JOSEPH PANCRAS and K. SUDHIR*

Advances in data collection and storage technologies have given rise to the customer data intermediary (CDI), a firm that collects customer data to offer customer-specific marketing services to marketers. With widespread adoption of customer relationship management (CRM) and one-to-one (1:1) marketing, the demand for such services continues to grow. Extant empirical research using customer data for CRM and 1:1 marketing tends to have an engineering emphasis and focuses on developing analysis techniques to implement CRM and 1:1 marketing optimally (i.e., the technology for the CDI). In contrast, this article focuses on marketing strategy issues that the intermediary faces, given the availability of the technology to implement such services. Specifically, the authors develop an empirical framework to evaluate the optimal customer (exclusive/nonexclusive), product (quality or accuracy of the 1:1 customization), and pricing strategy for a CDI. They illustrate the framework for one type of CDI—a 1:1 coupon service firm that caters to grocery manufacturers—using household purchase history data from the ketchup market. The authors find that selling on a nonexclusive basis using the maximum available purchase history data is the most profitable strategy for the CDI in the particular market. They also evaluate the potential impact of retailers entering the 1:1 coupon service business. Because 1:1 marketing can increase the retailer's profits from goods sold, it is optimal for the retailer to undercut the prices of a pure-play CDI that offers 1:1 coupon services.

Optimal Marketing Strategies for a Customer Data Intermediary

In recent years, a new type of firm has emerged—namely, the customer data intermediary (CDI)—that specializes in collecting customer behavior and demographic data and offers customer-specific marketing services. For example, in the grocery and drugstore markets, Catalina Marketing obtains purchase history data through cooperating retailers and provides targeted coupons on behalf of grocery manufacturers to households purchasing at that particular retailer. Currently, Catalina Marketing has penetrated approximately 21,000 of the nearly 34,000 supermarkets in the United States and records approximately 250 million transactions per week. The firm's targeted coupons are considerably more effective with redemption rates of 6%–9% relative to 1%–2% for free-standing-insert coupons. On the Internet, companies such as DoubleClick and TACODA Systems collect previous visit data from cooperating Web sites and use these to deliver targeted advertising for their advertising clients. In the catalog and specialty retailing industry, firms such as Abacus B2C Alliance and I-Behavior pool transactional data from more than a thousand catalog titles/retailers on approximately 90 million households to offer improved targeted direct marketing services to their members. A key difference between the two firms is that whereas Abacus uses purchase data only at the catalog level, I-Behavior uses purchase data at the stockkeeping unit level (Miller 2003).

With widespread adoption of customer relationship management and one-to-one (1:1) marketing, the demand for

*Joseph Pancras is Assistant Professor of Marketing, School of Business, University of Connecticut (e-mail: joseph.pancras@business.uconn.edu). K. Sudhir is Professor of Marketing, Yale School of Management (e-mail: k.sudhir@yale.edu). This article is based on an essay from the first author's 2005 dissertation at New York University. The authors thank Joel Steckel, Yuxin Chen, Paris Cleanthous, Russell Winer, Vicki Morwitz, and Jiwoong Shin for feedback and helpful discussions and the workshop participants at Carnegie Mellon University, Koc University, London Business School, New York University, Purdue University, State University of New York at Buffalo, University of Central Florida, University of California, Berkeley, University of California, Los Angeles, University of Connecticut, University of Georgia, University of Southern California, University of Texas at Austin, Washington University at St. Louis, and Yale University for their comments.

To read and contribute to reader and author dialogue on JMR*, visit http://www.marketingpower.com/jmrblog.*

ISSN: 0022-2437 (print), 1547-7193 (electronic)

CDI services continues to grow. Table 1 lists some of the major players in the CDI services business. For each of these players, we provide a brief description of their company and report their revenues, market capitalization, and growth rates. As Table 1 shows, the industry is gaining in importance, as reflected in its market valuations, revenues, and growth rates. Several companies in this industry have revenues in the hundreds of millions of dollars and valuations of more than a billion dollars.

Advances in data collection, data storage, and customer-specific promotion delivery technologies have fueled the rise and growth of CDIs. The use of scanners in offline

Table 1
ILLUSTRATIVE SET OF CDIs (LATEST AVAILABLE FIGURES)

Company and Division	*Revenue (in Millions of Dollars)*	*Total Company Revenue (%)*	*Market Cap (in Millions of Dollars)*	*Client Profile (Client Examples)*	*Revenue Growth Rate*	*Description*
Catalina Marketing[a]	431	100	1320	Manufacturers of packaged goods and grocery retailer (e.g., Nestlé and Safeway, respectively)	8% per year over 2000–2004, non–core divestments reduced revenue in 2005	Proprietary technology at point of purchase in grocery and pharmaceutical retailers to track customer transactions and deliver customized coupons. Tracks more than 250 million transactions per week across more than 21,000 supermarkets worldwide. Tracks purchase history of more than 100 million households in the United States. Delivers more than 4.5 billion customized promotional messages.
DoubleClick (Abacus B2C Alliance division)[b]	105	35	984	Catalog merchants (e.g., L.L.Bean, Sharper Image)	10% (2004), 15.1% (2003)	Consolidates the input from 1550 catalog, online, and retail merchants into a single database. Data on more than 4.4 billion transactions from catalog call centers, Web sites, and retail stores made by more than 90 million households, with household purchase data dating back five years. DoubleClick was acquired by Hellman & Friedman LLC (a private company) in the third quarter of 2005.
Experian (Z-24 division)[a]	520	23	N.A.	Catalog companies (e.g., Boston Proper, J. Jill, retailers, nonprofits)	4% in 2005, 6% per year over 2001–2004	The Z-24 database is similar to Abacus. Data from more than 755 catalogs with 38 million households that have purchased over the past two months. Experian is also a player in business-to-business targeting, with BizInsight's database of more than 15 million U.S.-based companies.
VT & NH Direct Marketing Group (I-Behavior)	N.A.	100	N.A.	Catalog companies (e.g., Gardener's Supply, Vermont Country Store)	N.A.	Competitor to Abacus and Z-24, but it uses transactional data at the stockkeeping unit level (in contrast to Abacus and Z-24, which use catalog-level data). There are 1000-plus contributors, mostly medium-sized catalog companies, and data on 103 million consumers and 89 million households.
Harte-Hanks (direct marketing division)[b]	695	61	2060	Retailers, finance sector, pharmaceuticals, telecommunications, and high-tech firms	8.4% (2005), 9% (2004), 2% (2003)	Specializes in providing targeting solutions in automotive, consumer products, financial services, insurance, high-tech, pharmaceuticals, retail, and telecommunications. Provides a suite of services, including constructing the database (Trillium Software System), accessing the data (Allink, inTouch), in-house data analytics, application, and execution of campaigns.
Q Interactive (owns CoolSavings Inc.)[b]	17.3 in the third quarter	N.A.	N.A.	Retailers, packaged goods manufacturers (e.g. Unilever, Land O' Lakes, Best Buy)	20% per year over 2001–2004, 92% increase in the third quarter of 2005 compared with the third quarter of 2004	CoolSavings maintains a network of Web sites that feature online coupons, targeted marketing (from 20 million online consumers), lead generation, e-mail marketing, and loyalty programs for more than 1000 companies in the retail, packaged goods, and media industries. CoolSavings was taken over in January 2006 by the privately owned Q Interactive, an online targeted marketing service provider for advertisers and publishers that uses demographic, behavioral, and transactional data.

[a]2005 figures.
[b]2004 figures.
Notes: N.A. = not available.

retailing and the intrinsic digital nature of online retailing have enabled the easy collection of purchase and other transaction histories, and the falling costs of digital storage and computation have made the recording and analysis of vast amounts of purchase history data for 1:1 marketing feasible. Furthermore, advances in promotion delivery technologies to individuals (in the store at the point of purchase, at home through direct mail, online through e-mail, and even on the move through cell phones) have increased the effectiveness and timeliness of 1:1 marketing strategies.

Despite the growing economic importance of CDIs, there is little empirical research addressing strategic issues of concern to this industry. To date, extant research on this industry has tended to be of an "engineering" nature, focusing on how firms should use individual browsing/purchasing data to personalize advertising or promotions. This research has occurred in marketing, information systems, and computer science (e.g., Adomavicius et al. 2005; Ansari and Mela 2003; Liu and Shih 2005) and is typically positioned as a means by which a firm can take advantage of its customer databases to improve marketing effectiveness. From the CDI perspective, this research develops the technology to create customer-specific marketing services.

In contrast to such engineering-based research, this article focuses on marketing problems that CDIs face. Specifically, we ask the following research questions: (1) Conditional on the availability of the 1:1 technology, what is the optimal customer and product strategy for the CDI? and (2) Conditional on the customer and product strategy, what price should the CDI charge for the service?

In practice, there is considerable diversity in the customer and product strategies of CDIs. Some sell their services on an exclusive basis, and others sell on a nonexclusive basis. For example, Catalina sells on an exclusive basis only to one grocery manufacturer in a particular category, in any given period, and within a geographic region. To define a market, Catalina divides a year into four 13-week periods, the United States into several regions, and retailer product offerings into more than 500 categories. Catalina focuses primarily on manufacturers for revenues; its retailer business revenues are less than 9% of its overall revenues (Catalina 10K, 2003). In contrast, Abacus and I-Behavior sell on a nonexclusive basis to any catalog marketer or specialty retailer that requests their services.

In addition, CDIs differ in their outlook toward increasing the accuracy of their targeting services. Catalina voluntarily restricts the length of transaction history used for couponing to a maximum of 65 weeks. Specifically, Catalina offers two types of targeting services: (1) Checkout Coupon, which is based on previous purchase data, and (2) Checkout Direct, which is based on 65 weeks of purchase history data. In contrast to Catalina, Abacus continues to expand the accuracy of its database. Abacus pools data from more than 1550 catalog marketers/specialty retailers on more than 90 million households and continues to increase the extent of household purchase information in its database. Abacus currently uses data for up to five years on each household in its database. When DoubleClick purchased Abacus in 1999, it attempted to improve the accuracy of the Abacus database further by combining the offline data from Abacus with online transaction behavior captured by DoubleClick. However, opposition from privacy advocates caused DoubleClick to back off from combining online and offline data.

It is possible that the existing strategies of an intermediary have arisen because of historical reasons but are not optimal in the current environment. For example, Catalina may have chosen an exclusive strategy because it served as a convenient sales pitch initially to prospective clients; that is, clients can have a competitive advantage by working with Catalina. However, with Catalina's current widespread acceptance by grocery manufacturers, exclusivity may no longer be necessary to win clients. In contrast, because Abacus uses a cooperative approach to collect data from its members, it may not be possible for Abacus to discriminate among its members by using a nonexclusive strategy. Similarly, Catalina's choice of restricting transaction histories to 65 weeks may have been due to the relatively high cost of storage two decades ago. Firms such as Abacus and I-Behavior may have been able to use longer histories because of their comparatively recent entry into these markets, by which time data storage costs had reduced considerably.

Can CDIs benefit from changing their current customer and product strategies? Currently, there is little research to guide them on what the optimal strategy should be. In this article, we offer an empirical framework to help an intermediary arrive at an optimal customer and product strategy. We illustrate the framework for a 1:1 coupon service firm, such as Catalina, using data from the ketchup market. Therefore, we tailor the details of the empirical modeling to the environment in which Catalina operates. However, the approach can be applied in other empirical contexts with appropriate modifications for the specific characteristics of that context. For example, the framework can be used to help answer whether DoubleClick should sell its targeted advertising services on an exclusive basis or a nonexclusive basis. For this, we need to calibrate the impact of advertising (as opposed to couponing) on the downstream firms' profitability, but the rest of the analysis would be similar. Such an approach would complement Iyer, Villas-Boas, and Soberman's (2005) theoretical analysis of targeted advertising.

The timeliness of our research questions is highlighted in a recent stock analysis report about Catalina by Deutsche Bank (Ginocchio, Chesler, and Clark 2005) on how this $1.3 billion market capitalization company can grow further, given that it has achieved virtually complete penetration at all major supermarkets in the United States. The report states (p. 16), "Categories are sold on four thirteen-week cycles with exclusivity (only one manufacturer can promote that category during that period). As Catalina believes that only approximately 20–25% of its customers want exclusivity, they are looking at ways to potentially sell more than one manufacturer in a category." Our approach provides Catalina an empirical basis on which to answer this critical business question.

In the grocery context, the retailer is the source of the customer purchase history data used for 1:1 coupon services. Catalina's intermediary business model is predicated on retailer cooperation. A natural question that arises is, What if the retailer chooses to disintermediate Catalina and offer the service itself? Large retailers with the appropriate infrastructure could easily implement such a targeting solution. Indeed, Tesco in the United Kingdom has been successfully collaborating with dunnhumby, a U.K.-based firm, in the development of 1:1 marketing services, including targeted couponing, over the last decade (Humby 2004; Humby, Hunt, and Phillips 2003). In the United States, dunnhumbyUSA is a joint venture between Kroger and dunnhumby that attempts to replicate dunnhumby's success in the United Kingdom with Tesco. We find that by providing targeting services, the retailer can also increase profits from the goods sold; therefore, the retailer has an incentive to undercut Catalina's price for the 1:1 service.

We also evaluate the profits for Catalina by providing 1:1 targeting services to retailers. We find that the profit from providing the 1:1 targeting service to the retailer is greater than that from providing the service to the manufacturers. This suggests that retailer services are potentially an underused revenue stream for Catalina. However, a practical problem in aggressively pricing retailer services is that retailers may balk at needing to provide the data and then paying for services using the same data. Therefore, Catalina may have only limited bargaining power to extract retailers' value from targeting compared with its power over manufacturers. This might explain why Catalina aggressively markets its manufacturer service compared with its retail services. Currently, retail services provide less than 9% of Catalina's total revenues, whereas manufacturer services provide more than 53% of revenues.

TRADE-OFFS IN CHOOSING THE OPTIMAL CUSTOMER AND PRODUCT STRATEGY

What are the trade-offs the CDI faces in deciding the optimal customer and product strategy? To fix ideas and to facilitate empirical work, we illustrate the trade-offs in the context of Catalina for the ketchup category, in which there are two main competitors: Heinz and Hunt's. Unlike standard products or services, for which the economic value to a customer is independent of who else uses it, the value of Catalina's 1:1 coupon service for Heinz depends on whether Heinz uses the service exclusively or whether Hunt's also uses it. This is because the effectiveness of a 1:1 coupon for Heinz in increasing sales is a function of whether Hunt's also offers targeted coupons.

Particularly notable is that the economic value of the service for Heinz may be higher or lower if Hunt's also uses the service (i.e., this service can have positive or negative externalities). If the service has positive externalities, it makes obvious sense for the firm to sell its service to both Heinz and Hunt's. If it has negative externalities, Catalina would need to evaluate whether the negative externalities for Heinz and Hunt's are sufficiently low to still sell to both; if not, it would sell the service on an exclusive basis only to one of them, depending on which company has the higher willingness to pay (higher economic value). Thus, the optimal customer strategy of whether to sell on an exclusive basis to Heinz or on a nonexclusive basis to multiple manufacturers is an empirical question for Catalina.

Thus far in this scenario, we have treated the "product" (i.e., the quality of the targeting that Catalina offers) as fixed. In the context of 1:1 marketing, the quality of the targeting is related to the accuracy with which a firm (e.g., Heinz) can identify the segment it wants to target. Catalina can increase the accuracy of targeting in several ways. It can (1) use demographic information, (2) increase the length of purchase history of households within a category at a cooperating retailer, (3) use information about purchasing behavior in other categories at the cooperating retailer to take advantage of cross-category similarities in purchase behavior (e.g., Ainslie and Rossi 1998; Iyengar, Ansari, and Gupta 2003), and (4) combine information about purchase behavior of households from other retailers. We consider only the first two options to improve targeting accuracy. We do not consider optimal targeting using cross-category purchase behavior, because it is computationally cumbersome and therefore beyond the scope of our analysis. We do not consider the option of pooling household purchase behavior across retailers, because Catalina is contractually obliged not to pool information across retailers. Catalina identifies households only by a retailer's internal identification number (e.g., from a loyalty program) and therefore cannot pool information across retailers.

What is the optimal product strategy for Catalina? For most products/services, firms would like to maximize their quality if doing so were relatively costless. However, 1:1 targeting is different in that increasing the quality of targeting may reduce the economic value of the service for the downstream clients. The idea is simple: If the targeting service is sold on an exclusive basis to only Heinz, the economic value of the targeting service for Heinz increases because Heinz can more effectively price discriminate its customers. However, if the targeting service is sold to both Heinz and Hunt's, the price discrimination effect of targeting can be overwhelmed by the more intense competition created by the targeting (e.g., Shaffer and Zhang 1995). Whether the price discrimination effect or the competition effect dominates is moderated by the level of targeting accuracy (Chen, Narasimhan, and Zhang 2001). At low levels of accuracy, price discrimination effects dominate competition effects, but at high levels of accuracy, competition effects dominate price discrimination effects. Thus, Catalina could potentially destroy economic value to downstream clients by increasing accuracy if it sold the product on a nonexclusive basis to both firms (Heinz and Hunt's). In such a case, to reduce the competition effect, Catalina may find it optimal to increase accuracy but sell on an exclusive basis only to one of the firms. Alternatively, it could reduce accuracy and sell to both firms and thus extract greater total revenues from both. Therefore, the customer strategy and the product strategy of a CDI are intertwined, and the empirical question of what is the optimal strategy for a particular CDI needs to be determined in the relevant empirical context.

Furthermore, note that theoretical models abstract away from many complexities of real-world markets, but these need to be accounted for in an empirical model. For example, theoretical models have typically allowed for household heterogeneity only on horizontal attributes, but in reality, households are also heterogeneous on vertical attributes. The empirical analysis needs to model the real-world demand and supply characteristics appropriate for the particular market to arrive at the correct product and customer strategies for the CDI.

LITERATURE REVIEW

This article is related to both theoretical and empirical research streams on 1:1 pricing. The literature refers to 1:1 pricing broadly, using terms such as "targeted couponing" and "behavior-based pricing." Using similar models, Thisse and Vives (1988) and Shaffer and Zhang (1995) show that in a competitive market, spatial discriminatory pricing or targeted coupons lead to a prisoner's dilemma compared with uniform pricing. These models assume symmetric firms. Shaffer and Zhang (2002) show that in the presence of firm asymmetry, higher-quality firms with larger market shares can improve volumes and profits due to gains in market share, though they continue to earn lower profit margins because of increased competition. Importantly, as we discussed previously, Chen, Narasimhan, and Zhang (2001) show that the level of targeting accuracy moderates the profits from 1:1 promotions. They show that there is an inverted U-shaped relationship between profitability and accuracy of targeting (personalization).

There is also a growing literature on behavior-based pricing (e.g., Chen 1997; Chen and Zhang 2004; Fudenberg and Tirole 2000; Shaffer and Zhang 2002; Villas-Boas 1999). In general, these studies find that behavior-based pricing leads to a prisoner's dilemma. Taylor (2003) and Villas-Boas (2004) highlight the effects of "strategic" consumers who alter purchasing behavior to avoid revealing their preferences. However, Shin and Sudhir (2006) develop a model in which customers are distinguished on the basis of their profitability (using the 80/20 rule), and behavior-based pricing does not lead to a prisoner's dilemma.

In terms of empirical research on 1:1 pricing, Rossi, McCulloch, and Allenby (1996) and Besanko, Dubé, and Gupta (2003) evaluate the profitability of targeted coupons. Rossi, McCulloch, and Allenby investigate how manufacturers can improve their profits with increasing levels of purchase history and demographic information. Unlike the current study though, they do not model the retailer or competition between manufacturers. Besanko, Dubé, and Gupta study only the profitability of targeting using previous visit data, but they model both competition and the retailer. However, unlike the current study, neither Rossi, McCulloch, and Allenby nor Besanko, Dubé, and Gupta investigate the 1:1 service provider's strategic decisions. Our analysis also finds that these two studies overestimate the profitability impact of personalization. This is because the models of consumer behavior used in computing profits with and without targeting are different. We discuss this issue in detail when we report the results on incremental profits. In terms of 1:1 advertising/communication, Ansari and Mela (2003) develop algorithms for how a firm should use consumer history to customize e-mail communications. Zhang and Krishnamurthi (2004) empirically study customized pricing in online stores.

The current study is also related to the theoretical literature on information suppliers and investments in information. Iyer and Soberman (2000) study how the marketing strategies of an information supplier are affected by downstream competition between firms that use product modification information. In a similar spirit, this article analyzes how the marketing strategies of the CDI are affected by how grocery manufacturers use the service downstream for 1:1 couponing. However, the model for how the data are used downstream is different between this study and that of Iyer and Soberman. Furthermore, our analysis is empirical and therefore requires the modeling of several specific features of the market (e.g., strategic interactions between manufacturers and retailers) that can be abstracted away in a theoretical analysis.

Other related theoretical works that address some issues in the current study are those of Chen and Iyer (2002) and Liu and Zhang (2006). Chen and Iyer study how firms differentially invest in customer addressability to avoid the negative effects of downstream competition. Liu and Zhang theoretically investigate the interaction between manufacturers and retailers in the presence of personalized pricing. Although a retailer's profits are lower in the presence of personalized pricing by either the manufacturer or the retailer, the retailer prefers to use personalized retail pricing to deter direct selling and personalized pricing by manufacturers, which leads to even worse outcomes for the retailer.

MODEL

Figure 1 represents a schematic of the grocery markets in which Catalina operates. There are four sets of agents involved in this market: (1) the CDI (e.g., Catalina), (2) the manufacturers, (3) a retailer, and (4) consumers.

Figure 1
SCHEMATIC OF THE MARKET

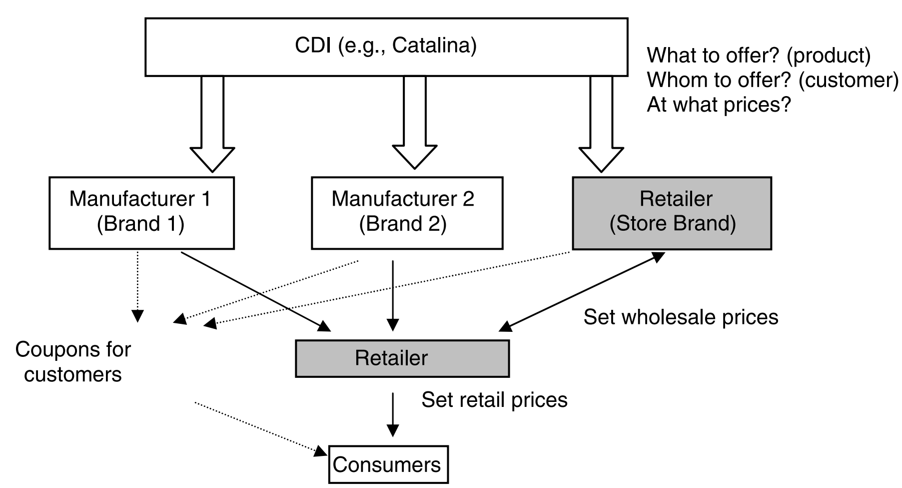

The model of manufacturers selling through a retailer to the consumer has been studied in previous research (e.g., Berto Villas-Boas 2007; Sudhir 2001; Villas-Boas and Zhao 2005). In these models, the pricing decisions of manufacturers and retailers are modeled as endogenous. Our model expands on this literature by endogenously modeling the decisions faced by a 1:1 coupon service provider (CDI) that facilitates targeted couponing to consumers in the market. Because Catalina is contractually obliged not to pool purchase history data across multiple retailers, the assumption that Catalina uses only data from one retailer for its targeting service is consistent with institutional reality. As in most previous research (e.g., Besanko, Dubé, and Gupta 2003; Besanko, Gupta, and Jain 1998; Sudhir 2001), we assume that the retailer is a local monopolist. Indeed, Berto Villas-Boas (2007) finds little evidence for cross-retailer competition at the single category level.

Figure 2 represents the schematic of the decision alternatives faced by a 1:1 coupon service provider, such as Catalina (CDI), regarding the sales of its services. We model the timing of the game into two phases: Phase 1, which involves the sale of 1:1 services, and Phase 2, which involves the sale of consumer goods.

Phase 1: Sale of 1:1 Services

Stage 1: Catalina's product choice decision: In this stage, Catalina decides on the length of purchase history it should optimally use for targeting. We consider three alternatives: (1) previous visit, which is similar to the concept of targeting used by Besanko, Dubé, and

Figure 2
DECISION TREE AND PAYOFFS

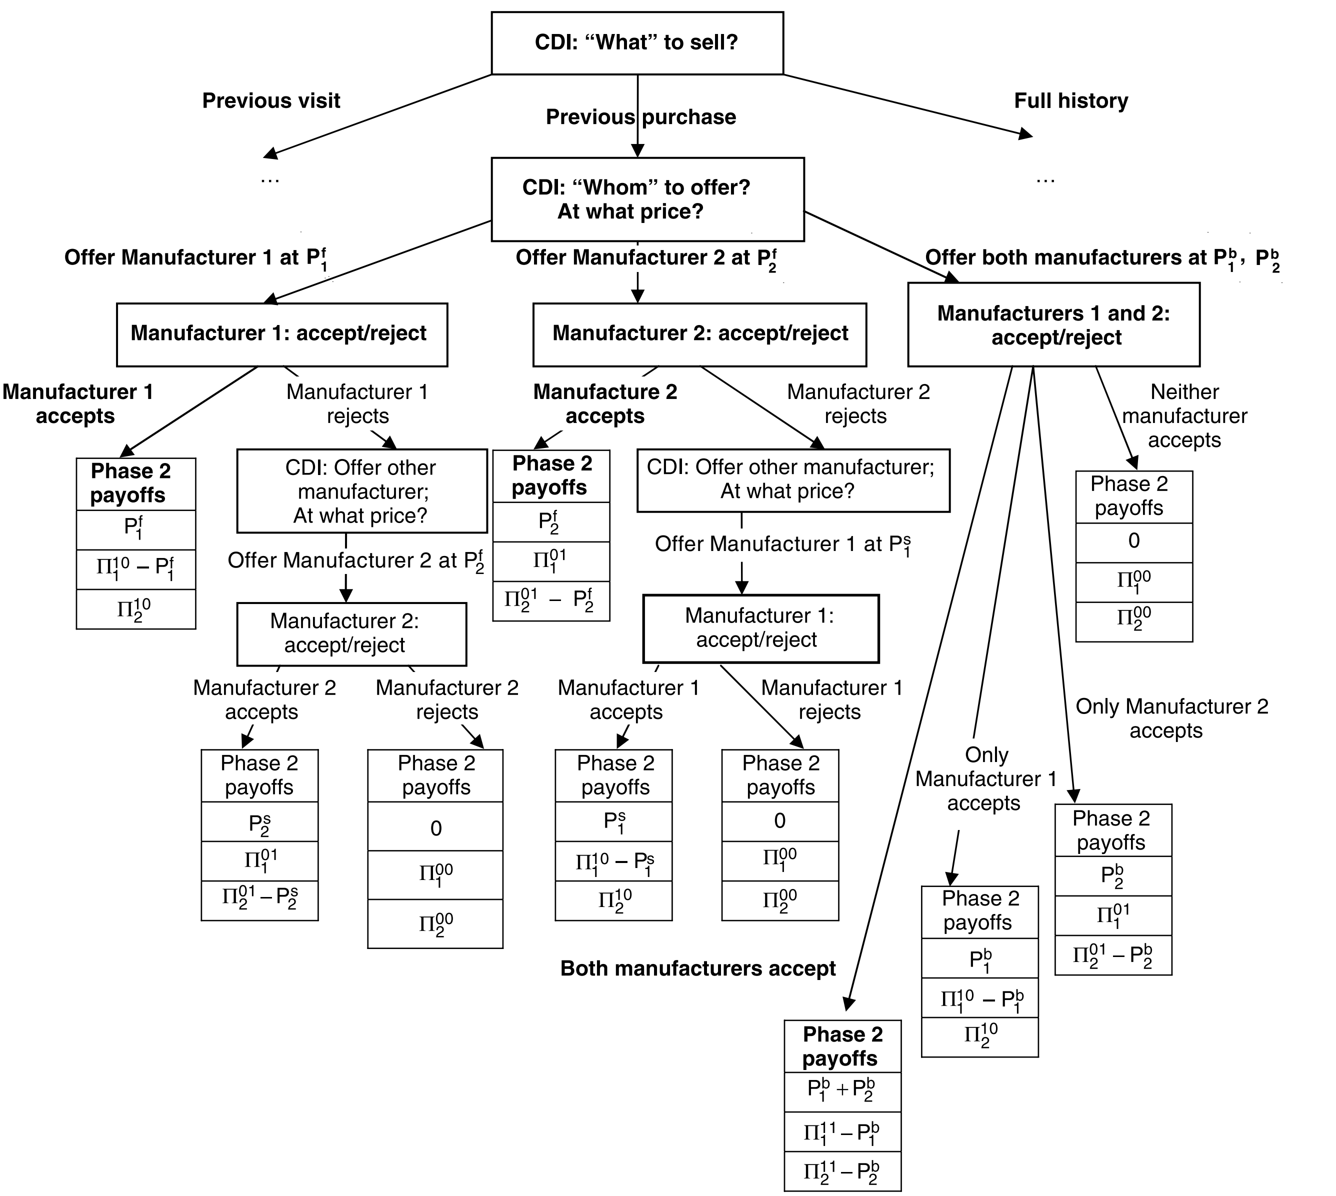

Gupta 2003; (2) previous purchase, as Catalina uses in its Catalina Coupon program; and (3) full purchase history, which is similar to what Catalina uses in its Catalina Direct program. However, rather than restricting the full purchase history to only 65 weeks as Catalina does, we evaluate different lengths of purchase history.

Stage 2: *Catalina's initial customer choice and price decision*: For ease of exposition, we consider a market with two national brand manufacturers. Catalina has three alternatives to make initial offers during this stage: (1) offer the 1:1 service to Firm 1 and set its price (p_1^f), (2) offer the 1:1 service to Firm 2 and set the price (p_2^f), and (3) offer the 1:1 service to both firms and set the prices to both firms (p_1^b, p_2^b).

The subscripts "1" and "2" on prices refer to the prices charged to Firms 1 and 2. The superscript "f" indicates that Firm 1 or 2 is "first" offered the service exclusively. The superscript "b" refers to the situation in which both firms are initially offered the service on a nonexclusive basis.

Stage 3: *Initial offer acceptance/rejection by manufacturers*: Manufacturers decide whether to accept or reject the offer of 1:1 services at the offered prices. In the case in which one firm is exclusively offered and accepts the offer, the manufacturers and retailers then move to the second "sales-of-goods" phase, and one of the firms has the capability to target. If both firms are offered initially, there are four possible outcomes: one of the two firms can accept, both can accept, and neither can accept. Given these outcomes, the manufacturers and retailers then move to the sales-of-goods phase, and the firms that have accepted the 1:1 service offers have the capability to target.

Stage 4: *Catalina offers 1:1 service to the "other" manufacturer at a second offer price*: If one firm is exclusively offered the 1:1 service first and rejects it, Catalina will offer the service to the other firm on an exclusive basis. For example, if Firm 2 receives the offer after Firm 1 rejects the initial offer of exclusive service, we denote the price to Firm 2 as p_2^s, where the superscript "s" indicates that Firm 2 was offered the service after Firm 1 refused.

Stage 5: *Second offer acceptance/rejection by manufacturers*: Manufacturers that received the second offer can either accept or reject the offer for the 1:1 service.

Given these decisions, the manufacturers and retailers then move to the second phase (sales of goods), and the firms that have accepted the 1:1 service offers have the capability to target. The payoffs realized after the second phase appear in three rows in Figure 2. We denote the profits from the sale of goods to manufacturer "f" as Π_f^{xy}, where x and y refer to the 1:1 service purchase decisions of Firms 1 and 2, respectively. A value of 1 (0) refers to whether the firm uses (does not use) the 1:1 services. The first row indicates the payoff to the 1:1 service provider (i.e., the price charged for 1:1 services), and the second and third rows indicate the payoffs to Firms 1 and 2, respectively, which show the net profits from the sale of goods and the fees paid (if any) to the 1:1 service provider.

Note that in this game of complete information, Stages 4 and 5 are in the off-equilibrium path, because Catalina will offer the right price in the initial offer so that whichever firm is offered initially will accept. We marked the equilibrium paths in bold. Thus, although there are ten payoff matrices shown, the only relevant payoffs in equilibrium are the three payoff matrices in which the firms that are initially offered the 1:1 service by Catalina accept the product. Nevertheless, the payoffs from the off-equilibrium paths are critically important for Catalina to determine what price it should charge the firms in Stage 2. This is because Catalina's offer price to the firms should account for the incremental profits a firm will make relative to the outcome in which the competitor obtains exclusive use of 1:1 services. Note that the price charged is not with respect to the situation in which there is no targeting at all. This is because the scenario in which neither firm purchases 1:1 coupons is not on the subgame perfect equilibrium path and therefore is not a credible alternative threat to either Firm 1 or Firm 2. This limits the amount of value that can be extracted from either firm by the 1:1 service provider. Thus, $P_1^f = \Pi_1^{10} - \Pi_1^{01}$, $P_2^f = \Pi_2^{01} - \Pi_2^{10}$, $P_1^b = \Pi_1^{11} - \Pi_1^{01}$, and $P_2^b = \Pi_2^{11} - \Pi_2^{10}$.

Phase 2: Sales of Goods

Stage 1: Manufacturers set wholesale prices and the coupon face values for individual households. If they have not purchased the 1:1 services, all households are assumed to have a coupon face value of zero.[1]

Stage 2: The retailer takes the information about wholesale prices and coupons issued into account when setting retail prices. Because the coupons are issued by the retailer, it is reasonable to assume that the retailer takes into account the coupons issued in setting retail prices.[2] We follow previous research (Besanko, Dubé, and Gupta 2003; Rossi, McCulloch, and Allenby 1996) in assuming that coupons are valid only for the week of issue.

Stage 3: Given the retail prices and coupons issued, the household makes buying decisions to maximize utility.

We now develop a detailed model of these three stages of Phase 2. We describe the decisions that each player—the consumer, the retailer, and the manufacturer—faces.

Consumer. A household i (i = 1, 2, ..., H) chooses one of J available brands (denoted by j = 1, ..., J) in the category or decides not to purchase in the category (j = 0, the no-purchase alternative or "outside good") on each household shopping occasion t = 1, 2, ..., n_i. Let the vector X_{ijt} denote all variables for brand j experienced by household i at shopping occasion t. This vector includes brand-specific indicators; marketing-mix variables, such as features and displays; and household-specific variables, which depend on the previous purchases, such as state dependence and household stock on occasion t.

Consumers choose the brand that offers the maximum utility. We specify the indirect utility of household i for brand j (j = 1, ..., J) on shopping occasion t as follows:

$$u_{ijt} = X_{ijt}\beta - r_{jt}\alpha + I_{it}\gamma + \xi_{jt} + \varepsilon_{ijt}, \tag{1}$$

[1]Technically, manufacturers set the wholesale prices, and Catalina decides whether to offer the coupon and what its face will be, but this distinction is unimportant for the results after the manufacturer makes the decision to purchase the targeting service.

[2]For brevity, we describe only the manufacturer Stackelberg model, though we also estimate the vertical Nash model.

where X_{ijt} includes all variables that affect household i's evaluation of brand j on occasion t (feature, display, and lagged brand choice) as well as time-invariant brand intercepts, r_{jt} is the price of brand j at occasion t, I_{it} is the inventory stock of household i in the category (across all brands) at time t,[3] ξ_{jt} is the brand j–specific effect on utility at shopping occasion t that affects all households but is unobserved by the econometrician, and ε_{ijt} is the unobserved utility of brands that vary over shopping occasions across households.

Because the indirect utility for any item in the choice set is identified only in terms of differences with respect to a base choice in the logit model, we treat the outside good as the base choice and normalize its utility as $u_{i0t} = \varepsilon_{i0t}$. We assume that the elements of the vector $\varepsilon_{it} = (\varepsilon_{i0t}, \varepsilon_{i1t}, \ldots, \varepsilon_{iJt})$ each follow an independent Gumbel distribution with a mean of 0 and a scale parameter of 1.

We model heterogeneity using a latent class framework (Kamakura and Russell 1989).[4] Consumers are probabilistically allocated to one of K segments, in which each segment k has its own parameter vector (α^k, β^k). The size of segment k is denoted as f^k, which can be interpreted as the likelihood of finding a consumer in segment k or the relative size of the segment in the population of consumers. The probability of household i that belongs to segment k choosing a brand j is given by

$$(2)\qquad S_{ijt}^k = \frac{\exp\left(X_{ijt}\beta^k - r_{ijt}\alpha^k + I_{it}\gamma^k + \xi_{jt}\right)}{\sum_l \exp\left(X_{ilt}\beta^k - r_{ilt}\alpha^k + I_{it}\gamma^k + \xi_{lt}\right)}.$$

Note that ξ_{jt} are the common demand shocks that affect all consumers. These are observable by the price-setting firms and consumers in the market but are unobservable by the researchers. Villas-Boas and Winer (1999) show that profit-maximizing firms account for ξ_{jt} when setting prices; therefore, price is correlated with ξ_{jt}. This causes a price endogeneity problem. Without correcting for endogeneity, the price coefficient will be biased toward zero. We discuss how we address this issue in the "Estimation and Solution Strategy" section.

Because f^k represents the likelihood of finding a consumer in segment k, we can compute the unconditional probability of choice for brand j by consumer i in period t as

$$(3)\qquad S_{ijt} = \sum_{k=1}^{K} f^k S_{ijt}^k = \sum_{k=1}^{K} f^k \left[\frac{\exp\left(X_{ijt}\beta^k - r_{ijt}\alpha^k + I_{it}\gamma^k + \xi_{jt}\right)}{\sum_l \exp\left(X_{ilt}\beta^k - r_{ilt}\alpha^k + I_{it}\gamma^k + \xi_{lt}\right)}\right].$$

When the estimates of the latent class model are obtained, we can then apply the Bayes' rule on the aggregate latent class estimates using each household's purchase history that is available. We obtain the posterior probability that a consumer i belongs to a segment k conditional on observed choice history H^i by revising the prior probability of membership f^k in a Bayesian manner (Kamakura and Russell 1989):

$$(4)\qquad \Pr(i \in k|H^i) = \frac{L\left(H^i|k\right)f^k}{\sum_{k'} L\left(H^i|k'\right)f^{k'}}.$$

Using different levels of household choice history will result in different levels of posterior probability for each consumer i. The posterior probability using the entire purchase history for the consumer i, which we denote as H_{FH}^i, plays an important role in our analysis, and we write the corresponding posterior probability as

$$(4a)\qquad \Pr\left(i \in k|H_{FH}^i\right) = \frac{L\left(H_{FH}^i|k\right)f^k}{\sum_{k'} L\left(H_{FH}^i|k'\right)f^{k'}}.$$

Retailer. The retailer's goal is to maximize category profits in period t, given the manufacturers' decisions to buy 1:1 services. Let x = 1(0) denote whether Manufacturer 1 has purchased (not purchased) the personalization service. Similarly, let y = 1(0) denote whether Manufacturer 2 has purchased (not purchased) the personalization service. Therefore, the retailer chooses retail prices $r_{1t}^{xy}, \ldots, r_{Jt}^{xy}$ conditional on which firms have purchased the 1:1 service to solve the following problem:

$$(5)\qquad \max_{r_{1t}^{xy}, \ldots, r_{Jt}^{xy}} \prod\nolimits_{Rt}^{xy} = \sum_{j=1}^{J} \sum_{i=1}^{N_t} \left[r_{jt}^{xy} - w_{jt}^{xy}\right] S_{ijt}\left(r_{jt}^{xy} - D_{ijt}^{xy}\right) M_t.$$

Note that we use the brackets for grouping terms and the parentheses for denoting arguments of function. For example, in Equation 5, the right-hand side consists of (1) the retail margin $(r_{jt}^{xy} - D_{ijt}^{xy})$; (2) the share $S_{ijt}(r_{jt}^{xy} - D_{ijt}^{xy})$, which is a function of the effective price net of individual-specific discounts $(r_{jt}^{xy} - D_{ijt}^{xy})$ the consumer faces; and (3) M_t, the total market size in time t. The shares $S_{ijt}(r_{jt}^{xy} - D_{ijt}^{xy})$ in Equation 5 are the weighted averages of the segment-specific shares across the k segments at the effective price faced by the consumer of $(r_{jt}^{xy} - D_{ijt}^{xy})$.[5]

[3]We calculated inventory as the stock of the relevant category (ketchup) that has accumulated at the household from previous purchases; the stock is depleted at the average consumption rate of the household for ketchup. The method of calculating inventory is similar to Gupta's (1988) approach. In our model, the utility of choosing the outside good, rather than being set to 0 as Chintagunta (2002) does, is parameterized by the ketchup inventory stock. In the profit simulations, the probability of purchase in the future is affected by the simulated purchase because the inventory variable is updated and affects the probability of the outside good choice in the future.

[4]The latent class model with discrete segments has considerable empirical validity and managerial relevance (Wedel and Kamakura 2000). Andrews, Ainslie, and Currim (2002) find that both the discrete and the continuous heterogeneity distributions fit the data about equally well, though some studies have argued that continuous heterogeneity coupled with discrete heterogeneity can fit the data better (Allenby, Arora, and Ginter 1998). We apply the latent class approach because of its computational tractability when solving for the equilibrium targeting prices after competitive and retailer reactions are incorporated into the model.

[5]We assume that the retailer uses the same level of information about consumers as that which the manufacturer buys from Catalina. Here, we follow the assumption made in the previous literature on targeting (e.g., Besanko, Dubé, and Gupta 2003) that both manufacturers and retailers use the same level of information. If we relax this assumption to allow the retailer to use customer data, when the manufacturer does not buy the CDI's service, this leads to a "common knowledge" problem because

Taking the first-order conditions of Equation 5 with respect to retail prices, we obtain the retailer's pricing equation for each product in the category in terms of wholesale prices. The details of the derivation appear in Part A of the Web Appendix (see http://www.marketingpower.com/content84060.php). The retailer price equation can be derived as

$$R = W - \left(\sum_{i=1}^{N} \Theta_R^i\right)^{-1} \times \left(\sum_{i=1}^{N} S_i\right), \tag{6}$$

where Θ_R^i is the matrix of first derivatives of all the (individual consumers') shares with respect to all retail prices (retail prices are common across consumers), with element (j, m) = $[\partial S_{im}(r_m)]/\partial r_j$; R is the vector of retail prices; W is the vector of wholesale prices (which are common across all consumers); and S_i is the vector of shares for each consumer i over all the brands:

$$R \equiv \begin{bmatrix} r_1 \\ \vdots \\ r_J \end{bmatrix}_{Jx1}, \; W \equiv \begin{bmatrix} w_1 \\ \vdots \\ w_J \end{bmatrix}_{Jx1}, \; S^i \equiv \begin{bmatrix} S_1^i \\ \vdots \\ S_J^i \end{bmatrix}_{Jx1}.$$

Manufacturer. A manufacturer m offering a subset $\aleph_m$ of brands in the market sets the wholesale price w_{jt}^{xy} (where $j \in \aleph_m$) and the coupon face values to individual households (D_{ijt}^{xy}) to maximize its profits. A manufacturer that has not been sold the personalization service will have coupon face values set to zero. The manufacturer accounts for the knowledge that retailer prices ($r^{xy}{}_{jt}$) are set to reflect the wholesale prices and the coupon face values that have been issued to individual households. The profit of manufacturer m at time t from the sales of goods is given by

$$\prod_{mt}^{xy} = \sum_{j \in \aleph_m} \sum_{i=1}^{N_t} \left(w_{jt}^{xy} - D_{ijt}^{xy} - c_{jt}\right) S_{ijt}\left\{\left[r_{jt}^{xy}\left(w_{jt}^{xy}, D_{ijt}^{xy}\right) - D_{ijt}^{xy}\right]\right\} M_t, \tag{7}$$

where c_{jt} is the marginal cost of the manufacturer for brand j in period t and $S_{ijt}^{xy}[r_{jt}^{xy}(w_{jt}^{xy}, D_{ijt}^{xy}) - D_{ijt}^{xy}]$ is the probability of household i buying brand j in period t given the decisions of Manufacturers 1 (denoted by x) and 2 (denoted by y) to buy the purchase history data. Note that the retailer sets the retail price after accounting for both the wholesale price (w_{jt}^{xy}) and the vector of discounts offered to all households (i.e., $D_{jt}^{xy} = \{D_{ijt}^{xy}\}_{i=1}^{H}$).

We can write the manufacturer profit equations at the individual level as follows:

$$\prod_{mt}^{xyi} = \sum_{j \in \aleph_m} \left(w_{jt}^{xy} - D_{ijt}^{xy} - c_{jt}\right) S_{ijt}\left[r_{jt}^{xy}\left(w_{jt}^{xy}, D_{ijt}^{xy}\right) - D_{ijt}^{xy}\right].$$

Taking the first-order conditions of Equation 6, with respect to $w_{jt}^{xy} = w_{jt}^{xy} - D_{ijt}^{xy}$, we can solve for the effective margin from each household. The wholesale price will be $w_{jt}^{xy} = \max_i w_{ijt}^{xy}$ and $D_{ijt}^{xy} = w_{jt}^{xy} - w_{ijt}^{xy}$.

From the manufacturer first-order conditions, we can write the manufacturer margin from a particular household i ($W_i - C$) as follows:

$$(W_i - C) = (O_W \times \Theta_W^i)^{-1} \times (-S^i), \tag{8}$$

where Θ_W^i is defined for each individual consumer such that it contains the first derivatives of all the (individual consumers') shares with respect to all wholesale prices (wholesale prices are common across consumers), with element (j, m) = $[\partial S_{im}(r_m - D_{im})]/\partial w_{ij}$. To account for the set of brands owned by the same manufacturer, we define the manufacturer's ownership matrix O_W such that element (j, m) is equal to 1 if the manufacturer that sells brand j also sells brand m and 0 if otherwise, where $[O_W \times \Theta^i{}_W]$ is the element-by-element multiplication of the two matrices, W_i is the vector of wholesale prices less the individual coupon values, C is the vector of marginal costs of the manufacturer (C is common across all consumers), and S^i is the vector of shares for each consumer i:

$$W_i = \begin{bmatrix} w_1 - D_{i1} \\ \vdots \\ w_J - D_{iJ} \end{bmatrix}_{JX1}, \; C = \begin{bmatrix} c_1 \\ \vdots \\ c_J \end{bmatrix}_{JX1}, \; S_i = \begin{bmatrix} S_{i1} \\ \vdots \\ S_{iJ} \end{bmatrix}_{JX1}.$$

We detail the derivation in Part A of the Web Appendix (see http://www.marketingpower.com/content84060.php). Note that though the assumed demand models entering the objective functions of the manufacturer and the retailer (and the chosen optimal wholesale prices, household discounts, and retail prices) reflect the level of information that is available to the market participants according to whether they have accessed the 1:1 marketing service, the actual demand and the profits resulting from such pricing strategies will reflect the "true" behavior of the consumer (which we approximate using estimates of the consumer's full purchase history, H_{FH}^i). We elaborate on this further when we report the profits to manufacturers and retailers from using 1:1 targeting. We specify manufacturer marginal cost as a function of factor prices, which assumes a fixed-proportions production technology:

$$c_{jt} = \lambda_j + \theta \times B_t + \upsilon_{jt}, \tag{9}$$

where B_t are the factor prices, λ_j are brand-specific intercepts, and υ_{jt} is the cost shock.

ESTIMATION AND SOLUTION STRATEGY

The solution strategy consists of five steps. The first two steps involve estimation to characterize the market, and the remaining three steps involve policy simulations to infer the optimal strategy for the CDI.

In Step 1, we estimate the demand and supply model discussed previously. The demand model is a latent class model of household preferences and responsiveness to marketing mix with alternative levels of purchase history lengths used to proxy for personalization quality from consumer information.[6] To account for potential price endo-

manufacturers do not know or cannot infer the demand model the retailer uses to set retail prices. Advances in game theory beyond the scope of this article are needed to solve this problem.

[6]Other aspects of consumer information, such as consumer demographics, could potentially improve the quality of the personalization service,

geneity concerns, we use the control function approach that Petrin and Train (2003) developed. The control function approach has similarities to that of Rivers and Vuong (1988) and Villas-Boas and Winer (1999). Essentially, we obtain residuals from a regression of prices of the different brands against their cost factors and include these residuals in the utility equation (Equation 1) in estimating the demand model. We explain more details of the control function approach in Part B of the Web Appendix (see http://www.marketingpower.com/content84060.php). Given the demand estimates, we can compute the wholesale and retail margins using Equations 6 and 8. Then, the cost estimation reduces to a linear regression, where the dependent variable is (retailer price – computed retail margin – computed wholesale margin) and the independent variables are the cost factors and the brand dummies.

In Step 2, we apply Bayes' rule on the aggregate latent class estimates using each household's purchase history (the length of history varies depending on the scenario being considered and the number of visits of the household during the estimation period) to obtain household-level probabilities of membership in each latent class. When purchase histories are short, the individual-level probabilities differ little from the aggregate probabilities, and as the purchase histories lengthen, the individual probabilities tend to become more different from the aggregate probabilities, thus reflecting more closely idiosyncratic household preferences. Manufacturers may use varying levels of information about consumers' purchase histories in targeting them. Note that the demand model is no longer as it is in Equation 3. The demand equation will replace the segment probability f^k in Equation 3 with the household-level segment probability in Equation 4 (or, for the full history case, in Equation 4a). Thus, the share equation a manufacturer uses to target a consumer i, conditional on observed choice history H_i, is as follows:

$$
\begin{aligned}
(10)\qquad S_{ijt}^{H_i} &= \sum_{k=1}^{K}\left\{\left[\frac{L\left(H^i|k\right)f^k}{\sum_{k'}L\left(H^i|k'\right)f^{k'}}\right]\times S_{ijt}^{k}\right\}\\
&= \sum_{k=1}^{K}\left\{\left[\frac{L\left(H^i|k\right)f^k}{\sum_{k'}L\left(H^i|k'\right)f^{k'}}\right]\right.\\
&\quad\left.\times\left[\frac{\exp(X_{ijt}\beta^k - r_{ijt}\alpha^k + I_{it}\gamma^k + \xi_{jt})}{\sum_{l}\exp(X_{ilt}\beta^k - r_{ilt}\alpha^k + I_{it}\gamma^k + \xi_{lt})}\right]\right\}.
\end{aligned}
$$

In Step 3, having thus characterized the household-level preferences using different lengths of purchase history data, we solve for the optimal prices and discounts under alternative targeting scenarios (exclusive and nonexclusive). To obtain steady-state profit estimates, we solve for prices and discounts over a large number of weeks, tracking both consumer prior purchases (to account for state dependence effects) and inventories (to account for inventory effects on category purchases) over this period. In solving for the equilibrium prices and discounts, we account not only for the pricing behavior of the manufacturers but also for retailers' equilibrium pass-through behavior. In this simulation, we use the same marketing-mix variables for features and displays as in the estimation data.

In Step 4, given the optimal prices and discounts computed in Step 3, we evaluate manufacturer profits on the basis of consumer choices at the optimal prices and discounts. Note that optimal prices and discounts will vary depending on the available purchase history and which firms do the targeting. However consumer behavior should be based on the same "true" preferences regardless of what data firms have. Thus, in predicting consumer choice, given the chosen prices and discounts, it is critical to always use the household-level estimates obtained with the full purchase history data because these are the best estimates of the true household behavior. The estimates obtained with shorter purchase histories should not be used at this stage, because this will grossly overstate the profitability of targeting. On first glance, this issue may appear a "mere detail," but we find that the improvements in profits in prior empirical studies (Besanko, Dubé, and Gupta 2003; Rossi, McCulloch, and Allenby 1996) can be overstated if we do not assume a true, stable consumer behavior based on the full purchase history.

In Step 5, given the profits obtained under alternative targeting scenarios of history length (full purchase history, previous purchase only, previous visit only, and no targeting) and client choice (exclusive and nonexclusive), we solve for the optimal customer and product strategy for the CDI.

EMPIRICAL ILLUSTRATION

Data

We use the ACNielsen scanner panel data on the ketchup category from the largest retailer in the Springfield, Mo., market for the empirical illustration. We restrict attention to the four largest brand sizes, which collectively account for 64% of the sales in this category—Heinz (32 ounces and 28 ounces), Hunt's (32 ounces), and the store brand (32 ounces)—and use 100 weeks of purchase history data from 1986 to 1988. We use a sample of 143 households that made at least five purchases of the chosen brand sizes during the 100 weeks of analysis. The 143 households bought ketchup in 1073 of 11,660 store visits. We provide a summary of brand shares (conditional on purchase) and prices in Table 2. We use the price of tomatoes as a cost factor. We obtained the price data from the Bureau of Labor Statistics. We obtained part of the data from the Web site and the rest through e-mail from bureau officials.

Estimation Results

The Bayesian information criterion suggests that a three segment latent class model is the best model.[7] The identifi-

but the incremental impact of demographics over purchase history was miniscule in our analysis. Thus, we focus on purchase history length as a measure of accuracy and omit demographics in further analysis. This is consistent with the findings in Rossi, McCulloch, and Allenby (1996).

[7]Because the Akaike information criterion and Bayesian information criterion were worse for the model that included demographic and seasonality variables, we report only results of the best-fitting model without demographic and seasonality variables. We also considered serial correla-

Table 2
DESCRIPTIVE STATISTICS FOR KETCHUP DATA

	Conditional Brand Share (%)	*Price (Dollars per 10 Ounces)*	*Feature*	*Display*
Heinz (32 ounces)	37	.41	.07	.11
Hunt's (32 ounces)	13	.42	.02	.01
Heinz (28 ounces)	22	.50	.04	.09
Store brand (32 ounces)	28	.28	.12	.12

cation of the latent class logit model with exogenous variables is standard. However, price is endogenous, and as discussed previously, we use a two-step control function approach to obtain unbiased estimates of the price coefficient.[8] We used a first-stage regression of ketchup prices with brand intercepts and factor costs (cost of tomatoes). The key identifying assumption is that the factor costs are independent of the demand shocks. The F-statistic for the tomato cost is 7.8, which is significant at the 5% level. Interacting tomato cost with the brand dummies, as in the work of Villas-Boas and Zhao (2005), caused the F statistics to become nonsignificant. Thus, we use a common slope coefficient across brands in the first-stage regression. We also considered other cost factors, such as wages and cost of packaging materials (glass and plastic), as instruments, but we did not find these to be effective instruments.

We use the residuals of the first-stage regression as an additional variable in the utility equation to estimate the demand model. The demand estimates appear in Table 3. Segment 2 is the least price sensitive, but according to the negative coefficients associated with the intercept, it also purchases least in the category. It is 24% of the market. Segments 1 and 3 are more price sensitive than Segment 2 and, together, constitute 76% of the market. However Segment 1 is relatively more loyal to Heinz (32 ounces). Segment 3's preferences are more diffused across all brands, and it is the most price sensitive segment in the market, suggesting the least amount of loyalty. It was also relatively insensitive to inventory levels. This suggests that Segment 3 does not purchase ketchup at regular intervals but opportunistically buys any brand when it is on sale.

tion in the effects of marketing variables and the error terms using the geometric decay approach that Seetharaman (2004) outlines, but these did not improve model fit.

[8]We tested for possible endogeneity of features and displays using Hausman's (1978) method and found that we cannot reject the null hypothesis that feature and display are exogenous even at the 10% level of significance. The test statistic is 14.2, and the critical value at a 5% (10%) significance level of the chi-square distribution with 27 degrees of freedom is 40.1 (36.7).

The price elasticities for the three-segment latent class demand model appear in Table 4. The own- and cross-price effects are as we expected. Hunt's (32 ounces) and the store brand (32 ounces) have higher own-elasticities than the two Heinz brand sizes. Heinz (28 ounces), the most expensive brand, has the lowest own-elasticity. Hunt's (32 ounces) and the store brand (32 ounces) have higher cross-elasticities, indicating that switching would be higher between these brand sizes. An increase in the price of the largest brand size, Heinz (32 ounces), will result in more substantial substitution to Hunt's (32 ounces) and the store brand (32 ounces) than to Heinz (28 ounces).

Given the estimates of the demand model, we now estimate the supply model. We test for the appropriate

Table 3
DEMAND MODEL ESTIMATES

Parameter	*Segment 1 (47%) Estimate (SE)*	*Segment 2 (24%) Estimate (SE)*	*Segment 3 (29%) Estimate (SE)*
Heinz (32 ounces)	1.90***	–2.28**	1.89***
	(.47)	(.13)	(.45)
Hunt's (32 ounces)	.60	–3.33***	3.14***
	(.56)	(.16)	(.50)
Heinz (28 ounces)	.80	–1.85***	2.95***
	(.62)	(.15)	(.59)
Store brand (32 ounces)	–.50	–5.66***	1.72***
	(.43)	(.26)	(.35)
Price	–13.23***	–2.89***	–16.91***
	(1.25)	(.19)	(1.22)
Feature	.90***	.74***	–.11
	(.14)	(.22)	(.12)
Display	.51***	.33*	.07
	(.14)	(.19)	(.12)
Inventory	–3.19***	–1.09***	–.16
	(.35)	(.10)	(.27)
State dependence	.62***	1.42***	1.24***
	(.18)	(.18)	(.17)
Price residual	.50***	–.04	–.37**
(Heinz [32 ounces])	(.18)	(.20)	(.20)
Price residual	.29	.07	1.70***
(Hunt's [32 ounces])	(.40)	(.66)	(.42)
Price residual	–.21	.10	.11
(Heinz [28 ounces])	(.29)	(.13)	(.26)
Price residual	.41	2.36	1.04***
(Store brand [32 ounces])	(.44)	(1.53)	(.31)

*$p < .1$.
**$p < .05$.
***$p < .01$.

Table 4
MEAN PRICE ELASTICITIES FOR THE THREE-SEGMENT MODEL

	Change in Share			
Change in Price	*Heinz (32 Ounces)*	*Hunt's (32 Ounces)*	*Heinz (28 Ounces)*	*Store Brand (32 Ounces)*
Heinz (32 ounces)	–3.52	.06	.04	.05
Hunt's (32 ounces)	.03	–5.18	.02	.09
Heinz (28 ounces)	.03	.03	–2.20	.03
Store brand (32 ounces)	.03	.11	.02	–4.14

manufacturer–retailer interaction (manufacturer Stackelberg and vertical Nash) and manufacturer–manufacturer interaction (Bertrand and collusion). The best-fitting model ($p < .01$) is the manufacturer Stackelberg model with manufacturers in Bertrand competition. For this supply model, we report the estimates of the cost factors in the cost equation in Table 5. The estimates suggest that Heinz and the store brand have lower marginal costs than Hunt's (though the differences are not significant). Not surprisingly, tomato prices have a significant effect on marginal cost of ketchup.

ANALYSIS OF THE CDI'S DECISIONS

Given the demand and cost estimates from the previous section, we now evaluate the profitability of the alternative decision scenarios from the 1:1 service provider's perspective using simulations. We simulate the market for 100 weeks, which is a sufficiently long period to obtain stable estimates of profits under alternative decision scenarios.[9] From the household-level demand model, we get the market share of the sample customers. We then scale this sample market share by the chain's volume of sales in the week to arrive at chain profits.

We first demonstrate how length of purchase history affects the ability to use 1:1 promotions. We then evaluate the manufacturers' (Heinz and Hunt's) profits from the sale of goods as a function of whether they used 1:1 coupons on either an exclusive or a syndicated basis; that is, we compute the payoffs (Π_1^{10}, Π_1^{01}, Π_2^{10}, Π_2^{01}, Π_1^{11}, Π_2^{11}) for different lengths of purchase history. Using these payoffs, we infer what price the CDI can charge under different scenarios and thus arrive at the optimal customer and product decisions of the 1:1 services vendor.

How Length of Consumer Purchase History Affects 1:1 Targeting

Figures 3 shows the distribution of posterior probabilities of households belonging to Segment 1 when using previous visit, previous purchase, and full history data for targeting. Figure 3, Panel A, shows that the marketer achieves little discrimination across consumers by using only information about the previous visit, because the vast majority of consumers are classified in the same quintile as the aggregate probability (f^k in Equation 7) (i.e., .47 for Segment 1). The previous purchase information enables more discrimination to be achieved between consumers, as Figure 3, Panel B, shows. We achieve much better discrimination among consumers by using 100 weeks of consumer purchase information, as the polarized probabilities in Figure 3, Panel C, show. With 100 weeks of information, approximately 40% of consumers are assigned with a high degree of probability (posterior probability in the highest quintile) to Segment 1, and more than 40% of consumers are not assigned to Segment 1 with a high degree of probability (posterior probability in the lowest quintile).

The Effect of 1:1 Coupons on Manufacturer Client Profits

We now assess the profitability of 1:1 targeting for manufacturer clients (Heinz and Hunt's). An important factor in estimating the profits of manufacturers is the assumption

[9]Average profits per week were stable with consumer choices simulated over 100 weeks. Increasing the period of simulation further had no effect on the results but simply increased computation time.

Table 5
COST EQUATION ESTIMATES

Parameter	Estimate (SE)
Heinz (32 ounces)	.036 (.072)
Hunt's (32 ounces)	.088 (.072)
Heinz (28 ounces)	.038 (.073)
Store brand (32 ounces)	–.041 (.072)
Tomatoes	.152* (.063)

*$p < .01$.

Figure 3
HOW PURCHASE HISTORY LENGTH AFFECTS ESTIMATED SEGMENT 1 PROBABILITIES

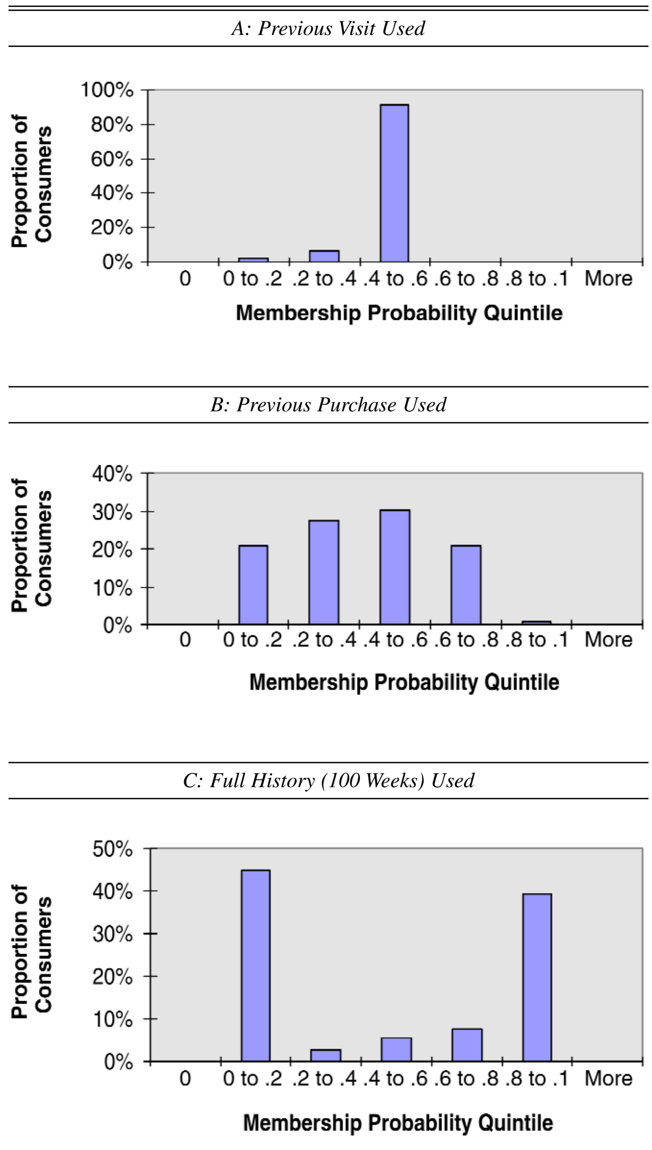

about retailer behavior. We compare the results using two assumptions about the retailer: (1) The retailer is a category profit maximizer, and (2) the retailer charges a simple constant markup over wholesale prices (e.g., Silva-Risso, Bucklin, and Morrison 1999); we illustrate this with a markup of 25%. The improvement in profits from targeting for Heinz is much greater when the retailer uses a constant markup strategy (9%) than when the retailer uses an optimal category profit maximization strategy (2%). However, for Hunt's, the increase in profits from targeting is low (under 1%) under both retailer strategies. In practice, retailers are expected to be somewhere in between the two extremes in their pricing sophistication; therefore, we can expect the true benefits of targeting for manufacturers to lie between these bounds.

In the rest of the analysis, we assume that the retailer follows the optimal category profit maximization strategy. We report the profits under the different scenarios in Table 6. Several insights emerge.

First, 1:1 promotions by both firms increase profits compared with no targeting for all levels of data length (previous visit, previous purchase, and full purchase history). As the reported t-statistics show, these increases are statistically significant. Furthermore, increasing accuracy (from previous visit to previous purchase and from previous purchase to full history) improves profits. The increases are also statistically significant.[10]

Thus, the positive price discrimination effect of targeting dominates the negative competitive effect of targeting in this market. Even with the full purchase history of 100 weeks and competitive targeting, we have not reached the peak of the inverted U-shaped relationship between targeting accuracy and profitability (see Chen, Narasimhan, and Zhang 2001).

Second, we compare the case in which only one firm exclusively targets versus the case in which both firms target. Under 1:1 targeting using full purchase history, both Heinz and Hunt's make more profits when both firms target than when either firm targets alone. Thus, there is a positive externality from the use of 1:1 targeting for both Heinz and Hunt's in this market.

Finally, we examine the magnitudes of the improvements in profits from the use of targeting. The maximum profit gain that any firm obtains by using targeted pricing in the ketchup category is approximately 2%. An improvement of gross margins by 2% can be a substantive increase in net profits. For example, according to Hoover's Online, Heinz had a gross margin of 40% and a net margin of 10% in 2003. A 2% increase in gross margin can then translate into an increase of approximately 8% in net margins. As we discussed previously, the 2% increase is a conservative lower bound in the presence of a sophisticated retailer maximizing category profits. The profits can be greater if the retailer is less sophisticated in its pricing.

1:1 Targeting Profits: Measurement Issues

The profit increases from targeting we report are smaller than the profit increases that Rossi, McCulloch, and Allenby (1996) and Besanko, Dubé, and Gupta (2003) report. Using full purchase history data (without demographics), Rossi, McCulloch, and Allenby find an increase of 5% for one item in the tuna category. Besanko, Dubé, and Gupta find improvements of 4% for Heinz and 37% for Hunt's in the ketchup category, with only previous visit data. We detail three key modeling issues that can explain these differences. First, we include inventory in the demand model, whereas Rossi, McCulloch, and Allenby and Besanko, Dubé, and Gupta do not. Although they do not have inventory data, Besanko, Dubé, and Gupta find suggestive evidence that inclusion of inventories can reduce the potential incremental gain in profits significantly. Category purchase will be overestimated when the effect of inventory is not included in the demand model. In other words, the absence of inventory in their model implies that consumers who purchase during the previous period are still likely to purchase at the same level in the current period. This overestimates the benefits of accurate price targeting. Rossi, McCulloch, and Allenby use a conditional choice model, so they do not model inventory issues.

Second, the assumption about retailer pricing behavior has an impact on profitability of targeting. Rossi, McCulloch, and Allenby (1996) do not consider competitive manufacturer or retailer reaction to targeting. As we discussed previously, the retailer reaction has an effect on the benefits of targeting; when the retailer charges a constant markup, we found that Heinz profits increase by approximately 9%, a magnitude comparable to that in Rossi, McCulloch, and Allenby. Therefore, this issue needs to be explored further.

Third, a consistent standard of consumer purchase behavior should be maintained in the computation of targeting profits. Besanko, Dubé, and Gupta (2003) compare profits with no targeting using aggregate data and targeting using previous visit data. However when computing profits under the two scenarios, they assume different consumer behavior that is consistent with the level of detail of data available for targeting. However, because consumer behavior should be invariant to the level of data used to estimate preferences, we use the estimates obtained with full history data as our best approximation of true consumer behavior for targeting with different levels of purchase history.

Table 7 illustrates the magnitude of the bias in the estimates of targeting profits when a consistent standard is not adopted. The first two rows illustrate that using just the information about consumers in characterizing consumer response can result in an increase in profit estimates by 10.02% for Heinz and .56% for Hunt's.[11] These two rows are for situations in which neither firm targets. The difference is purely a bias introduced as a result of posterior allo-

[10]We use bootstrapping to compute the standard errors and t-statistics. We take 30 draws from the distribution of the demand estimates and compute the difference in profits under targeting and no-targeting scenarios for each draw. We perform a paired t-test based on the difference in profits for each draw under the targeting and no-targeting scenarios. Because the profits from targeting are better than the profits from not targeting for most draws, the t-statistics are relatively high, even when the profit increases are small. We also find that the profits for both Heinz and Hunt's using 100 weeks of data are higher (and the difference is statistically significant at $p < .01$) than the profits using 65 weeks of history (which Catalina currently uses in Checkout Direct). Note that longer purchase histories can be better for 1:1 marketing, provided that consumer preferences do not change over time. Nonstationary consumer preferences could actually make longer purchase histories a liability, an extreme example of which could be consumer preferences in the fashion industry.

[11]This is consistent with the finding of Shaffer and Zhang (2002) that the benefits of price discrimination are greater for the larger market share firm.

Table 6
INCREMENTAL PROFITS FROM 1:1 COUPONS

		Previous Visit		*Previous Purchase*				*Full History*			
Targeting by		*Heinz*	*Hunt's*	*Heinz*	*t-Statistic Relative to Previous Visit*	*Hunt's*	*t-Statistic Relative to Previous Visit*	*Heinz*	*t-Statistic Relative to Previous Purchase*	*Hunt's*	*t-Statistic Relative to Previous Purchase*
Neither	Profits	73,301	4,139	73,301		4,139		73,301		4,139	
Hunt's only	Profits	73,317	4,139	73,301	–7.80*	4,139	–.99ns	73,260	–2.92*	4,154	8.22*
	Increase relative to no targeting (%)	.02	.00	.00		.00		–.06		.36	
	t-statistic relative to no targeting	8.06*	–.74ns	20.52*		–1.79ns		–2.84*		8.13*	
Heinz only	Profits	73,304	4,139	73,620	38.41*	4,145	15.09*	74,534	27.64*	4,139	–7.27*
	Increase relative to no targeting (%)	.00	.01	.43		.16		1.68		.00	
	t-statistic relative to no targeting	3.43*	8.82*	37.36*		16.19*		35.84*		1.39ns	
Heinz and Hunt's	Profits	73,338	4,139	73,620	47.15*	4,145	13.53*	74,700	34.92*	4,174	13.02*
	Increase relative to no targeting (%)	.05	.01	.44		.15		1.91		.84	
	t-statistic relative to no targeting	10.94*	5.35*	37.39*		16.81*		40.79*		16.45*	

$*p < .01$.
Notes: ns = not significant.

Table 7
HOW CUSTOMER BEHAVIOR ESTIMATES AFFECT TARGETING PROFITS

	Heinz Profits ($)	*Hunt's Profits ($)*
No firm targets (aggregate behavior)	66,628	4,116
No firm targets (true individual behavior)	73,301 (10.02%)	4,139 (.56%)
Both firms target with previous visit data	73,338 (10.07%)	4,139 (.56%)
Both firms target with previous purchase data	73,620 (10.49%)	4,145 (.70%)
Both firms target with full history data	74,700 (12.12%)	4,174 (1.41%)

cations based on consumer history, leading to different shares being estimated for the brands. Note that the profit increases from targeting in Table 7 are much higher than the figures we report in Table 6 and similar to the profit increases that Rossi, McCulloch, and Allenby (1996) report and to a few of Besanko, Dubé, and Gupta's (2003) estimates. Further research on targeting needs to account for this potential oversight when computing profits from targeting.

Profile of Consumers Targeted

Figure 4 shows the posterior segment probabilities of households targeted by Heinz (32 ounces) and Hunt's (32 ounces). For the sake of exposition, we label each segment with its most striking characteristic. Thus, we label Segment 1 as "Price Sensitive Heinz (32 Ounces) Loyals," Segment 2 as "Light Users, Heinz (28 ounces) Loyals," and Segment 3 as "Price Sensitive Heavy Users." In equilibrium, because Heinz targets almost entirely through coupons for its popular 32-ounce product, we profile only households receiving coupons for this product. Heinz (32 ounces) targets price sensitive households. Segment 3 (price sensitive heavy users) receives the most coupons from Heinz; households receiving a Heinz (32 ounces) coupon have a 67% probability of being in Segment 3 and a 32% probability of being in Segment 1. Thus, Heinz increases its profit margins from likely Segment 1 households (47% of market size) but competes aggressively with lower prices for likely Segment 3 households (29% of market size). Overall, Heinz offers targeted coupons to approximately 32% of households in the market.

Consumers targeted by Hunt's (32 ounces) are predominantly from Segment 3 (the most price sensitive segment of consumers who marginally favor the cheaper Hunt's brand). Given the lack of strong loyalty, Hunt's uses coupons to defend market shares in this segment. Hunt's (32 ounces) offers coupons relatively infrequently to households belonging to the other two segments. Overall, Hunt's offers targeted coupons to only approximately 9% of households in the market.

Identifying Sources of Targeting Profits

The increase in profits from 1:1 targeting arises from three sources: higher margins, higher brand shares, and consumption expansion. Ketchup consumption is unlikely to expand much as a result of couponing; indeed, category purchase expansion due to targeting is only .2%. We report the effect of targeting on each brand's shares and profit margins (in the full purchase history case) in Table 8. We calculated the average margins across all households by appropriately weighting the margins using household-level brand shares.

The gain in profits for Heinz (32 ounces) is essentially from price discrimination. Its average margins increase by approximately 2.8%, whereas brand share increases by

Figure 4
AVERAGE SEGMENT PROBABILITIES OF HOUSEHOLDS BY TARGETING STATUS

A: Heinz Targets/Does Not Target

B: Hunt's Targets/Does Not Target

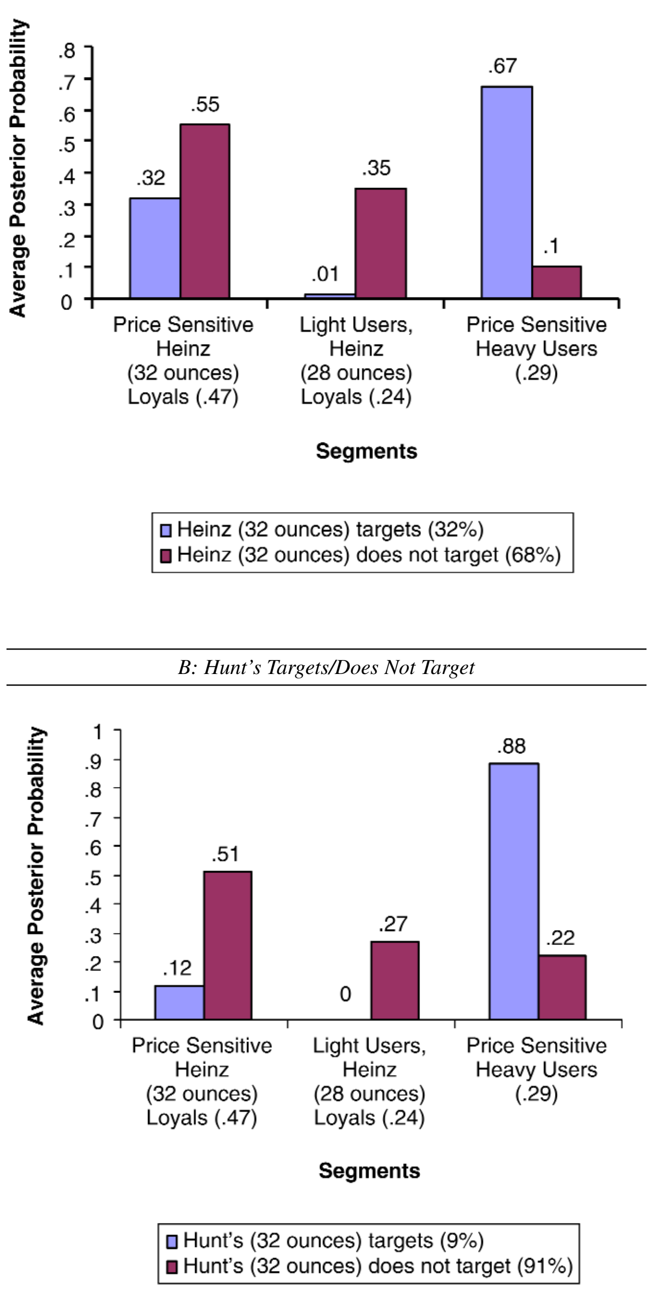

Table 8

EFFECT OF 1:1 COUPONS ON SHARES, MARGINS, AND CATEGORY PURCHASE

	Both Target Using Full History	
	Heinz	*Hunt's*
Average increase in share	32 ounces: +.3% 28 ounces: +.4%	32 ounces: +2.9%
Increase in (share weighted) margins	32 ounces: +2.8% 28 ounces: +.6%	32 ounces: +.7%
Average increase in category purchase	+.2%	

approximately .3%. In contrast, Hunt's share goes up by 2.9%, and margins increase by .7%. As the smaller brand, Hunt's takes advantage of the increase in prices by Heinz to increase its share (on its smaller base), even with a price increase. Thus, Heinz prices less aggressively than Hunt's, because as the larger brand, it can gain more from 1:1 pricing.

Targeting by Heinz (32 ounces) is more extensive; approximately 32% of consumers are targeted compared with only 9% of consumers targeted by Hunt's (32 ounces). The depth of discounts Heinz issues is also greater than that of Hunt's (32 ounces), but the aggregate prices of Heinz (32 ounces) increase because of this selective discounting.

Evaluating Strategic Options for the CDI

We next evaluate the optimal strategies for the CDI. Because the CDI always gains by selling to either Heinz or Hunt's, the price it can charge from a given client is the difference in profits of the client in the particular scenario being evaluated compared with the scenario when only one of the other clients receives the targeting service. For example, the price the CDI can charge from selling to Heinz (denoted as Firm 1) exclusively when selling the full purchase history is $P_1^f = \Pi_1^{10} - \Pi_1^{01} = 74{,}534 - 73{,}260 =$ \$1,274. Table 9 shows the price that Catalina will charge and its profits (assuming zero costs) in each of the targeting scenarios. The table shows that the profit for the CDI is greatest when both Heinz and Hunt's target using the full purchase history (\$1,475). Therefore, the firm will sell the targeting service to both firms ("whom to sell to?" or the optimal customer strategy) using the full purchase history of 100 weeks available ("what to sell?" or the optimal product strategy) at a price of \$1,440 to Heinz and \$35 to Hunt's ("for how much to sell?" or the optimal pricing strategy).

The results suggest that the total profits for the CDI using merely previous visit/previous purchase–based 1:1 targeting are small compared with the profits it obtains from using the full history. For example, with both firms targeting, the 1:1 vendor makes only \$21 in profits from previous visit–based targeting, and it makes \$1,475 from full visit history–based targeting. Another notable aspect of the results is that though most of the profits for the CDI come from Heinz, offering the service to Hunt's (even for free) can increase the price the CDI can obtain from Heinz. This is because of the positive externality for Heinz when Hunt's uses the service. Heinz profits increase by \$1,274 when it alone uses the service, but if Hunt's also uses the service, Heinz profits increase by \$1,440.

Thus, in this category, Catalina would maximize profits by selling its service on a nonexclusive basis to both vendors. It should reevaluate its current strategy of offering the service to only one firm. Furthermore, as we increase the length of purchase data even up to 100 weeks, the profitability of downstream clients continues to increase. Thus, restricting the data used for targeting to 65 weeks is suboptimal. Specifically, Catalina can improve its profits by increasing the data used from 65 weeks to 100 weeks by 16%. The main reason is that in infrequently purchased categories, such as ketchup, the information obtained from purchases over 65 weeks of data is not that large (the median number of purchases in 65 weeks is five). Catalina can improve its profitability by increasing the length of purchase history used in targeting. As data storage continues to become cheaper, this should be technologically feasible.

THE RETAILER'S PERSPECTIVE ON CUSTOMER DATA INTERMEDIATION

The retailer is the point of purchase when the consumer purchase data are collected, when customized coupons are printed and delivered, and when the coupons are redeemed. The retail loyalty card is most often the means of identifying the consumer, and the coupons are usually redeemable only in the same retail chain where purchases are made. Thus, a plausible scenario is one in which a retailer disintermediates the intermediary. Another issue of interest is what the value of targeting services is to a retailer that does not have the targeting infrastructure and whether a firm such as Catalina can benefit from providing the service to retailers. Therefore, we examine the roles of the retailer as both a 1:1 service provider and as a client of a 1:1 service provider.

The Retailer as a CDI

The retailer has two sources of increased profits as a CDI: (1) from the sales of the 1:1 service and (2) from the more efficient sales of goods (ketchup) at the retail store. Table 10 reports these two sources of profits. The profit

Table 9

PRICE AND CDI PROFITS UNDER ALTERNATIVE 1:1 MARKETING SCENARIOS

	Previous Visit–Based Targeting (\$)			*Full History–Based Targeting (\$)*		
	Price for Heinz	*Price for Hunt's*	*Total Profits*	*Price for Heinz*	*Price for Hunt's*	*Total Profits*
No firm targets	0	0	0	0	0	0
Hunt's only targets	0	0	0	0	15	15
Heinz only targets	0	0	0	1,274	0	1,274
Both firms target	21	0	21	1,440	35	1,475

Table 10
INCREMENTAL RETAILER PROFITS FROM CUSTOMER DATA INTERMEDIATION

	Full History–Based Targeting ($)		
	Profits from Ketchup Profits	*Profits from CDI*	*Total Retailer Profits*
No firm targets	0	0	0
Hunt's only targets	1,040	15	1,055
Heinz only targets	0	1,274	1,274
Both firms target	1,133	1,475	2,608

increase from both sources is greatest when both manufacturers target using full history. Retailer profits from sales of ketchup increase by $1,133, and profits from sales of 1:1 services increases by $1,475.[12]

Because the retailer profits from sales of ketchup also go up when manufacturers target, the retailer could forgo some proportion of its profits from the 1:1 services business to benefit from increase in ketchup profits due to targeting. This provides a compelling economic rationale for why retailers cooperate with Catalina, especially when they do not have the technological infrastructure for targeting. The analysis also implies that retailers could be formidable competitors to a company such as Catalina not only because such retailers may withdraw themselves from the CDI network (e.g., the Catalina Marketing Network) but also because they can price their targeting services more aggressively than a "pure" CDI. However, there are two main disadvantages for retailers in entering the CDI business. First, Catalina has patented several aspects of the personalization technology. Second, although a firm such as Catalina can provide manufacturers with one-stop shopping for 1:1 coupon services across the country, retailers can provide the 1:1 service only at their chain. Thus, manufacturers will need to negotiate for the targeting services with multiple retailers if the retailer provided the services. Indeed, in the retail industry, there is other evidence that manufacturers appreciate the benefits of one-stop shopping. For example, News America, a division of News Corp., currently owns the rights to contract out in-store feature and display advertising at approximately 35,000 food, drug, and mass merchandisers nationwide with revenues estimated at approximately $300 million (Neff 2006).

The Retailer as a Client of a CDI

If Catalina provides the 1:1 targeting service to retailers, retailer profits with full history data increase by approximately 1.88%, suggesting that targeting services to retailers can be an important source of revenue for Catalina. However, Catalina's revenues from retail targeting services is currently only approximately 9% of its revenues, compared with manufacturer targeting services, which account for 53% of its revenues.

It seems surprising that Catalina has not taken advantage of this potential revenue stream. A plausible reason Catalina has not aggressively marketed the service to retailers is that, given that retailers are the source of the data, it may be difficult for Catalina to extract the surplus created by targeting from retailers. Therefore, the prices Catalina can charge from retailers cannot be as high as those from manufacturers. It also raises broad questions about the property rights with respect to the data and how profits from the use of data should be shared. These issues require more detailed examination in future work.

CONCLUSION

The potential for CDIs has been growing as a result of advances in data collection and analysis technologies as well as advertising and promotion delivery technologies. In contrast to extant research on this topic, which has an "engineering" orientation, this article develops an empirical approach to answer strategic questions of interest to CDIs.

Our analysis enables us to obtain substantive insights into a CDI such as Catalina. First, as we discussed in the introduction, Catalina is currently reevaluating its policy of offering targeting services on an exclusive basis to manufacturers. Given the reservations expressed in the theoretical literature about the negative externalities induced by competitive targeting, Catalina must be careful in shifting from its extant policy of selling its targeted couponing services only on an exclusive basis. Our analysis shows that in the category we analyze, Catalina can increase its profits by selling to multiple manufacturers. By performing such an analysis on a category-by-category basis, Catalina can identify categories in which it can improve profits by changing its exclusive selling policy.

Second, we offer the insight that the retailer is likely to be a potent competitor to Catalina. Ginocchio, Chesler, and Clark (2005) suggest that a major threat to Catalina is the growing market share of Wal-Mart in groceries. Because Wal-Mart does not offer targeted coupons and is not part of Catalina's network, this can hamper Catalina's growth. According to the report, a second major threat is from Valassis Communications (currently in the business of offering coupons in free-standing inserts), which is considering entry into Catalina's targeted couponing business. However, the report suggests that Valassis will find it difficult to replicate Catalina's success, given Catalina's strong relationship with retailers.

Our analysis suggests that the major threat to Catalina may not be from Wal-Mart or Valassis but rather from large retailers themselves; retailers can effectively subsidize the price of the CDI service because this considerably increases retail profits due to 1:1 marketing. This threat should be salient given that many retailers (e.g., Tesco in the United Kingdom, Kroger in the United States) are developing their own technologies for offering 1:1 coupons to customers. In informal conversations, we learned that some supermarket chains offer targeted coupon services for free to all manufacturers, and a few currently charge for the service. Indeed, the retailer might be the most powerful potential competitor to Catalina in the future.

[12]Our result on retailer profits differs from the theoretical analysis of Liu and Zhang (2006), who find that retailer profits fall in the presence of personalized pricing. There could be several reasons for this: (1) They do not consider vertical differentiation in their model of consumer preferences; (2) they do not model the store brand, which the retailer owns directly; and (3) they assume consumer preferences are fixed, whereas empirically, we use a logit model in which consumer preferences are stochastic. Shin and Sudhir (2006) show analytically that vertical heterogeneity in quantities purchased (similar to the heavy-user/light-user segment in this article) and stochastic consumer preferences over time are required for firm profits to increase by using personalized pricing.

What would happen if there was competition among CDIs because either the retailer entered the market or the retailer supplied the data to multiple CDIs? In such a scenario, a syndication strategy by both 1:1 service providers would not be optimal because neither firm would make any profit given that they would sell homogeneous goods. An exclusive strategy according to which they sell to different downstream clients is likely to be optimal because it would create product differentiation.

Finally, we find that Catalina can benefit from increasing the length of purchase history it uses in its targeting services from the current self-imposed limit of 65 weeks. Even if storage costs are a reason for the current limit of length of purchase history used, the declining costs of storage and computing speeds should make it possible for Catalina to increase the length of history used for 1:1 marketing in the future profitably.

There are several ways this research can be extended. It would be useful to investigate the robustness of our results across multiple categories. We chose the ketchup category to compare our results with those of Besanko, Dubé, and Gupta (2003). Although the gains from targeting in the ketchup category are low, a Catalina executive told us that the firm expects substantially higher gains in categories such as snack foods, in which there is also potential for category expansion in profits due to targeting. It is possible that alternative channel structures exist in other product categories and other retail chains, and it would be useful to study CDIs in product categories or retail chains in which channel structures, such as the retailer Stackelberg or the Vertical Nash, are more appropriate.

Further research should expand on the nature of personalization data used for targeting. Because we did not find demographics to be useful, we treat purchase history length as the proxy for data quality. As we discussed, quality may be increased through greater "breadth" of the data by integrating purchase behavior from other categories. Although estimation and optimization is computationally much more difficult across categories, we believe that this is an important area for future work. In general, this approach can identify the potential profitability of cross-selling services.

Further research should also investigate the impact of greater flexibility in couponing and pricing approaches offered by the 1:1 service provider. In this article, we follow previous research (Besanko, Dubé, and Gupta 2003; Rossi, McCulloch, and Allenby 1996) in assuming that coupons are valid only for the week of issue. In practice, however, coupons are valid for multiple weeks, and consumers can time when they use the coupon; thus, firms should account for such dynamic behavior in issuing the coupon. Therefore, modeling timing of coupon redemption requires empirically modeling dynamics and forward-looking behavior on both the consumer side (e.g., Gonul and Srinivasan 1996) and the firm side (e.g., Gonul and Shi 1998). In addition, we model the upper threshold of the fee that Catalina can charge manufacturers, given our objective of identifying the best strategy for a 1:1 service provider. This is also consistent with the two-part pricing that Catalina currently uses, in which the fixed price for the service varies considerably across categories. Nevertheless, a systematic investigation of alternative pricing schemes based on the number and value of coupons issued or redeemed would be a worthwhile area for further research.

As in previous research, we assume that retailers use the same level of information as that purchased by the manufacturer from the CDI. We make this assumption because of the need for "common knowledge" between the manufacturer and the retailer to solve the manufacturer–retailer pricing game. The implications of relaxing this common knowledge assumption need to be studied in further research, but this requires methodological advances in the game theory literature.

Finally, we hope that that our approach will inspire additional research to facilitate decision making in other CDI contexts, such as that of durable goods markets, financial services, catalog marketing, and targeted advertising. In the contexts of durable goods or financial services, there will be shorter purchase histories but greater information across categories that can be used for 1:1 marketing. In the context of targeted advertising services, the empirical model needs to calibrate the impact of advertising (rather than couponing) on consumer purchasing decisions. The analytical approach we developed can also be applied to domains other than 1:1 marketing. For example, the optimal licensing/selling of a patented innovation (e.g., Katz and Shapiro 1985) could be analyzed. Licensing of an innovation to multiple downstream firms may create greater competition downstream and thus reduce the total value of the innovation compared with the strategy of exclusively selling the patent to one firm. In short, although appropriate changes are needed for the model to accommodate institutional details appropriate for each context, the general framework of understanding the trade-offs involved in improving quality and selling to exclusive/multiple clients will continue to be relevant. More broadly, we hope that this approach spawns similar complementary research to game-theoretic analysis on other marketing institutions to help decision makers and managers obtain empirically driven answers to their business strategy questions.

REFERENCES

Adomavicius, Gediminas, Ramesh Sankaranarayanan, Shahana Sen, and Alexander Tuzhilin (2005), "Incorporating Contextual Information in Recommender Systems Using a Multidimensional Approach," *ACM Transactions on Information Systems*, 23 (1), 103–144.

Ainslie, Andrew and Peter E. Rossi (1998) "Similarities in Choice Behavior Across Product Categories," *Marketing Science*, 17 (2), 91–106.

Allenby, Greg M., Neeraj Arora, and James L. Ginter (1998), "On the Heterogeneity of Demand," *Journal of Marketing Research*, 35 (August), 384–89.

Andrews, Rick L., Andrew Ainslie, and Imran S. Currim (2002), "An Empirical Comparison of Logit Choice Models with Discrete Versus Continuous Representations of Heterogeneity," *Journal of Marketing Research*, 39 (November), 479–87.

Ansari, Asim M. and Carl Mela (2003), "E-Customization," *Journal of Marketing Research*, 40 (May), 131–46.

Berto Villas-Boas, Sofia (2007), "Vertical Contracts Between Manufacturers and Retailers: Inferences with Limited Data," *Review of Economic Studies*, 74 (2), 625–52.

Besanko, David, Jean-Pierre Dubé, and Sachin Gupta (2003), "Competitive Price Discrimination Strategies in a Vertical Channel Using Aggregate Retail Data," *Management Science*, 49 (September), 1121–37.

———, Sachin Gupta, and Deepak Jain (1998), “Logit Demand Estimation Under Competitive Pricing Behavior: An Equilibrium Framework,” *Management Science*, 44 (11), 1533–47.

Chen, Yongmin (1997), “Paying Customers to Switch,” *Journal of Economics and Management Strategy*, 6 (4), 877–97.

Chen, Yuxin and Ganesh Iyer (2002), “Consumer Addressibility and Customized Pricing,” *Marketing Science*, 21 (2), 197–208.

———, Chakravarthi Narasimhan, and Z. John Zhang (2001), “Individual Marketing with Imperfect Targetability,” *Marketing Science*, 20 (3), 23–43.

——— and Z. John Zhang (2004), “The Benefit of Behavior-Based Targeted Pricing with Strategic Consumers,” working paper, Marketing Department, New York University.

Chintagunta, Pradeep K. (2002), “Investigating Category Pricing Behavior at a Retail Chain,” *Journal of Marketing Research*, 39 (May), 141–55.

Fudenberg, Drew and Jean Tirole (2000), “Customer Poaching and Brand Switching,” *RAND Journal of Economics*, 31 (4), 634–57.

Ginocchio, Paul, Matthew Chesler, and David T. Clark (2005), “Catalina Marketing: No Coupons at Wal-Mart,” Deutsche Bank Global Equity Research Report, (March 2005).

Gonul, Fusun and Meng Ze Shi (1998), “Optimal Mailing of Catalogs: A New Methodology Using Estimable Structural Dynamic Programming Models,” *Management Science*, 44 (9), 1249–62.

——— and Kannan Srinivasan (1996), “Estimating the Impact of Consumer Expectations of Coupons on Purchase Behavior: A Dynamic Structural Model,” *Marketing Science*, 15 (3), 262–79.

Gupta, Sunil (1988), “Impact of Sales Promotions on When, What, and How Much to Buy,” *Journal of Marketing Research*, 25 (November), 342–55.

Hausman, Jerry A. (1978), “Specification Tests in Econometrics,” *Econometrica*, 46 (November), 1251–71.

Humby, Clive (2004), “The Future of Communications: Is Consumer Marketing Irrevocably Broken?” *Admap*, (October), 2–4.

———, Terry Hunt, and Tim Phillips (2003), *Scoring Points: How Tesco Is Winning Customer Loyalty*. London: Kogan Page.

Iyengar, Raghuram, Asim Ansari, and Sunil Gupta (2003), “Leveraging Information Across Categories,” *Quantitative Marketing and Economics*, 1 (4), 425–65.

Iyer, Ganesh and David Soberman (2000), “Markets for Product Modification Information,” *Marketing Science*, 19 (3), 203–225.

———, J. Miguel Villas-Boas, and David Soberman (2005), “The Targeting of Advertising,” *Marketing Science*, 24 (3), 461–76.

Kamakura, Wagner A. and Gary J. Russell (1989), “A Probabilistic Choice Model for Market Segmentation and Elasticity Structure,” *Journal of Marketing Research*, 26 (November), 379–90.

Katz, Michael L. and Carl Shapiro (1985), “On the Licensing of Innovations,” *RAND Journal of Economics*, 16 (4), 504–520.

Liu, Duen-Ren and Ya-Yueh Shih (2005), “Hybrid Approaches to Product Recommendation Based on Customer Lifetime Value and Purchase Preferences,” *Journal of Systems and Software*, 77 (2), 181–91.

Liu, Yunchuan and Z. John Zhang (2006), “The Benefits of Personalized Pricing in a Channel,” *Marketing Science*, 25 (1), 97–105.

Miller, Paul (2003), “Lists and Prospecting: Who’s Who Among the Consumer Co-Op Databases,” *Multichannel Merchant*, (December 1), [available at http://multichannelmerchant.com/catalogage/database/marketing_cheat_sheet_whos/].

Neff, Jack (2006), “Murdoch Plan Locks Up the Store,” *Advertising Age*, (January 30), 1–2.

Petrin, Amil and Kenneth Train (2003), “Omitted Product Attributes in Discrete Choice Models,” National Bureau of Economic Research Working Paper No.W9452, Cambridge, MA.

Rivers, Douglas and Quang H. Vuong (1988), “Limited Information Estimators and Exogeneity Tests for Simultaneous Probit Models,” *Journal of Econometrics*, 39 (3), 347–66.

Rossi, Peter E., Robert McCulloch, and Greg M. Allenby (1996), “The Value of Purchase History Data in Target Marketing,” *Marketing Science*, 15 (4), 321–40.

Seetharaman, P.B. (2004), “Modeling Multiple Sources of State Dependence in Random Utility Models: A Distributed Lag Approach,” *Marketing Science*, 23 (2), 263–71.

Shaffer, Greg and Z. John Zhang (1995), “Competitive Coupon Targeting,” *Marketing Science*, 14 (4), 395–417.

——— and ——— (2002), “Competitive One-to-One Promotions,” *Management Science*, 48 (9), 1143–60.

Shin, Jiwoong and K. Sudhir (2006), “Competitive CRM: Should We Punish or Reward Our Current Customers?” working paper, Yale School of Management.

Silva-Risso, Jorge M., Randolph E. Bucklin, and Donald G. Morrison (1999), “A Decision Support System for Planning Manufacturers’ Sales Promotion Calendars,” *Marketing Science*, 18 (3), 274–300.

Sudhir, K. (2001), “Structural Analysis of Manufacturer Pricing in the Presence of a Strategic Retailer,” *Marketing Science*, 20 (3), 244–64.

Taylor, Curtis R. (2003), “Supplier Surfing: Price-Discrimination in Markets with Repeat Purchases,” *RAND Journal of Economics*, 34 (1), 223–46.

Thisse, Jacques-Francois and Xavier Vives (1988), “On the Strategic Choice of Spatial Price Policy,” *American Economic Review*, 78 (1), 122–37.

Villas-Boas, J. Miguel (1999), “Dynamic Competition with Customer Recognition,” *RAND Journal of Economics*, 30 (4), 604–631.

——— (2004), “Price Cycles in Markets with Customer Recognition,” *RAND Journal of Economics*, 35 (3), 486–501.

——— and R. Winer (1999), “Endogeneity in Brand Choice Models,” *Management Science*, 45 (10), 1324–38.

——— and Ying Zhao (2005), “Retailer, Manufacturers, and Individual Consumers: Modeling the Supply Side in the Ketchup Marketplace,” *Journal of Marketing Research*, 42 (February), 83–95.

Wedel, Michel and Wagner A. Kamakura (2000), *Market Segmentation: Conceptual and Methodological Foundations*, 2d ed. Boston: Kluwer Academic Publishers.

Zhang, Jie and Lakshman Krishnamurthi (2004), “Customizing Promotions in Online Stores,” *Marketing Science*, 23 (4), 561–78.