484
	LME(t)	RE	0.747	0.517	-0.251	1.708
	LME_Volatility(t)	RE	-1.294	1.056	-3.331	0.670
<i>Price equation (price_lb)</i>						
	Intercept	FE	3.540	0.070	3.404	3.671
	Std dev	FE	1.104	0.254	0.613	1.577

*Posterior means and standard deviations are calculated across retained MCMC iterations

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

*** Bold estimates are significant at the 95% posterior interval

Table 5: 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(t-1)	RE	0.052	0.018	0.017	0.086
	Loss(t-1)	RE	-0.105	0.018	-0.140	-0.072
State2	intercept,	RE	2.514	0.043	2.431	2.594
	Gain(t-1)	RE	-0.505	0.018	-0.539	-0.471
	Loss(t-1)	RE	0.712	0.019	0.675	0.748
Initial membership						
	Pai(t)	RE	0.697	0.030	0.639	0.752
Distributional Parameters						
	Std dev. for the quantity model, log scale	FE	0.126	0.052	0.027	0.222
	State 1 Shape parameter for interpurchase time, log scale	FE	0.116	0.019	0.079	0.152
	State 2 Shape parameter for interpurchase time, log scale	FE	0.079	0.012	0.055	0.102
	Frank copula parameter for interpurchase and quantity	FE	0.606	0.178	0.264	0.937
	Frank copula parameter for quote request and bid acceptance	FE	-10.023	0.563	-11.110	-8.975

Table 6: Posterior Mean of Transition Matrix Across Customers

	10% Price Decrease		Average price		10% Price Increase	
	Relaxed (t+1)	Vigilant (t+1)	Relaxed (t+1)	Vigilant (t+1)	Relaxed (t+1)	Vigilant (t+1)
Relaxed (t)	0.895	0.105	0.864	0.136	0.795	0.205
Vigilant (t)	0.096	0.904	0.078	0.922	0.043	0.957

Interdependent Decisions. The parameters of the Frank Copula in Table 5 indicate significant and substantial positive correlation between quantity and inter-purchase-event time, and negative correlation between quote request and bid acceptance. The fact that these correlations are significant in magnitude is consistent with the improvements in fit and prediction of our model relative to a model that does not capture interdependence among these components (see Benchmark model 2 in Table 3). The negative correlation between the quote request and

bid acceptance rates indicates selectivity (Heckman 1981) and indicates that factors that increase the likelihood of quote requests also result in lower bid acceptance rates, a kind of a “double jeopardy” for the firm. The positive correlation between the inter-purchase-event times and quantity decisions can be partly explained by the presence of inventories. Finally, the shape parameters for the inter-purchase-event time log-logistic distribution in Table 5, suggest that the baseline hazard for a purchase first increases and then decreases with time.

Investigating Price Endogeneity. Table 7 lists the estimated price parameters under the proposed full model vs. the benchmark model 4, to assess the extent of price endogeneity and its impact on the parameters. The bottom of Table 7 shows the correlation of each one of the four decisions, with regard to price, inferred from the terms in the control function to correct for endogeneity⁷. The results show that with respect to the unobserved characteristics in the price generating mechanism, quantity and bid acceptance behavior exhibit mild positive correlations, whereas inter-purchase-event time and quote request behavior exhibit mild negative correlations. For example, an unobserved factor that would raise the price would also tend to raise the quantity ordered, shorten the inter-purchase-event time, lower the probability of a quote request, and increase the probability of a bid acceptance. These correlations structure might imply that 1) despite the predominant cost-plus pricing practice, the seller does occasionally target price-insensitive[sensitive] customers by charging them higher[lower] prices, and that 2) the seller charges customers higher prices during an economic expansion, but the same economic boom would create more demand, resulting in a seemingly favorable purchasing behaviors. Hence, not properly accounting for price endogeneity would overestimate the effects of price gains, and underestimate the effects of price losses, which is what we observe in Table 7. Model comparison between the proposed model with the model uncorrected for endogeneity results in a better fit.

⁷ i.e. residuals and error components of the control function. For parsimony of composition, these terms are left out of Table 7 but are available from the authors upon request.

Table 7: Price Parameters - Bayesian Control Function Approach

	Parameter	With Control Function	Std Dev.	Uncorrected	Std Dev.
<i>Quantity</i>					
Relaxed	Gain(t-1)	0.004	0.021	0.007	0.016
	Loss(t-1)	-0.035	0.013	-0.028	0.011
Vigilant	Gain(t-1)	0.133	0.019	0.150	0.013
	Loss(t-1)	-0.096	0.014	-0.072	0.010
<i>Inter-purchase-event time</i>					
Relaxed	Gain(t-1)	-0.062	0.014	-0.079	0.013
	Loss(t-1)	0.046	0.014	0.047	0.013
Vigilant	Gain(t-1)	-0.016	0.016	-0.025	0.014
	Loss(t-1)	0.037	0.010	0.032	0.008
<i>Quote request vs. direct order behavior (I= quote request)</i>					
Relaxed	Gain(t-1)	-0.081	0.026	-0.122	0.025
	Loss(t-1)	0.096	0.024	0.118	0.019
Vigilant	Gain(t-1)	-0.312	0.015	-0.390	0.018
	Loss(t-1)	0.768	0.017	0.689	0.018
<i>Bid acceptance vs. rejection behavior (I= accept)</i>					
Relaxed	Gain(t)	0.211	0.015	0.253	0.014
	Loss(t)	-0.180	0.016	-0.182	0.017
Vigilant	Gain(t)	0.133	0.015	0.145	0.017
	Loss(t)	-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		
Bid acceptance	Relaxed	-0.108	0.059		
	Vigilant	-0.144	0.061		
Log-likelihood at convergence			-62,531		-62,956

4.5.4. Disentangling the Effects of Pricing

Isolating the impact of pricing is difficult in our model because of the complex feedback structures among the four customer decisions. These decisions are interrelated via explicit conditionality and via the copula approach. Furthermore, the use of a HMM, adds additional complexity in the interpretation of parameters as price can have both a short-term impact conditional on a state and a long-term impact via the transition matrix. In addition, the decomposition of the reference price effects into gains and losses results in an asymmetric impact of prices, both in the short and long run. Thus, the interpretation of the price parameters in Tables 4 and 5 is not immediately obvious. To analyze the overall impact of pricing on each decision component, and to compute the short- and long-term price elasticities, we performed a series of “what-if” simulations.

We computed the price elasticities for each of the four decision components, as well as for the HMM state memberships. Elasticities for price increases are computed numerically using the formula $e_{y,p}^+ = (\Delta y_+ / y) / (\Delta p_+ / p)$, where y corresponds to the average level for each decision variable, p corresponds to the average unit price, Δp_+ corresponds to increase in unit price, and Δy_+ corresponds to the change in the level of the decisions in the presence of price increase. Similarly, elasticities for price decreases are calculated using the formula $e_{y,p}^- = (\Delta y_- / y) / (\Delta p_- / p)$, where Δp_- corresponds to decrease in unit price, and Δy_- corresponds to the change in the level of the decisions in the presence of price decrease. To compute the changes in each of the above decision levels (Δy_+ and Δy_-), we use the parameter estimates from the HMM to simulate the effect of increasing or decreasing the unit price by a one-time shock of 10% across all customers. We first compute the base levels for each decision when there is no price shock (y_0) by simulating the next 20 purchase-event periods at the current price. Next, we compute the outcomes for each decision during the next 20 purchase-event periods following a one-time 10% price hike (y_+), and a one-time 10% price drop (y_-). Thus,

$\Delta y_+ = y_+ - y_0$, and $\Delta y_- = y_- - y_0$. To compute the short-term elasticities, we consider the effect of price hike and price drop in the first purchase-event period only; whereas, the long-term elasticities are computed using the effects in the remaining 19 periods. The simulated short- and long-term elasticities for each of the four components of the buying process and the vigilant state membership probability are reported in Table 7. Similarly, Figures 1 and 2 depict the asymmetric percentage increase and decreases in each decision component over the simulated 20 periods following a single period 10% price increase or decrease.

Table 8: 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
Bid 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

Figure 1: Duration of the Asymmetric Price effects on Quote Request, Bid Acceptance, Quantity, and Inter-purchase-event Time

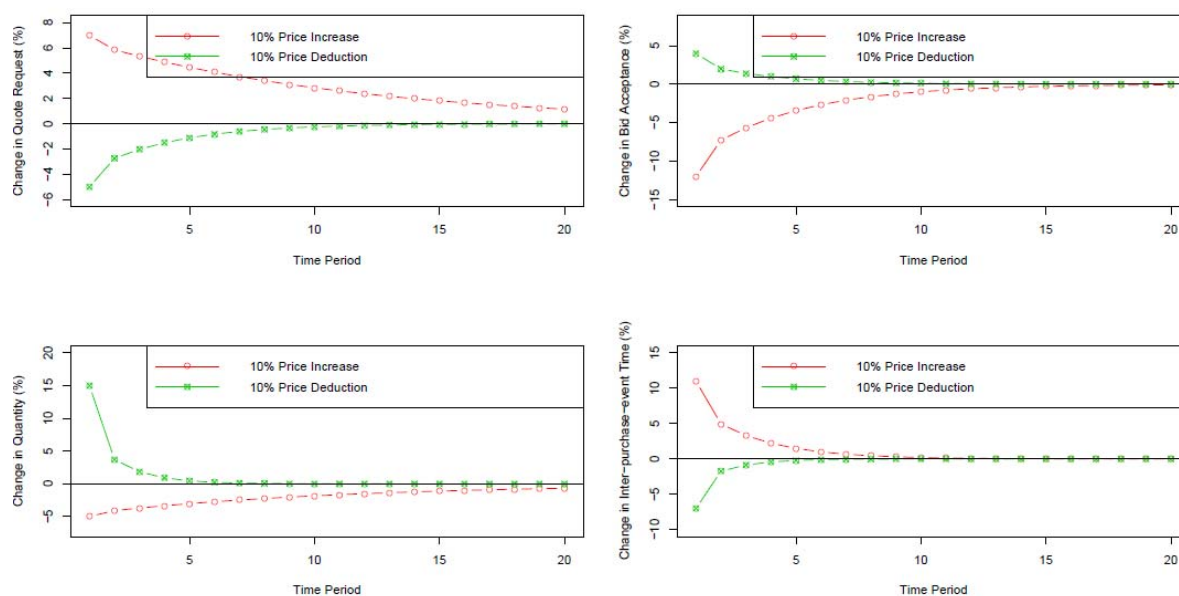



Figure 2: Duration of the Asymmetric Price Effects on Vigilant State Membership



Several insights can be derived from Table 8 and Figures 1 and 2. First, for all four decisions and for the vigilant state membership, the magnitudes of the long-term elasticities are generally much larger than those for the short-term elasticities. The average short-term quantity elasticities are consistent with those of Jedidi et al. (1999), and fall within the range of the price elasticities found by Tellis (1988) in a comprehensive meta-analysis of price elasticities. On average, the short-term price elasticities are only 10%-50% of the total price elasticities. Second,

factors that are directly related to customers' attitudes and relationships with the seller (i.e., quote request and vigilant state membership) exhibit stronger long-term elasticities relative to other factors. Third, consistent with loss aversion, we find that the detrimental effects caused by price increases (loss) are stronger than the positive effects caused by price decreases (gain) for all decisions but quantity. The quantity decision exhibits gain seeking. Fourth, Figures 1 and 2 further depict that the negative effects of a price hike (loss) persist longer, and take a longer time for customers to adapt to, than the positive effects of price drop (gain). We find support for this result in the hedonic adaption literature, which states that individuals' adaptation to improvements is faster than their adaptation to deteriorations (Frederick and Lowenstein 1999; Strahilevitz and Lowenstein 1998). To the best of our knowledge, this is the first paper to investigate the long-term effects of asymmetric reference price effects, and demonstrate that effect of price losses are not only stronger in magnitude than the effect of price gains, but also last longer.

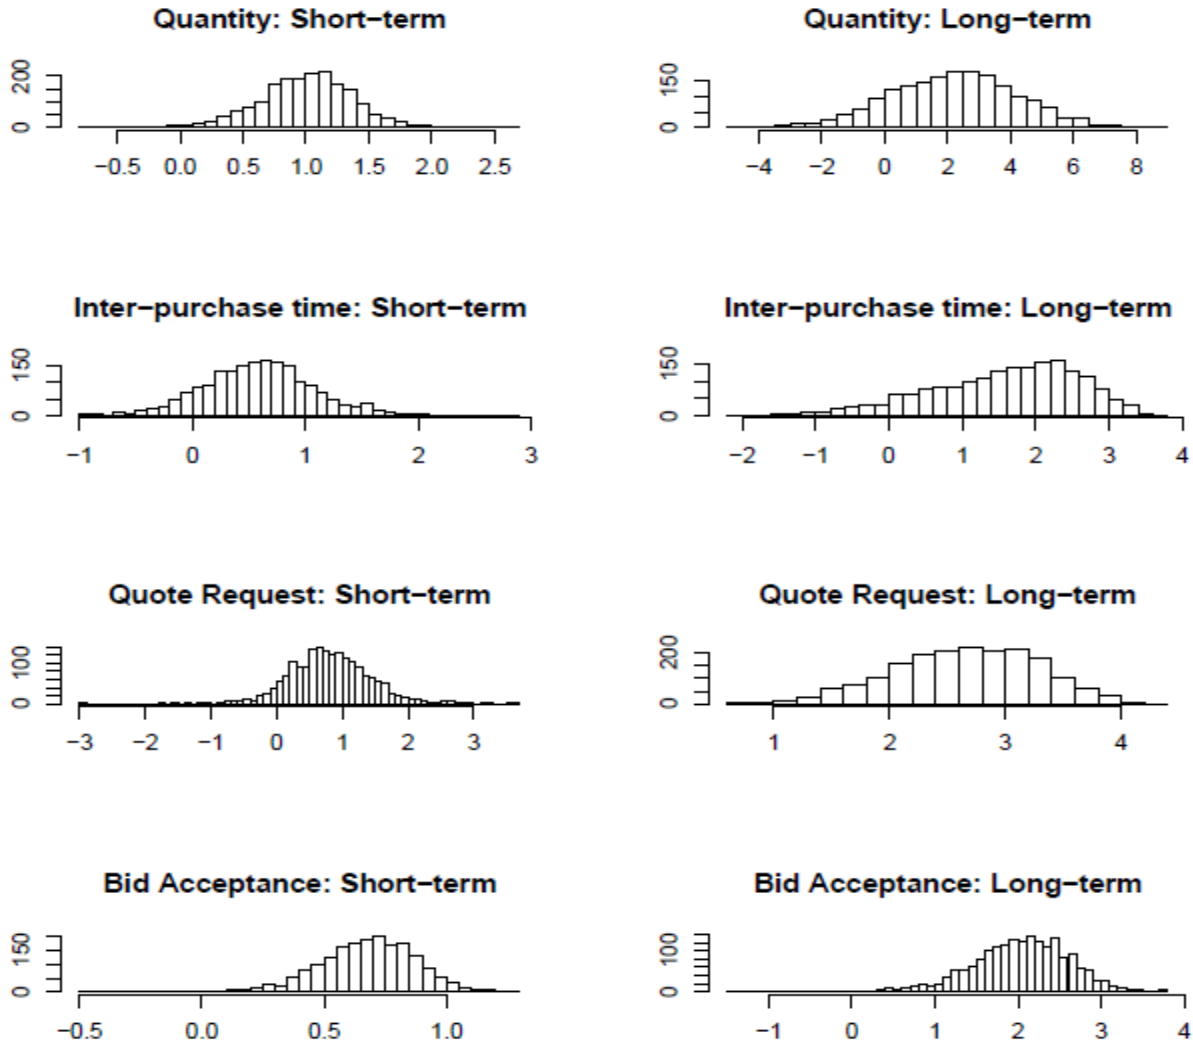
Overall, these results suggest that 1) in the context of B2B, studies that consider only the short-term effects of pricing can significantly and substantially underestimate the overall impact of pricing, and 2) studies that do not separately examine the, possibly asymmetric, price effects caused by gains and losses can miss important insights into how customers behave, both in short and long run. From a managerial perspective, firms need to be very wary of increasing prices and hence causing losses for the customer. Such a pricing policy can not only have an immediate negative impact, but also the damage the relationship with the customer in the long run, and recovery may be slow. Next, we demonstrate how price effects on the different components of the decision can be linked together to allow for an assessment of the overall impact of pricing on firm's profitability through customer lifetime value.

4.5.5 Heterogeneity across Customers

Our modeling framework incorporates full heterogeneity across customers in the dynamics and reference price effect parameters. For example, a set of transition matrices, as depicted in Table 6, can be derived for each customer. Similarly, the MCMC estimation framework provides us

with the complete distribution of elasticities across customers for each of the average price elasticities calculated previously. Figure 3 below depicts the heterogeneity in the short- and long-term price elasticities across customers for each decision component.

Figure 3 Heterogeneity in Price Elasticities Across Customers



The price elasticities exhibit high degree of heterogeneity. This result is consistent with the importance of heterogeneity for model fit and prediction reported in Table 3. Accounting for the unobserved heterogeneity is an important component of the B2B individually targeted pricing strategy.

In the next section, we discuss how to utilize the parameters estimates and the insights on customers' behaviors from our proposed non-homogeneous HMM, to dynamically and individually price target each customer for long-term profitability.

5. Price Policy Simulation Procedure

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:

$$LCV_i = \sum_{j=1}^{J_i} \frac{(\text{Price}_{ij} - \text{Cost}_{ij})q_{ij} \left[(1 - \delta_{ij}^b) + \delta_{ij}^b \delta_{ij}^w \right]}{(1+r)^{\tau_{ij}}} \quad (22)$$

where, δ_{ij}^b is a binary indicator that equals 1 when a quote is requested on event j for customer i , and is zero for a direct order, δ_{ij}^w is another binary indicator that equals 1 when the bid 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 9 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 2008.⁸ We then use the parameter estimates from the calibration dataset and conduct the price optimization over the holdout period of 9 months. The optimization is done for a sample of 300 customers who experienced between 6 and 16 purchase events over the calibration period.⁹ An average customer has an average of ten purchase events in the 9-months 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 period¹⁰. 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

⁸ 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.

⁹ We use this criterion to choose those customers whose purchase frequencies are representative of the company's customer base.

¹⁰ We choose to stay within each customer's historical range of experienced prices, to avoid a drastic change in the price regime. This specification would also account for the price variance experienced by each customer.

purchase events of the customer in the 9-months optimization 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 probabilities. At each simulated purchase event, profits are computed by weighing the HMM latent state-specific profits by the state membership probabilities. For each customer we simulated 200 random sequences of purchase events associated with each path such that the cumulative inter-purchase event time does not exceed 9 months. We then calculate the LCV over the 9-months horizon by computing for each price path the average net present profits, assuming a discount rate of 12% (a weekly discount rate r of 0.22%) over the Monte Carlo draws.¹¹ This yields an optimal dynamic price policy for each customer. Full details of the price simulation process are available in Appendix B.

6. Price Policy Simulation Results

Following the simulation procedure outlined in Section 5, we generate for each customer nine months of purchase events¹², and find the price path that maximizes the LCV. We assume a weekly discount rate r of 0.22%, which corresponds to an annual discount rate of 12%. The cost for the seller is mildly associated with the current level of the aluminum market price (LME), and we set the “cost per pound” to be $1.34+0.21*LME$, which is similar to the average unit cost of \$1.60 per pound observed in the dataset. We compare the performance of our proposed individual dynamic pricing policy to the performances of six competing policies: (1) an individual static pricing policy, (2) a segment dynamic policy, (3) a segment static policy, (4) an aggregate static policy, (5) a myopic individual dynamic policy, and finally, (6) the seller’s current pricing policy. These competing policies will help us to assess the impact of properly accounting

¹¹ 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 promise between precision and computational time, as it offers 8% improvement in profits over 100 draws, but only underperforms 300 draws by 1%.

¹² We chose a time horizon of nine months because that is the duration of the current pricing policy in the validation period (Jan 2008 to September 2008)

for individual targeting, dynamics, and a long-term perspective. Under the individual policies, different unit prices are charged to each customer, fully utilizing each customer's individual price parameters. The segment policies provide a single optimal price for those customers that are in the vigilant state, and a single optimal price for those that are in the relaxed state. The aggregate policy charges only one optimal price for all customers. With regard to dynamics, the dynamic policies update the prices charged to each customer at each purchase occasion, whereas the static policies will charge the same price to each customer (or each segment) throughout the simulation horizon. Specifically, the static policies identifies each customer's buying state at the beginning of the simulation and treat these initial state memberships as fixed indefinitely, hence neglecting the possibility that customers' buying states can change at each purchase occasion throughout the planning horizon. In other words, static policies might initially identify a particular customer as a valuable customer, and will repetitively price the same way on each purchase occasion, failing to acknowledge that this customer might have reacted to the prices and shifted his behavior to become less valuable.

With respect to having a long-term planning horizon, policies (1) to (4) all use LCV maximization as the criterion when deciding which prices to charge, whereas the individual myopic dynamic policy considers only the short-term effect of pricing in each period. The myopic policy optimizes the prices to maximize only the current period profit, given the customer's updated state-membership probabilities.

Table 9 summarizes the performances of the six price policies. The proposed policy yields the highest LCV per customer at \$4,795 over nine months. By incorporating both individual and dynamic perspectives, there is a 73% gain in profitability compared to the aggregate static price policy. There is a 40% improvement due to individual targeted pricing (\$3,855 vs. \$2,759), and an 24% improvement due to accounting for dynamics (\$4,795 vs. \$3,855). The relative sizes are consistent with the existing literature in marketing, which highlights the importance of accounting for heterogeneity. Similarly, by having a long term perspective, the proposed policy improves the myopic policy by 17%.

Examining the seller's current pricing policy during the nine-month holdout period suggests that the company is not doing too poorly – it beats the segment static policy, but falls slightly short of the segment dynamic policy. It appears that the seller might acknowledge that there are

“good customers” and “bad customers”, and treats them differently, but does not update very well its belief of how customers transition between these two states. Hence, the current policy does not properly account for the individual and the dynamic aspects of pricing. Compared to the current policy, the proposed policy offers a 52% improvement in profitability. On a base of 1,800 customers, an average increase of \$1,700 in profit over nine months can translate into a potential profit improvement of approximately \$4 million annually.

To visualize the performance of each policy, we depict in Figure 4 the effect of implementing the proposed individual dynamic policy, the individual static policy, and the individual myopic policy, over the nine-month planning horizon. The plots show the average monthly profit per customer under each pricing policy. Alongside our policy simulations, we also track the profit under company’s current policy during this period. Several interesting patterns emerge. First, all individual-level policies (static, dynamic and myopic) start at a higher per-customer profit level than the current policy, highlighting the immediate reward of individual price targeting. Second, the proposed policy immediately jumps and achieves a higher projectile relative to the static policy, showing the impact of accounting for customer dynamics. Third, the proposed policy initially underperforms the myopic policy, as it carefully balances the interplay between several forces that govern customer buying behaviors, namely, 1) charging lower prices to increase bid-acceptance; 2) charging lower prices to keep customers in and toward the relaxed state, 3) charging higher prices to increase profit, and 4) charging higher prices to not decrease customers reference prices. The myopic policy, on the other hand, ignores points 2) and 4), and as shown in Table 9, charges lower prices than the proposed policy in order to convert the quote-requests into immediate sales. However, doing so will drastically lower customers’ reference prices, make future prices more likely to be perceived by customers as losses, and subsequently increase the likelihood of customer migration to the vigilant state, which, as we have shown in Figure 1 and 2, could have long lasting negative effects. The lowered reference prices create more downward pressure for the seller when setting prices, establishing a vicious cycle of lowered prices and hence lowered profits. Within the first two to three months, the proposed policy begins to dramatically outperform the myopic policy, demonstrating the significance of keeping the long-term reference-price effects in mind when setting prices. This findings consistent with those of Montoya et al. (2010), who show that in the context of physician detailing and targeting, myopic policies initially perform better than the forward-looking policy,

as the forward-looking policy needs to first invest in moving the physician base to higher buying states.

Furthermore, Figure 4 shows the decline in myopic policy profit crosses over the static policy around the fifth month, suggesting that in the long run, accounting for the long-term effects of pricing, even if the prices are set statically for each customer, is more impactful than dynamically yet myopically varying prices. In contrast with the findings of Montoya et al. (2010), who highlight the superior profit performance of a dynamic policy over a static policy even if such policy is myopic in nature, we show that in our pricing context, where reference-price updating plays a crucial and a long-lasting role, one needs to first and foremost be aware of these reference-price effects and the long-term impact of pricing. We argue that a dynamic pricing policy is useful only if the seller keeps a long-term perspective. Myopia, in the world of B2B where pricing can affect both state transitions as well as reference-price updating, is very dangerous and can be a slippery slope, literally.

Figure 4: Policy performance Comparison Over 9 Months

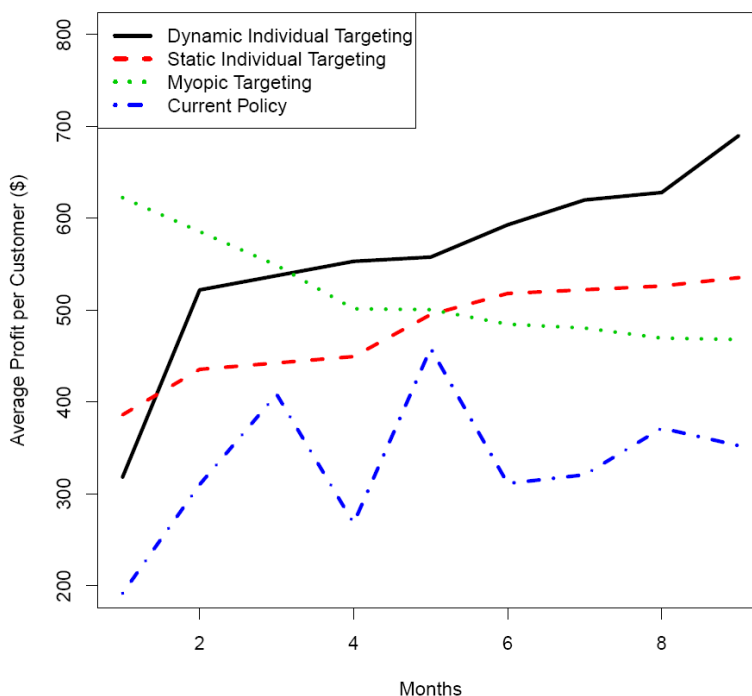



Table 9: Policy Performance Comparison Over 9 Months

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	NA	NA	\$3.10

6.1. Comparison of Prices Charged Under Different Price Policies

We compare the unit price per pound charged under each policy. The results are shown in Table 9 and can be visualized in Figure 5. Compared to the company's current pricing practice, both the myopic and the proposed policies recommend charging higher prices, with multi-modal distributions, to flexibly account for heterogeneity in price elasticities among customers. The myopic policy on average recommends lower prices than the proposed policy (mean prices of \$3.39 vs. 3.27, and median prices of \$3.29 vs. \$3.10). These lower prices are the result of focusing only on gaining a customer's business today, with no regard for how the lower prices will adversely impact this customer's references prices going forward.

Figure 5: Comparison of Price Distributions



6.2. How do We Value Vigilant and Relaxed Customers Differently?

We have shown in Section 4 how customers behave differently with respect to the decision components in the B2B buying process, when they are in vigilant vs. relaxed buying states. We now synthesize the myriad of differences in behavioral parameter estimates, translate them into the contexts of LCV, and to examine how much vigilant and relaxed customers differ in terms of their long-run profitability to the firm. To perform this comparison, we track, for the entire nine-month holdout period under the proposed dynamic targeted pricing policy, each customer's buying behavior and state-membership probability on each purchase occasion. Then, given the state that a customer is in, we assign his profit to that particular state. The distributions of the state-specific LCVs are shown in Figure 6. Each customer has an overall nine-month LCV of \$4,795. Customers who are in the vigilant state have a LCV of \$3,551 on average, whereas those who are in the relaxed state have an impressive LCV of \$6,137, almost doubling the profit of their vigilant state counterparts. As Figure 6 illustrates, the stark contrast in profitability among the two states results in a bimodal pattern of the overall LCV distribution. This analysis shows that not only will customers in the relaxed state exhibit buying behaviors that are “favorable” to

the seller, as portrayed in Section 4, these “favorable” behaviors can translate into substantial financial gains.

To see how state membership can directly impact profitability, we graph the monthly LCV with the percentage of customers in relaxed state on the same time-series plot in Figure 7. We see that as the relaxed state membership increases during the nine-months planning horizon, the profitability directly increases – first in almost lock-step during the initial two months, and then experiences a lagged increase from the third month onward.

The results monetize the phenomenon of the slow long-run behavioral adaption shown in Figure 1 and 2. When customers are transitioned into the relaxed state, the favorable buying behaviors become long lasting as the state-memberships are sticky, resulting in consistently higher profit, which over the span of nine months could be as much as double the profit in the vigilant state.

Figure 6: Long-term Customer Value (9 months)





Figure 7: Impact of Relaxed State Membership on Profitability

6.3 How Much to Charge, and to Whom?

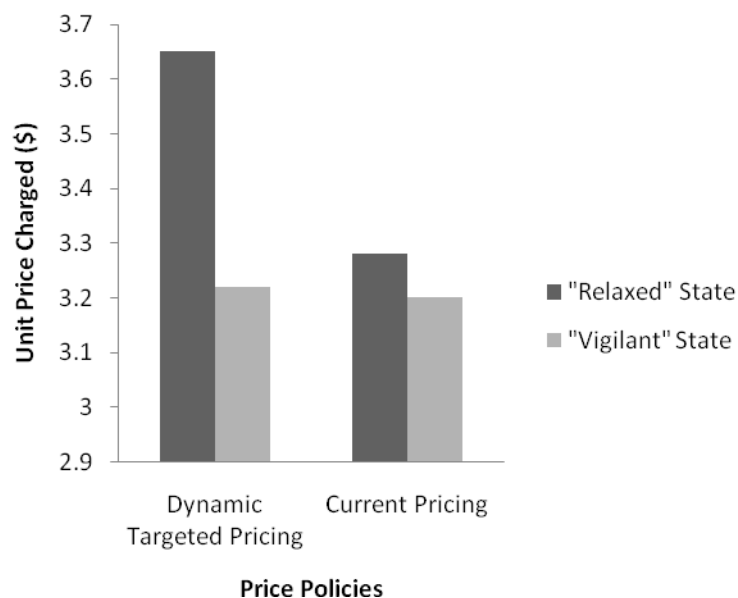
Figure 5 has shown that compared to the current pricing practice, price distribution under the proposed policy has a higher mean, is more diffuse, and has multiple local modes to accommodate customer heterogeneity. An examination of how much customers in each state are priced, shown in Figure 8, tells us that whereas the seller's current policy does not vary prices between vigilant and relaxed customers, the proposed price policy recommends charging a much higher price for customers that are in the relaxed state vs. those that are in the vigilant state (\$3.65 vs. \$3.22). The reasons for the current pricing's lack of price differentiation could be a mixture of 1) the seller not properly accounting for customer heterogeneity, 2) the seller not properly assigning the correct state membership to each customer, and 3) the seller not properly tracking customer's state transition dynamically and hence adjust its prices in a timely fashion.

These empirical results suggest the following pricing strategy: when customers are in the

relaxed state, the seller should charge a higher price, as doing so will 1) increase immediate profit and 2) increase the customer's reference price. Although a higher price might increase the probability of transition into the vigilant state, the stickiness of the state probability and the already increased reference prices as a "shield" against price losses empirically suggests that the benefits of a higher price outweigh the potential negative impact. On the other hand, when customers fall into the vigilant state, the seller should price very aggressively to achieve the following two goals. First, for the customers on the margin, the seller needs to try to get them out of the vigilant state, as this state has a drastically lower profit, and as shown in the adaptation plot in Figure 2, customers tend to stay much longer in the vigilant state if there is no intervention. Due to customers' loss aversion in the long-run, the positive gains in state transition probability towards the relaxed state outweigh the negative effect of lowered reference prices. Second, for those customers who are persistently in the vigilant state, who habitually order via quote requests and with little chance of transitioning back to the relaxed state, lowering prices (above the seller's cost, of course) would at least increase the immediate probability of the quote requests getting successfully accepted, thus ensuring some profits from these "unprofitable customers". In short, the overall pricing strategy for our seller would be – 1) milk your best customers with high prices; 2) turn your marginally vigilant customers into good customers by first lowering and subsequently increasing prices; and 3) charge low prices for your "worst" customers to salvage some profits from them.

In summary, the improved profit performance of our individual dynamic targeted price policy, relative to the current policy can be attributed to 1) understanding each customer's inherent price elasticities and setting prices to fully account for heterogeneity, 2) accurately identifying and tracking customers' buying states and dynamically changing the prices accordingly, charging higher prices for those customers in the relaxed state, and 3) a long-run perspective to ensure the optimal tradeoff between the short and long run impacts of pricing on immediate profit, delayed profit, reference-prices effects, and customer buying state migrations.

Figure 8: Optimal and Current Unit Price per Pound by State



6.4 Pricing in a Volatile Economic Environment - the Role of External Reference Price

In this section we examine the price strategies that the seller should use to manage volatile economic conditions. Specifically, we look at how fluctuations in the raw aluminum prices would impact the seller and its customers.

For the duration of our dataset, the aluminum prices on LME have fluctuated between US\$2,393 to US\$3,318 per tonne. Initially stable for the first half of 2007, the price of aluminum experienced a sharp drop starting in August of 2007 when the world felt the seriousness of the sub-prime meltdown, followed by a rapid increase starting in January of 2008, reaching the peak in July, and then eventually declined for the periods leading up to the Great Recession. Although the price of aluminum is perhaps the most important factor for our seller as it directly influences the wholesale price, the impact of LME fluctuation on seller's profitability is not readily obvious. To illustrate, rising raw material prices could increase the seller's costs, providing downward pressure on profitability; but at the same time, customers might view firm's pricing to be

attractive in reference to the high material price, and the seller might be able to pass the increased costs onto customers. On the other hand, a drop in material prices might lower the seller's cost, but at the same time customers might wish to pay lower prices in light of the cheaper commodity.

To provide an empirical answer to the question “how should the seller price in the presence of a volatile economic environment?”, we use the proposed individual dynamic price policy to perform price simulations in the presence of a 20% increase and a 20% decrease in LME, and compare the resulting prices and profits to those of the original LME level.¹³ In general, a 20% change in LME translates to approximately 10% change in seller's cost.

Figure 9 shows the distribution of the optimal prices under each LME regime. For a 20% increase in the LME, 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 10 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 firm to pass through this increase in replacement cost by increasing the unit prices for customers in the vigilant state by 12%. In comparison, a modest increase of 4.6% is optimal for customers in the relaxed state. This makes sense as vigilant customers are more sensitive to changes in the LME prices, relative to relaxed customers (see Table 4). Figure 10 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 corresponding lowering of the internal reference price, and this can have long terms consequences for the seller's profitability.

Figure 9: Price Distributions at Different LME Levels

¹³ Our use of a 20% shock in LME falls within the range of fluctuations experienced in this business. The maximum daily, monthly and 3-monthly LME fluctuation in our data are 5.6%, 22%, and 31% respectively.





Figure 10 – Optimal and Current Unit Price per Pound by State, under Different LME Regimes



So what governs this asymmetry in pricing? There are two opposing mechanisms at work – buyers’ internal vs. external reference prices. First, when LME increases, as shown in Table 4, customers respond by treating it as a higher external reference price, and reacts more enthusiastically to firm’s current prices in the forms of larger quantity, higher direct order and bid acceptance probability. This behavior is more prevalent amongst vigilant customers, who pay more attention to the external reference prices. The usage of LME as an external reference price grants the seller the license to increase prices to offset the rising cost, but only by so much, as too high of a price increase would be punished by customer’s focus on internal reference price, which would represent the price increase as a “loss”. However, when the LME drops, although the seller is now “penalized” by exposing customers to a lower external reference price, it has no strong incentive to lower prices because doing so will lower customer’s internal reference prices, which appears to have a much more drastic impact than the external reference price effect.

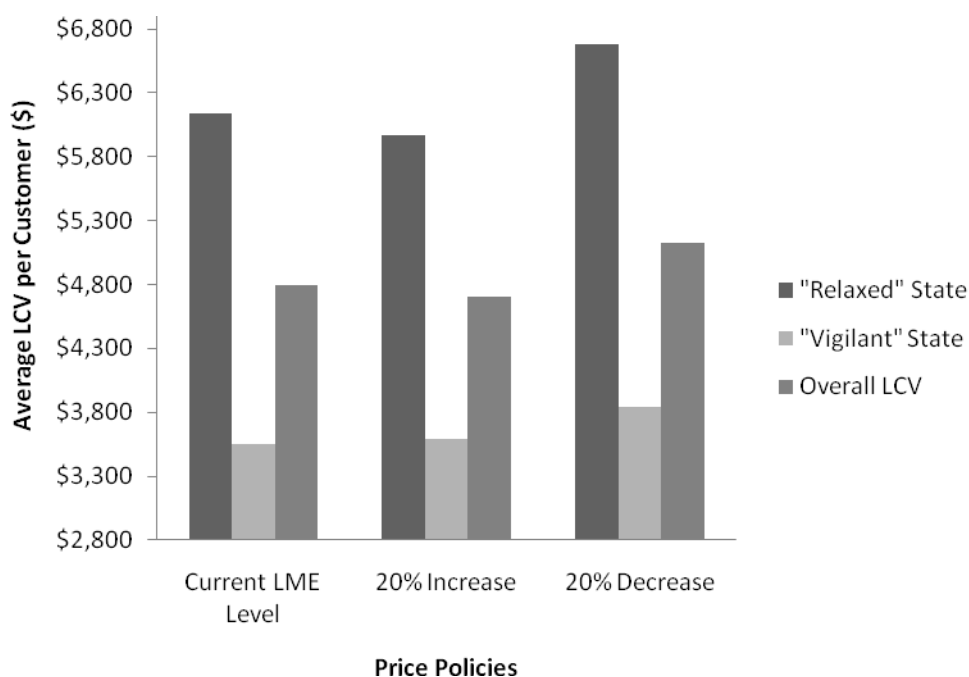
This price strategy of passing on the cost increase, and “hoarding” the benefit of cost decrease is consistent with the economic research on price fairness (Akerloff 1979, Kahneman et al. 1986, Okun 1981, Urbany et al. 1989). In their seminal work, Kahneman et al (1986) formulate a principle of *dual entitlement*, which governs community standards of fairness in the

market place. 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 would be glad to pay their past reference prices. In our framework, price fairness is captured by the external, internal reference prices, and loss aversion – when costs increase, the seller is justified to increase prices to a certain degree, when costs decrease, the potential of the lowered reference prices and the subsequent “price losses”, amplified by loss aversion, suggest that it is best to keep prices the same.

How does this pricing strategy translate into long-term profitability? Figure 11 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 \$4795 to \$4701). However, the optimal price response to a drop in LME prices increases the LCV substantially, from \$4795 to \$5103. This is an increase of about 7% increase resulting from proportional increases within both states.

To the best of our knowledge, this is the first paper to empirically test the dual-entitlement principle, and to measure the extent of it in a B2B transaction setting. Therefore, to profitably manage the fluctuation in economic conditions, sellers should pass on the cost increases, especially for vigilant customers. To justify price increases, sellers should effectively communicate the reason for the price increases to customers, and make them aware of the high external reference on LME. On the other hand, cost decreases present a good opportunity for managers to enjoy a period of heightened profitability by keeping prices the same, at least in the short-run.

Figure 11 – LCV per State, under Different LME Regimes



7. General Discussion

Understanding and managing the impact of pricing on customer buying behavior is critical for B2B firm profitability in the long-run. Dynamic and targeted pricing at the individual level provide firms with great flexibility in extracting additional profits from their customers. To do so successfully, and to ensure long-term profitability from each customer, the firm needs to 1) effectively leverage individual-level information and consider heterogeneity in customers' responses, 2) understand and appropriately capture the industrial buyer buying process and the interdependencies therein, 3) understand customer's sensitivities to prices and the asymmetric reference price effects, 4) measure the dynamics in customer's buying behavior, 5) understand both the short-term and long-term impact of the pricing policy, and 6) operate in a business environment where individually and inter-temporally targeted pricing are feasible.

In this paper, we present such an integrative framework to capture the individual-level and long-term impact of pricing decisions in B2B settings. We identify several insights by applying our modeling framework in the under-researched area of B2B pricing. First, we

uncover two latent and theoretically interesting buying behavior states which characterize the buyer-sellers relationships. Second, we find that pricing in B2B 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 impact of pricing, may severely underestimate the impact of pricing. Third, despite the view that B2B customers are rational, we find significant asymmetric reference effects in industrial customer responses to prices. Fourth, the impact of a price hike is stronger than the impact of a price drop, and takes much longer for customers to adapt to, particularly for quote request behaviors, bid acceptance, and vigilante state membership - factors that are more closely related to customers' relationship with the sellers. Therefore, firms should be cautious of the possibly long-term negative impact of charging a customer a higher price on any particular order.

Through a series of price policy simulations and customer life-time value analyses, we demonstrate that the proposed dynamic targeted price policy can offer a 56% improvement in profitability over the current policy. We decompose the effects of individual pricing, dynamic pricing and pricing with a long-term perspective, and show that in the presence of reference-price effects, a dynamic pricing policy itself is not sufficient in improving profit, if it does not properly account for the long-term effect of pricing on future reference-price updating. Furthermore, we show that relaxed customers on average are worth twice as much as vigilant customers, and that the optimal price policy is a "Milk, Transition, and Salvage" strategy.

When implementing our proposed dynamic targeted pricing strategy in a volatile economic environment with cost fluctuations, we generally recommend the seller to pass on the increased costs to customers but keep the prices the same when costs decrease. This active management will help the seller maintain existing level of profitability during inflationary periods, and enjoy increased profitability during deflationary periods. These empirical results are consistent with the dual-entitlement principle.

Methodologically, we jointly model the multiple interrelated decisions via copulas, capture the dynamics in the interrelated buying decisions and the long-term effect of pricing decisions via a non-homogeneous HMM, and fully account for potential price endogeneity in a Bayesian control function framework. To the best of our knowledge this is the first paper to incorporate copulas in a HMM, and the first paper to flexibly incorporate price endogeneity in this modeling setting. Substantively, this research offers a first step in providing B2B managers

with a comprehensive modeling framework to manage their customer base through an effective and profitable dynamic targeted pricing policy. As many B2B firms are still applying cost-based pricing strategies (Anderson et al. 1993), we demonstrate that there is substantial value in instituting a value-based pricing policy, taking advantage of the flexibility offered by the ability 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 firm's pricing decisions. One could extend our framework to incorporate customer's expectation about the price changes (Lewis 2005). Second, similar to typical B2C and B2B CRM datasets, our data include no information about competitive pricing. Obtaining such information one could extend our framework to a competitive setting. Alternatively, one could augment the competitive pricing effect using an approach similar to that by Moon et al. (2007). Third, while we focus in this research on B2B pricing decisions, our methodological framework is appropriate for many B2C settings in which the customer buying process is composed of several interrelated decisions, and the firm has the opportunity to price discriminate and varying degrees. Finally, although this is the first paper to demonstrate empirically the dynamic persistence of reference-price effects, future research could examine the dynamic patterns of reference-price effects in other industries, and to consider alternative models of reference prices when comparing prices among incomparable products.

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.
- 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. 2009. Modeling Multivariate Distributions Using Copulas: Applications in Marketing. *Forthcoming in Marketing Science*.
- 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., and Loewenstein, G. 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 and 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.
- 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. and O'Shaughnessy, J., 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., R. McCulloch, G. Allenby. 1996. The Value of Purchase History Data in Target Marketing. *Marketing Science*, 15(4), 321-340.
- Schweidel, D., E. Bradlow, P. Fader. 2009. Portfolio Dynamics for Customers of a Multi-Service Provider. *Working Paper*. University of Pennsylvania.
- 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.
- Stock J., M. Watson. 2003. *Introduction to Econometrics*. Reading, Mass: Addison-Wesley.
- 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. and 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
- Villas-Boas, M. J. and 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 A. Priors and Full Conditionals Distributions for Hierarchical Bayesian Estimation

We denote by $\theta_i = \{\beta_{is}, \gamma_{is}\}_{s=1}^S$ the set of random-effect parameters and by $\phi = (\theta_{qt}, \theta_{bw}, \sigma, \{\alpha_s^0\}_{s=1}^S, \{\beta_s^0\}_{s=1}^S)$ the set of fixed-effect (including state-specific and non state-specific) parameters. See equation (2) - (20).

Priors

Random-effect parameters θ_i : $\theta_i \sim N(\mu_\theta, \Sigma_\theta) \Rightarrow P(\theta_i) \propto \exp\left(\frac{1}{2}(\theta_i - \mu_\theta)' \Sigma_\theta^{-1} (\theta_i - \mu_\theta)\right)$

$$\mu_\theta \sim N(\mu_0, V_0) \Rightarrow P(\mu_\theta) \propto \exp\left(\frac{1}{2}(\mu_\theta - \mu_0)' V_0^{-1} (\mu_\theta - \mu_0)\right)$$

$$\Sigma_\theta \sim IW(df_0, df_0 S_0)$$

Fixed-effect parameters ϕ : $\phi \sim N(\mu_\phi, \Sigma_\phi) \Rightarrow P(\phi) \propto \exp\left(\frac{1}{2}(\phi - \mu_\phi)' \Sigma_\phi^{-1} (\phi - \mu_\phi)\right)$

Likelihood: $L(\text{data}, \{\theta_i\}, \phi, \mu_\theta, \Sigma_\theta) = P(\text{data} | \{\theta_i\}, \phi) P(\{\theta_i\} | \mu_\theta, \Sigma_\theta) P(\phi) P(\mu_\theta) P(\Sigma_\theta)$

Full conditionals: $P(\theta_i | \mu_\theta, \Sigma_\theta, \phi, \text{data}_i) \propto \exp\left(\frac{1}{2}(\theta_i - \mu_\theta)' \Sigma_\theta^{-1} (\theta_i - \mu_\theta)\right) P(\text{data}_i | \theta_i, \phi)$

$$P(\phi | \mu_\phi, \Sigma_\phi, \{\theta_i\}, \text{data}) \propto \exp\left(\frac{1}{2}(\phi - \mu_\phi)' \Sigma_\phi^{-1} (\phi - \mu_\phi)\right) P(\text{data} | \{\theta_i\}, \phi)$$

$$\mu_\theta \sim N(\mu_n, V_n)$$

$$\Sigma_\theta \sim IW(df_1, df_1 S_1)$$

where,

$$V_0^{-1} = [V_0^{-1} + N \Sigma_\theta^{-1}]$$

$$\mu_n \sim V_n [\mu_0 V_0^{-1} + N \bar{\theta} \Sigma_\theta^{-1}]$$

$$df_1 = df_0 + N$$

$$df_1 S_1 = \sum_{i=1}^N (\theta_i - \mu_\theta)(\theta_i - \mu_\theta)' + df_0 S_0$$

N is the number of individuals, $P(\text{data} | \{\theta_i\}, \phi)$ and $P(\text{data}_i | \theta_i, \phi)$ are given by Equation (31), and quantity, inter-purchase-event time, quote request, bid acceptance and price-generating equation respectively follow lognormal, log-logistic, logistic, conditional logistic and normal distributions.

The MCMC procedure generates a sequence of draws from the posterior distribution of the model's parameters. Because the full conditionals for θ_i and ϕ do not have closed forms, the Metropolis-Hastings algorithm is used to draw the samples. In particular, we use a Gaussian random-walk Metropolis-Hastings where a proposal vector of parameters $\varphi^{(t)}$ for θ_i or ϕ , at iteration t is drawn from the Gaussian proposal distribution $N(\varphi^{(t-1)}, \sigma_\varphi^2 \Delta)$ and accepted using the Metropolis-Hastings acceptance ratio. The tuning constant σ_φ^2 and Δ are updated using the method proposed by Atchade (2006) to yield an acceptance rate of approximately 25%.

We use the following uninformative prior hyper-parameters for the two-state HMM:

For random effect parameters:

$\mu_0 = [\tilde{0}]$, $V_0 = 100I_{n\theta \times n\theta}$, $df_0 = n\theta + 5$, $S_0 = I_{n\theta \times n\theta}$, where $n\theta$ is the number of parameters in θ for each individual, and $I_{n\theta \times n\theta}$ is an identity matrix of rank $n\theta$.

For fixed-effect parameters:

$\mu_\phi = [0.3, 0.3, 0.5, 0.5, 0.5, \eta]$, where η is a 1x13 vector of zeros.

Σ_ϕ is a diagonal matrix of rank $n\phi$, with the first 5 diagonal elements set to 2, and the next 13 diagonal elements set to 100.

Because some parameters are transformed to an exponential scale, we accordingly chose prior hyper-parameters for μ_θ , Σ_θ , and μ_ϕ such that their prior means and prior variances are diffuse in the *transformed* space.

Appendix B – 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) . We use the same random draws for all price points at the purchase event.
2. Using the Frank Copula (Nelsen 2006) we 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} denotes the Frank copula correlation parameter that governs the interdependence structure between quantity and inter-purchase event time.

3. Generate state specific quantity q_c such that $q_{sij} = F_q^{-1}(u_1)$ with mean $\mathbf{x}_{ij}\boldsymbol{\beta}_{qsi}$ and parameter σ .¹⁴
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 $\mathbf{x}_{ij}\boldsymbol{\beta}_{qsi}$ and parameter α_s .
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 $b_{sij}^* = \mathbf{x}_{ij}\boldsymbol{\beta}_{bsi} + e_{bij}$, where

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

8. Calculate the latent utility for quote request, where $w_{sij}^* = \mathbf{x}_{ij}\boldsymbol{\beta}_{wsi} + e_{wij}$, 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. Given the 4 state specific behaviors, calculate state-conditional profits. The purchase event unconditional profit is calculated as the weighted-average of the state-conditional profits, weighted by state-membership probability for purchase event j .
10. Update the state membership for each customer as described by Equation (18).
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 reaches the planning horizon (9 months). The 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 event for customer i in the planning horizon.
13. Repeat step 1-12, 200 times to generate 200 profit functions for each possible price paths 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's management: (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

¹⁴ As quantity is used to determine inter-purchase event time and reference price going forward, to avoid bad draws of quantity that will make a purchase path explode into unreasonable regions and create computational issues, we resample q_{sij} for a customer if the current draw is less than half of the minimum quantity, or more than double of the maximum quantity experienced by the customer during the calibration period.

administrative costs (4) we use the actual LME commodity prices in each day in the planning horizon.