

Should Price Increases Be Targeted?—Pricing Power and Selective vs. Across-the-Board Price Increases

Aradhna Krishna, Fred M. Feinberg

Stephen M. Ross School of Business, University of Michigan, Ann Arbor, Michigan 48109
{aradhna@umich.edu, feinf@umich.edu}

Z. John Zhang

The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104,
zjzhang@wharton.upenn.edu

Firms in many industries experience protracted periods of *pricing power*, the ability to successfully enact price increases. In these situations, firms must decide not only whether to raise prices, but to whom. Specifically, in a competitive context, they must determine whether it is more profitable to increase prices across-the-board or to a specific segment of their customer base. While selective price decreases are ubiquitous in practice (e.g., better deals to potential new customers by phone carriers; better deals to current customers by various magazines), to our knowledge selective price increases are relatively rare.

We illustrate the benefits of targeted price increases, and, as such, we expand the repertoire of firms' promotional policies. To that end, we explore a scenario where two competing firms must decide whether to increase prices to the entire market or only to a specific segment. Targeted price increases (TPI), i.e., being offered an unchanged price (selectively) when others are subject to price increases, can be offered to Loyals (those who bought from the firm in the previous period) or Switchers (those who did not). The effects of TPIs are estimated through a laboratory experiment and an associated stochastic model, each allowing for both rational (Loyalty, Switching) and behaviorist (Betrayal, Jealousy) effects.

We find that TPIs can indeed yield beneficial results (greater retention for Loyals or greater attraction of Switchers) and greater profits in certain circumstances. Results for TPI are additionally benchmarked against those for targeted price decreases and are found to differ. The range of effects stemming from the experiment can be used in a competitive analysis to yield equilibrium strategies for the two firms. In this case, we find that—depending on the magnitude of the price increase, market shares of the two firms, and price knowledge across consumer segments—a firm may wish to embrace targeted price increases in some situations, to institute across-the-board price increases in others, and to not enact any price increases in still others. We show that a firm can sacrifice considerable profit if it settles on a suboptimal pricing strategy (e.g., wrongly instituting an across-the-board increase), favors the wrong segment (e.g., Switchers instead of Loyals), or ignores “behaviorist” effects (Betrayal or Jealousy).

Key words: game theory; stochastic models; targeted pricing; pricing power; promotions

History: Accepted by Wallace J. Hopp, marketing; received May 3, 2004. This paper was with the authors 1 year and 3 months for 2 revisions. Published online in *Articles in Advance* July 20, 2007.

1. Introduction: Pricing Power and Targeted Price Increases

Firms in many industries experience protracted periods of *pricing power*, the ability to successfully enact price increases. While pricing power has received little direct attention from academic researchers, the business press has consistently emphasized its importance. *Industry Week*, *The Wall Street Journal*, and *Fortune* have all reported on upswings in pricing power, while *The New York Times* placed it among the main qualities sought by savvy investors.¹ When firms have

substantial pricing power, they must grapple with the prospect of instituting price increases and whether such increases should be targeted in some way. In this article, we investigate whether, and under what competitive conditions, firms should consider imposing price increases on some segment as opposed to across-the-board.

Even within popular press accounts of targeted price increases (TPI), opinions differ on their suit-

¹ See, for example, *Industry Distribution*, February 1, 2006, Vol. 95, Issue 2 (“2006: The Year of Pricing Power?”); *The Wall Street Journal*,

August 2005, p. B3A (“Pricing Power Lifts Some Sectors, but Others Cannot Pass on Costs”); *Industry Week*, April 2004, p. 71 (“Pricing Power Returns”); *The New York Times*, April 3, 2004 (“Investors Eye Firms with ‘Pricing Power’”); *Fortune*, September 15, 2003, p. 52 (“Pricing Power Ain’t What It Used To Be”).

ability and sustainability. For example, a study of British television mergers concluded that targeted price increases, defined explicitly as “giving less good terms to some customers and better terms to others,” were not feasible because any incremental revenue would be “outweighed by the disincentive in terms of potential loss of business” (http://www.competition-commission.org.uk/rep_pub/reports/2003/fulltext/482c6.pdf). Even the Federal Trade Commission (FTC), concerned about abuse of pricing power in certain industries to enact targeted price increases, investigated the phenomenon; yet it reached precisely the opposite conclusion. As part of a major investigation of vacation package pricing policies, the FTC’s analysis indicated it is indeed possible for firms to benefit by imposing a TPI policy. In fact, it singled out targeted price increases among important areas for additional empirical research, calling for “detailed industry data.” The Boston Consulting Group evidently agrees that TPI can be a useful tactic, claiming a 1% price increase to specific segments can boost incremental profits fourfold over comparable cuts in overhead and fixed costs. Operations Management Consultants Inc.’s “top recommendation” for increasing profitability is that firms learn to employ TPI policies.

Based on such anecdotal and economic evidence, one might wonder why TPI policies are not more broadly adopted. Amazon.com offers the classic cautionary tale, having faced allegations that it imposed higher prices on some customers based on their purchase and site visit histories; when consumer and advocacy groups complained vociferously, the practice was abandoned. Garbarino and Lee (2003) point out that while TPI-like policies can intrinsically enhance firm profitability, consumers nevertheless react strongly against them. Their experimental results show such policies erode consumer trust, largely because they are considered unfair. Some firms apparently know this and keep their targeting pricing policies under wraps. For example, a federal lawsuit against Rite-Aid revealed its practice of imposing unadvertised surcharges on certain customers, including those who had not shopped at the store before. Apparently the company believed this practice would be viewed negatively, even though it was (loyal) continuing customers who, relatively speaking, benefited from such price increases.

As these varied examples suggest, there is little consensus on whether TPI is a reasonable or sustainable practice. Compounding such questions about the practice itself are contingent ones regarding which consumers, if any, should be thus targeted. At the heart of the matter is whether consumers will react in a “rational” way to having their prices raised

(or seeing it happen to others), with rationality narrowly construed as making the best economic choice amongst available options. Here, we take up this question, providing a model that addresses whether (if at all) TPI should be instituted and, further, which segment(s) should be thus favored. Our model comprises not only the “rational” effects of “Loyalty” (to one’s current provider) and “Switching” (to another), but also “behaviorist” effects born of emotional reactions to how one is treated: “Betrayal” when one’s own firm treats other customers better and “Jealousy” when another firm treats its own customers better than your firm treats you (Feinberg et al. 2002).

We study environments in which pricing power allows firms to potentially enact price increases. As such, a major goal of this paper is to expand the repertoire of promotional policies available to firms when pricing power can be profitably exercised. Specifically, we attempt to answer the following questions:

- Under TPI, will consumers perceive “same price as before” as a price deal to them, just because a price to the other segment is raised, i.e., will there be a positive impact of imposing a price increase selectively on the other segment?
- Are TPI effects similar to those of targeted price decrease (TPD) policies?
- In an environment of increasing prices, if a firm offers “some consumers (and not others) the same price as before” versus offers “all consumers get the same price as before,” would the response of the favored consumers be the same in the two situations?
- How does TPI affect firm profits, that is, how do the various effects stemming from TPI act in concert to alter profit levels?
- What can be said about competitive equilibria? Specifically, when they propose price increases, should firms offer deals to Loyals, to Switchers, or to neither segment?
- Will equilibria differ if “behaviorist” (Betrayal, Jealousy) effects are not accounted for?
- How does the degree of market share asymmetry affect a firm’s decision to impose selective or across-the-board increases and, if the former, which segment should be favored?
- How does level of market knowledge about prices offered to the other segment affect resulting equilibria?

Procedurally, our approach integrates multiple methodologies to better understand the impact of targeted pricing policies. We show that although the “behaviorist” effects of Jealousy and Betrayal exist in both settings, they cannot be presumed symmetric and such differences can affect a firm’s strategic choices. Our analysis further demonstrates that across-the-board price increases are not always the best way to exercise a firm’s pricing power. Indeed,

price-increase environments can offer opportunities to profitably implement targeted pricing, and we explore the tactical imperatives of a firm's being able to avail itself of them. We also find that competitive equilibria vary substantially with the proportion of "aware and care" consumers, those who find out about price deals offered to others and are potentially influenced by them.

The remainder of the paper is organized as follows. We first examine prior literature from marketing, economics, psychology, and behavioral decision theory which speaks to the issue of targeted price increases. This allows us to develop a Markovian framework to examine the purchase and profit impacts of various targeted pricing policies, and we estimate the model's parameters based on data collected in a laboratory experiment. These estimated parameter values are then used within a game theory model built upon the Markovian framework; together, these allow for an analysis of competitive equilibria for a two-firm market. It is this competitive analysis which offers specific responses to each of the questions raised above. We conclude with a discussion of how the model applies to real-world promotional policies and how it could be further enhanced relative to current retailing practice.

2. Prior Literature

Much prior research on targeted pricing has focused on whether it tends to increase or decrease price competition. Specifically, it has examined welfare implications of customized pricing and whether changes in consumer behavior induced by targeted pricing lead to socially optimal equilibria (Thisse and Vives 1988, Shaffer and Zhang 1995). Markets with high consumer switching costs are especially concerned with implications of targeted pricing: Firms in these markets must trade off the relative merits of charging a high price to everyone (greater unit profits with some potential attrition among the customer base) against charging a low price to everyone and thereby attracting new customers (Klemperer 1987, 1995). Targeted promotions provide a means to bypass such a trade-off because they facilitate charging different prices to these two segments of consumers. Subsequent research has suggested that when following a policy of targeted pricing, a firm should invariably target Switchers (Chen 1997, Taylor 2003). While Shaffer and Zhang (2000) suggest that targeting Switchers need not be optimal, their results indicate that it is never prudent for all competing firms to target loyal customers unless the cost of targeting prompts competing firms to resort to randomized promotions (Shaffer and Zhang 2002).

All these papers consider an environment where prices are stable and there are no imperatives to increase price. They also assume that a consumer's

brand preference is dependent only on absolute prices offered to him and not on prices offered to other consumers; in short, that relative prices have negligible effects on consumer behavior. However, literature from social psychology and behavioral decision theory suggests that "relative" prices (i.e., price to you versus price to someone else) may also affect behavior. More specifically, research on relative deprivation, perceived fairness, and equity (e.g., Adams 1965, Greenberg 1986, Stark and Taylor 1989) suggests that perceptions of fairness in a broad spectrum of economic interactions do not take place in a vacuum of perfect objectivity. Rather, they are assessed relative to standard comparison values derived, in large part, from how others are treated.

Consumers' perceptions of fairness, specifically, have been addressed by research in marketing and psychology. For example, even if objective price levels are within the range considered normal, if they are not deemed equitable, then consumers may consider them "unfair" and refuse to accept them (Kahneman et al. 1986a, b; Martins and Monroe 1994; Urbany et al. 1989; Campbell 1999). Bolton et al. (2003) study price unfairness directly, finding that consumers often view prices as lying well above what is "fair" and are overly sensitive to such referents as prior and competitor price levels. Nunes and Park (2003) present a great deal of evidence that perceptions of price levels are strongly frame dependent. In line with these findings, Feinberg et al. (2002) formulate a "behaviorist" approach to targeted promotions which allows for consumers being influenced by prices that other consumers are offered and of which they themselves cannot avail. They show that behaviorist effects of Betrayal and Jealousy result in decidedly different promotion policies than a "rationalist" model which does not allow for such effects. Their paper examines environments of constant prices where firms consider to whom they should offer price decreases. As such, their analysis does not apply when a firm considers price hikes and can allow some (segment of) customers to remain at their current price level.

In short, based on these prior analyses, nothing can be said about the profit implications of targeted price increases or whether they can be more profitable than across-the-board increases. In the next section, we build a Markov model for consumer promotional response to address these issues.

3. Model

Similar to many prior studies in marketing and economics, we analyze a market consisting of two brands, X and Y (i.e., sold by Firms X and Y, respectively). We adopt a (first-order) Markovian framework and operationalize market segments based on consumers' most recent purchase. Thus, "Loyals" are those who purchased from the firm in the last period,

while “Switchers” purchased from the other firm in the last period. In this manner, there are two market segments for each firm, Loyals for Firm X (who are Switchers for Firm Y) and Switchers for Firm X (who are Loyals for Firm Y). Note that this definition of Switchers differs from another fairly common use of “Switcher” (e.g., Lal 1990), where a segment of consumers always buys from Firm X (Loyals for X), another always buys from Firm Y (Loyals for Y), and a third segment switches between the two firms based on price (Switchers). Also, in our model, there is no absolute Loyalty—all consumers can potentially switch. These terms thus act as labels for the immediately prior purchase and do not refer to an intrinsic propensity to switch between the two brands.

So that terminology is not open to interpretation, we adhere to the following conventions. Phrases such as “charges a higher price” mean that a firm charges a price compared to its rival, not compared to a base or past price for that same firm. Similarly, “less likely to purchase” compares likelihoods for buying from a specific firm when a condition holds versus when it does not. Thus, “consumers are less likely to buy from their firm if it charges a higher price to them than to Switchers” means that consumers’ probability of buying from their firm is less when the firm charges a higher price than it does to (potential) Switchers versus when it does not.

Each firm (X or Y) has four options for offering a *price special* in the current period: to No One (price increase to both segments, i.e., an across-the-board increase), only to Switchers (price increase to Loyals only), only to Loyals (price increase to Switchers only), or to both Loyals and Switchers (i.e., no price increase to either segment). Thus, the two firms generate a strategy space with 16 possible pricing scenarios. To simplify references to these 16 possible pairwise pricing scenarios, we use the symbols N , S , L , and A to stand for possible actions by each of the firms, so that $\{S, N\}$, for example, means that Firm X favors Switchers (that is, Firm X imposes no price increase for Y’s customers) while Firm Y offers no price breaks at all (i.e., Firm Y implements an across-the-board price increase). Similarly, $\{L, A\}$, for example, means that Firm X favors Loyals (that is, Firm X imposes no price increase for its own customers) while Firm Y offers price breaks to both its Loyals and its potential Switchers (i.e., Firm Y implements an across-the-board price special—it does not implement the price increase).

3.1. Consumer Choice and Price Specials

When there are no price specials of any sort, we allow for consumers being first-order purchasers. For consecutive purchase occasions (periods $t - 1$ and t),

$P[X_t | X_{t-1}] = \alpha$ and $P[Y_t | Y_{t-1}] = \beta$; that is, the transition matrix is

$$\begin{array}{c} X_t \quad Y_t \\ \begin{array}{|c|c|} \hline X_{t-1} & \begin{array}{c} \alpha \quad 1 - \alpha \end{array} \\ \hline Y_{t-1} & \begin{array}{c} 1 - \beta \quad \beta \end{array} \\ \hline \end{array} \end{array}$$

The parameters $\{\alpha, \beta\}$, which we take to be stationary, reflect intrinsic preferences for the brands in terms of repurchase probabilities. If the choice process is zero order, $\beta = 1 - \alpha$; if not, i.e., if variety-seeking (Givon 1984) or inertial (Jeuland 1979) tendencies exist, then $\beta \neq 1 - \alpha$. It is important to note that this specification accounts for consumers’ switching costs: in the absence of promotion, reluctance to change from one to the other will be reflected, all else equal, in higher values of α and β .

We account for promotion-induced shifts away from these baseline preference levels through four basic effects: *Switching* (s), *Loyalty* (l), *Betrayal* (b), and *Jealousy* (j) (Feinberg et al. 2002), as discussed earlier. Switching occurs when a rival firm offers some inducement to its competitors’ customers, while *Loyalty* involves doing the same to its own customers. Both are well-known to modelers; *Betrayal* and *Jealousy*, however, are considerably less familiar. *Betrayal* occurs when a firm treats its own customers less well than it treats some other group, while *Jealousy* occurs when customers perceive that they would be treated better if they were loyal to a firm other than their own. In addition to these four basic effects (with their four accompanying parameters), we have two additional parameters, l' and s' , that capture the Loyalty and Switching effects when firms offer price specials (no price increase) to both segments— l' when your own firm has no price increase to either its Loyals or its potential Switchers, and s' when the other firm has no price increase to either its Loyals or Switchers. The parameters l' and s' are not presumed identical to l and s , so that promotional effects can differ across the *targeted* and *indiscriminate*² price special cases.

When multiple deals are being offered, effects can combine (as specified in the forthcoming model and detailed in Table 1, which specifies both which segment is favored and whose prices are increased).³ The parameters are meant to account for effects over and above baseline levels (α and β) in the sense that if

² We use the term “indiscriminate” when firms do not discriminate, that is, they offer the price special to both segments.

³ Heterogeneity in α and β is accounted for in our empirical analysis using fixed effects, although other specifications are possible. Because we found no evidence of heterogeneity in any of the “effects parameters” $\{s, l, b, j\}$ across brands, we omit brand subscripts on all model parameters.

Table 1 Conceptual Model of Effects

Your firm favors [Price increases to]	Other firm favors			
	No One, <i>N</i> [All]	You, <i>S</i> [Others]	Others, <i>L</i> [You]	All, <i>A</i> [No One]
No One, <i>N</i> [All]	(Baseline)	Switching, <i>s</i>	Jealousy, <i>j</i>	Switching, <i>s'</i>
You, <i>L</i> [Others]	Loyalty, <i>l</i>	<i>l</i> and <i>s</i>	<i>l</i> and <i>j</i>	<i>l</i> and <i>s'</i>
Others, <i>S</i> [You]	Betrayal, <i>b</i>	<i>b</i> and <i>s</i>	<i>b</i> and <i>j</i>	<i>b</i> and <i>s'</i>
All, <i>A</i> [No One]	Loyalty, <i>l'</i>	<i>l'</i> and <i>s</i>	<i>l'</i> and <i>j</i>	<i>l'</i> and <i>s'</i>

consumers have higher Switching costs (reflected in higher α and β), then the same promotion would have a smaller effect on choice.

We encode the four effects in a manner consistent with prior literature (Kahn and Raju 1991), specifically, that any effect operates as a *proportion* of the relevant affected set of the market. An example helps clarify the principle. Take the case where Firm X entices the consumers of Firm Y by offering them an inducement: For the consumers of Firm Y, this will induce Switching effects, whereas for the consumers of Firm X, it will induce Betrayal effects. So, the *repurchase* probabilities change from their “prior” values of α and β to $\alpha(1 - b)$ and $\beta(1 - s)$. Hence, one can simply replace the labels in Table 1 with corresponding formulas to obtain the entries of Table 2. Note that the expressions in Table 2 are taken from the perspective of a customer of Firm X; a corresponding table exists for Firm Y’s customers.⁴

In the course of formulating Table 1, a wide variety of alternate specifications suggested themselves. Many were compelling in some way, like parsimony, but less so in others, like eventual empirical fit. It is instructive to examine some of these other specifications. For example, in the (All, No One) cell, one could specify *l* and *b*, which would reflect a mixture of the two cells directly above it. Although this would have a certain symmetry and elegance, it is less reasonable on its face from a psychological perspective: If Loyals receive a special deal, then giving that same deal to Switchers would not entail betrayal, a highly specific effect arising when one is treated in a seemingly unfair, detrimental manner. Removing the Betrayal effects would leave only the *l* effect for this cell, forcing there to be no modeled differences between the rows for “All” and “You.” While this would be parsimonious, it would make little sense. The posited effect *l'* allows for differences between each row (and,

in fact, cell, as no two are formulaically identical); although it could be supplemented by specifying “*l* and *b'*,” this not only lacks parsimony (expanding the range of “behaviorist” effects), but fails to yield a superior fit for the data in our forthcoming application. It also seemed important not to presume that the (All, All) cell reverted to “baseline” simply because all groups were treated equally, as in the (No One, No One) cell. Similar logic held throughout Table 1, particularly so for the (No One, All) cell, where the parameter *s'* appears as the sole effect. Of course, any particular specification entails caveats and trade-offs; Table 1 represents insights in the simplest manner, while capturing empirical features of our data as well as any alternate specification.

Note that in Table 2, if all hypothesized effects parameters are set to zero, we are left with the standard inertial/variety-seeking specification, $P[X_t | X_{t-1}] = \alpha$, $P[Y_t | Y_{t-1}] = \beta$. In some cells with two effects—for example, (You, You) and (You, Others)—one effect impacts your (i.e., the focal firm’s) consumers (those who bought from you in the last period) and, hence, your base repurchase probability, whereas the other effect impacts the other firm’s customers (those who bought from the other firm in the last period) and, hence, the other firm’s repurchase probability. In such cases, the two effects can be jointly accounted for by simple addition, without jeopardizing logical consistency (Naert and Bultez 1973). However, in other cells with two effects, for example, (Others, You) and (Others, Others), both effects impact your customers (those who bought from you in the last period) and, hence, your base repurchase probability. Taking the example of the (Others, You) cell, a purely additive model would suggest a repurchase probability of $\alpha(1 - b - s)$; this, in turn, requires an exogenous constraint that $b + s \leq 1$, a constraint that would need to hold throughout Table 2 even in cells with single effects. It is difficult to argue why this should be the case, and one can imagine scenarios in which the parameters can be independently manipulated. We thus adopted a convention assuring both logical consistency and lack of such constraints, whereby effects enter multiplicatively, for example, $\alpha(1 - b)(1 - s)$. (The chosen multiplicative formulation reflects consumer behavior well: We also tested the purely additive model and found that it provides a slightly worse fit to the data presented in our application. Model estimation results for this linear specification are available in the online appendix, which is provided in the e-companion.⁵) We can therefore estimate the parameters over an unconstrained space, allowing

⁴ The terminology and notation of Table 2 are, for consistency with prior literature and later tests, oriented toward which segment is favored. Note that the row and column labels can be readily translated into which group has a price increase imposed on them, as in Table 1: “No One” being favored corresponds to a price increase being imposed on *everyone*; similarly, favoring “You” means raising price to others, favoring “Others” means raising price to you, and favoring “All” means not raising price at all.

⁵ An electronic companion to this paper is available as part of the online version that can be found at <http://mansci.journal.informs.org/>.

Table 2 Specified Model

Your firm favors	Other firm favors			
	No One, N	You, S	Others, L	All, A
No One, N	α	$\alpha(1-s)$	$\alpha(1-j)$	$\alpha(1-s')$
You, L	$\alpha + l(1-\alpha)$	$\alpha(1-s) + (1-\alpha)l$	$\alpha(1-j) + (1-\alpha)l$	$\alpha(1-s') + (1-\alpha)l$
Others, S	$\alpha(1-b)$	$\alpha(1-b)(1-s)$	$\alpha(1-b)(1-j)$	$\alpha(1-b)(1-s')$
All, A	$\alpha + l'(1-\alpha)$	$\alpha(1-s) + (1-\alpha)l'$	$\alpha(1-j) + (1-\alpha)l'$	$\alpha(1-s') + (1-\alpha)l'$

standard errors their usual interpretations. Note that in some cells, such as those in which both firms target exactly the same individuals, it would appear reasonable to allow all effects to simply cancel, yielding “baseline.” While this may be more parsimonious, the pattern of effects achievable under Table 2 is intrinsically richer, and moreover, we found all simpler models (i.e., parametric restrictions) to offer demonstrably inferior empirical results.

We next present an experiment designed to allow estimation of these various effects, $\{s, l, b, j, s', l'\}$.

4. Experiment

4.1. Purpose

The parameters of the model specified in Table 2 can be estimated using field data. Although conducting such tests in a lab setting raises issues of external validity, natural field experiments where a pair of firms in the same product class have adopted multiple targeted policies are at best exceedingly rare, let alone multiple concurrent targeted price-increase policies. The laboratory experiment allows us to do this while attempting to replicate market conditions and inducements to the best extent possible. One cautionary note exists, however, laboratory experiments cannot evoke the kind of negative emotions that one would expect in the real world when, for example, loyal customers of a firm are “betrayed” and learn this from a third-hand source. As such, our results could err on the side of underestimating behaviorist effects and could be viewed in this light. However, one could also argue that our results may be biased upward, since subjects in the experiments may focus more on price that they would in the real world.

4.2. Design

Consistent with our model, the experiment has a 4-by-4, between-subjects design. That is, each of the two firms can choose to offer a price deal to Switchers, to Loyals, to Neither, or to All (both segments). Specifically, the price deal is implemented through a price increase to Loyals (i.e., deal to Switchers who have a constant price), price increase to Switchers (deal to Loyals), an increase to both segments (deal to neither), or no increase to either segments (deal to both

segments). In broad outline, the experiment consists of three phases: collecting “prior” relative preference for two firms based on full descriptions of the services they offer; presenting information about various upcoming price increases and whom they affect; and collecting “posterior” relative preference for the two firms. These individual-level prior and posterior relative preference measures will allow model calibration.

A total of 560 undergraduate business school students took part in the experiment to fulfill subject pool requirements. Subjects were asked to make choices involving “audio content delivery” services similar to Rhapsody, eMusic, and Apple’s iTunes. This industry was chosen because it is poised to become the premier distribution medium for new music: All subjects, in fact, indicated a high level of familiarity with Web services and digital music, and appeared to find the study highly engaging. Descriptions of the two competing services were modeled on those being presently marketed, and were deliberately designed to be as balanced as possible so that few subjects would be a priori strongly opposed to switching, irrespective of pricing policy. All “numerical” differences between the two firms (e.g., number of songs available, number of downloads allowed), as well as percentages for price deals, were set at 20%.⁶ To render the descriptions as authentic as possible, details on over a dozen relevant dimensions were presented, among them technical specs (e.g., download speed, simultaneous channels, file types), extramusical services (e.g., contact lists, chat room capabilities), and contract conditions (e.g., annual length).

After being provided with description of the two firms, subjects expressed their relative preference by splitting 100 points between the two; these served as the individual-specific “prior” preferences for the firms. Respondents failing to express a preference were uncommon, under 3% of the total. However, to avoid misinterpreting their a priori favorite among the brands, we excluded them from further analysis. Following a filler task, subjects were then told to imagine that they had engaged their preferred firm for the

⁶ A pilot study showed 20% differences to be substantial enough to influence preference and subsequent choice, but not overwhelming in the sense of being insurmountable by attractive promotions.

full annual contract period, with good results, and also that customers from the rival firm reported similar reactions. Subjects (in all conditions) were then made aware of impending price increases; specifically, that “there will be a 20% increase” in service prices “from the current value of \$150 to \$180 for the coming year.” In addition, each subject then came by new, “pertinent information” about the two digital music services for the coming (annual) contract consistent with the condition he was in, as described in detail below.

4.3. Segments and Price Deal Patterns

Knowledge of prices had to be handled with great care so that, as in reality, firms would attempt to make consumers aware only of deals that they were offering to these consumers and not of deals that they were offering to other consumers (and not to them). However, as commonly happens, consumers could learn of these deals to other consumers from alternate sources or happenstance. Subjects were made aware of possible price deals through “flyers” offering them 20% off next year’s price, that is, the same price they had already been paying. For example, subjects in the $\{L, N\}$ condition (who got a no-increase price deal from their own firm, whereas the other firm increased price to both segments) received a flyer from their own firm offering continuing customers (only) “last year’s prices this year,” meaning an additional year of the \$150 fee level as opposed to \$180. Similarly, subjects in the $\{S, N\}$ condition (whose firm gave a better price to Switchers, whereas the other firm would increase price to both segments) “happened to come across” a flyer from their firm offering new customers (only) “last year’s prices this year,” meaning an additional year of the \$150 fee level as opposed to \$180. In the $\{N, S\}$ condition, subjects got a flyer from the other firm offering them (new customers) a price of \$150; and those in the $\{N, L\}$ condition came across a flyer where the other firm offered their continuing customers a price of \$150.

In the conditions where both firms made some type of offer ($\{L, S\}$ and $\{S, L\}$), a combination of the situations described earlier occurred, with subjects receiving two flyers. However, we counterbalanced which flyer the subject received first; order effects were not significant and we do not refer to them again. And, in the cases where all consumers got an offer from at least one of the firms (the seven cells in the “All” column and row of Table 1), participants were told that the offering firm was “pleased to announce the following special offer to all our customers, both current and new.” Regardless of the pattern of price increases and deals in the market, subjects were also made aware that by “asking around” they learned that the deals they came across were the only ones being offered. Our subjects reported no difficulties in accepting or

understanding the pricing scheme, perhaps due to its ubiquity. Consumers constantly get e-mails and flyers offering new/continuing consumers better prices. Prasad et al. (2003) have, in fact, used a similar scenario, where some consumers paid a higher and some a lower price in an e-shopping setting.

The experiment concluded with subjects being asked again for relative preference information (i.e., a 100-point split) and these served as the individual-specific “posterior” preferences. Debriefing indicated no skepticism or questioning regarding the cover story.

4.4. Model Estimation

Recall that each subject was asked to divide 100 points between the firms; after presentation of the promotional scenario (for each subject’s condition), they were asked for the same information. Thus, data consist of “prior” and “posterior” preference allocations for each subject. Estimation proceeds as follows. Table 2 presents expressions for each of the 16 promotional setups. Since we take each subject’s “prior” as his initial preference (α or β , depending on which firm he preferred), we are left with six parameters to estimate, $\{s, l, b, j, s', l'\}$. We do so based on minimizing the weighted least squares error between the stated posterior preference allocations and those predicted by the model.⁷ These, in turn, allow for tests of the form $H_0: \zeta = 0$ for $\zeta = \{s, l, b, j, s', l'\}$ or any subset thereof.

4.5. Results

We test to see whether the posited effects are supported for these data, and compare parameter estimates accordingly. Table 3 reports the results for three types of models: a “Full” model (i.e., behaviorist, with all four effects estimated); models where Jealousy, Betrayal, Loyalty, and Switching are each set *individually* to zero; and a “Strongly” rational model where behaviorist effects (Betrayal and Jealousy) are set to zero, so that only Switching and Loyalty are estimated. The last line of the table reports the results of Wald tests that compare each of the models to the “Full” model, and suggests which effects lead to additional explanatory power.

Table 3 indicates that all effects are very strongly significant.⁸ We first consider the Loyalty and Switch-

⁷ As a check, we had subjects also indicate which firm’s service they would choose; all chose the firm to which they had allotted more preference points. Importantly, the 100-point split preference allocations allow assessment of (upward or downward) preference changes even when the same brand would have been “chosen” twice. Note as well that many interactions among the various effects are addressed directly by the specifications of Table 2.

⁸ Because the six focal parameters are bounded, dividing by standard errors in the usual manner does not guarantee an asymptotic t distribution. Instead, all tests are Wald, whereby a “fuller” model is compared to restricted variants, requiring that a model be re-estimated subject to every parametric restriction one wishes to test.

Table 3 Estimated Parameters and Tests in TPI Scenario

	Switching "Full"	Switching $s_i, s'_i = 0$	Betrayal $b_i = 0$	Loyalty $l_i, l'_i = 0$	Jealousy $j_i = 0$	"Strong" $b_i, j_i = 0$
s_i	0.206	—	0.247	0.069	0.173	0.216
b_i	0.162	0.297	—	0.248	0.207	—
l_i	0.378	0.137	0.425	—	0.312	0.347
j_i	0.139	0.024	0.186	0.001	—	—
s'_i	0.287	—	0.327	0.159	0.257	0.301
l'_i	0.321	0.085	0.371	—	0.239	0.272
Parameters	6	4	5	4	5	4
$F_{[6-Params], 554}$	—	85.4	28.8	63.2	26.6	40.8
p vs. Full	—	$<10^{-32}$	$<10^{-6}$	$<10^{-24}$	$<10^{-6}$	$<10^{-16}$

ing effects under targeted promotions, i.e., l and s . We subsequently examine whether there are differences in the Switching parameters (i.e., $s_i = s'_i$) and Loyalty parameters (i.e., $l_i = l'_i$).

4.5.1. Loyalty and Switching Effects Under Targeted Promotions. The largest among all the effects is that of Loyalty under the targeted promotion scenario ($l_i = 0.378$; l subscripts are used throughout to denote an Increase-price scenario). This indicates that if a firm’s Loyals get a constant price when that firm’s Switchers face an increased price, the likelihood of repurchase rises from α to $\alpha + 0.378(1 - \alpha)$. In other words, approximately 38% of the Loyals who were (stochastically) not going to repurchase the brand will now do so. The Switching effect is also large ($s_i = 0.206$); this suggests that in a price-increase environment, when another brand offers one a constant price, the likelihood of repurchasing one’s own brand decreases by a factor of 0.206 (i.e., repurchase probability is lowered from α to 0.794α). Significantly positive values for the Loyalty and Switching effects provide support for our argument: In a price-increase scenario, consumers perceive constant price as a deal to them, and further, that there is a positive impact of imposing a price increase selectively on the other segment. While these two effects have been tested repeatedly for price-decrease type promotions (and thus lend support for the standard “strong rationality” effects), they have not to our knowledge been validated in the context of price increases.

Each effect need not, of course, operate in isolation. A particularly interesting case occurs if both firms offer a deal to the same particular segment, both trying to entice them: for example, if Firm X promotes to the Switching segment (i.e., to Firm Y’s Loyals) and Firm Y promotes to its own Loyals ($\{S, L\}$), or if Firm X promotes to its Loyals and Firm Y to its Switchers ($\{L, S\}$). These effects need not “wash out,” and in fact, the model makes specific predictions for such cases. Our analysis suggests that in the $\{L, S\}$ case, the repurchase probability for Firm X should change from α to $[\alpha(1 - 0.206) + (1 - \alpha)0.378]$, an increase if $\alpha < 0.647$. That is, when the repurchase

probability for Firm X is small, a deal of the same size offered to its Loyals by both firms will be in its favor.

4.5.2. Betrayal and Jealousy Effects. The two behaviorist effects, Betrayal and Jealousy, are also significant and sizeable. The Betrayal effect ($b_i = 0.162$) indicates that if a firm offers Switchers a constant price, but increases price to its own Loyals, the repurchase rate among Loyals decreases by a factor of 0.162 (i.e., from α to 0.838α). In a like manner, if the other firm offers a deal to its Loyals, but not to its potential Switchers, the Jealousy effect ($j_i = 0.139$) reduces the focal firm’s customers’ repurchase probability from α to 0.861α .

For simplicity, we have thus far considered Betrayal and Jealousy effects in isolation, but more than one effect may be operational at any given time. For example, if your firm offers its Switchers a deal (evoking Betrayal effects) but the other firm offers its Loyals a deal (evoking Jealousy effects), then the detrimental effects for your firm are huge: its repurchase probability can drop from α to $\alpha(1 - 0.206)(1 - 0.139)$, that is, by approximately 31%.

4.5.3. Loyalty and Switching Effects When the Firm Offers a Price Special to All Segments (i.e., No Price Increase). We note that the Loyalty effect when the firm gives a deal to all ($l' = 0.321$) is smaller than that when it offers a deal to its Loyals only ($l = 0.378$). A Wald test (against constraining the two parameters to be equal) shows that this difference is indeed significant ($l = l' = 0.347$; $F_{1, 554} = 4.68$; $p \approx 0.03$). This indicates that Loyals may “feel more special,” and may thus be more inclined to repurchase the brand when the price special is not an across-the-board strategy. On the other hand, we find that the Switching effect is larger when the other firm gives a deal to all ($s' = 0.287$) versus only to Switchers ($s = 0.208$). This may result since consumers are even more inclined to switch to a firm that does not increase prices to anyone (in an environment of increasing prices). This difference is also significant, comparing against a constrained version of the model ($s = s' = 0.253$; $F_{1, 554} = 8.97$; $p \approx 0.003$). As might be expected, these results hold against the jointly constrained model as well ($l = l' = 0.350$ and $s = s' = 0.254$; $F_{2, 554} = 4.98$; $p \approx 0.007$). Finally, as might be expected given the strong significance of the various individual effects, the “full” (behaviorist) model provides significantly better fit to the data than the “strongly rational model ($b_i, j_i = 0$), which ignores the (highly significant) Betrayal and Jealousy effects.

4.6. Testing Model Predictions and Robustness

It is reasonable to question whether the proposed model actually makes good predictions. That parameters may be statistically distinguishable from zero or one another may indicate only modest effects and a

Table 4 Directional Effects by Condition and Model Hit Rates

Your firm	Other firm	Case	Predicted				Repurchase probability
			Down		Up		
			Down	Up	Down	Up	
No One (N)	No One (N)	{N, N}	0	0	0	0	α
No One (N)	You (S)	{N, S}	18	1	0	0	$\alpha(1 - s)$
No One (N)	Others (L)	{N, L}	17	3	0	0	$\alpha(1 - j)$
No One (N)	All (A)	{N, A}	37	2	0	0	$\alpha(1 - s')$
You (L)	No One (N)	{L, N}	0	0	0	17	$\alpha + l(1 - \alpha)$
You (L)	You (S)	{L, S}	4	0	1	7	$\alpha(1 - s) + (1 - \alpha)l$
You (L)	Others (L)	{L, L}	0	0	2	7	$\alpha(1 - j) + (1 - \alpha)l$
You (L)	All (A)	{L, A}	14	4	2	9	$\alpha(1 - s') + (1 - \alpha)l$
Others (S)	No One (N)	{S, N}	10	1	0	0	$\alpha(1 - b)$
Others (S)	You (S)	{S, S}	18	0	0	0	$\alpha(1 - s)(1 - b)$
Others (S)	Others (L)	{S, L}	18	2	0	0	$\alpha(1 - j)(1 - b)$
Others (S)	All (A)	{S, A}	34	0	0	0	$\alpha(1 - s')(1 - b)$
All (A)	No One (N)	{A, N}	0	0	4	30	$\alpha + l'(1 - \alpha)$
All (A)	You (S)	{A, S}	8	4	4	12	$\alpha(1 - s) + (1 - \alpha)l'$
All (A)	Others (L)	{A, L}	9	3	4	14	$\alpha(1 - j) + (1 - \alpha)l'$
All (A)	All (A)	{A, A}	7	2	0	4	$\alpha(1 - s') + (1 - \alpha)l'$
							Hit rate (%)
"Full" behaviorist			194	22	17	100	88.3
$s = s', l = l'$			191	24	20	98	86.8
"Strong": $b = 0, j = 0$			141	31	25	85	80.1
$b = 0, j = 0, s = s', l = l'$			138	33	28	83	78.4

large sample. Here, we take up the issue of how well the "Full" model fares against restricted variants and its performance within the various cells of the design.

Evaluating the accuracy of directional predictions is complicated by the fact that respondents' posterior preferences could be greater, lesser or the same as their prior ones; the class of models estimated here makes strong directional predictions in each cell and only one, the {N, N} baseline cell, posits no change at all. For example, the Switching, Jealousy, and Betrayal effects each posit that all posterior preferences decrease and, consequently, interactions between those effects will posit even greater decreases. For example, when your firm offers Others a deal (Betrayal) and the other firm offers You a deal (Switching), we have a repurchase probability of $\alpha(1 - s)(1 - b)$, which must be smaller than α . In some cases, however, such as where both firms offer You a deal (generating Switching and Loyalty), the corresponding expression $\alpha(1 - s) + (1 - \alpha)l$ can be larger or smaller than α . Because all these expressions are continuous in parameters, only in the baseline cell will any model make a prediction of exactly no change. Consequently, we consider those cases where the predicted and the actual response each went up or went down, resulting in four possibilities and allowing a calculation of hit rate. These counts and summary measures appear in Table 4.

The tabulated figures tell a clear story. Overall, the "Full" behaviorist model does very well, with a hit rate of 88.3%. Restricting the Switching ($s = s'$) and Loyalty ($l = l'$) parameters to equality (not to zero) lowers the hit rate to 86.8% (*ns*); by contrast, the "Strong"ly rational model entails a large decrement, to 80.1% ($p < 0.005$); and further imposing $\{s = s', l = l'\}$ reduces the hit rate still further, to 78.4% ($p < 0.001$).

The repurchase probabilities listed in the last column of Table 4 can be classified as those that specify no change (only the first, {N, N}), two that specify an increase, seven that specify a decrease, and six that can go either way. We look at each of these subsets in more detail. The "Full" model fares especially well within the cells that make strong downward directional predictions, those in which your firm offers a deal to No One or to Others; in both these cases, your own firm is doing nothing for you and, in the latter set, is actively Betraying you. For these seven cells, the model predicts correctly 152 of 161 times, for a 94.4% hit rate. In the two cells where the prediction is always for a Loyalty-induced increase (via l and l'), the hit rate is $47/51 = 92.2\%$. As might be expected, performance was more mixed for the other six cells, which do not make unidirectional predictions. For these, the hit rate was $95/121$, or 78.5%. Note that in all these cells subjects were presented

with highly complex deal patterns, which may have also led to some confusion. Still, the “Full” model’s directional performance was quite good even in this subset of cells—on par, in fact, with the Behaviorist model overall.

In sum, we conclude that subjects reacted to price-increase deals in a manner largely consistent with the model’s predictions, even when multiple deals were offered by the two firms. We next examine whether price increases induce reactions similar to those of price decreases reported in prior literature.

4.6.1. Benchmarking Against the Targeted Price Decrease (TPD) Scenario. As an illustrative benchmark against our results, we ran a closely matched experiment for the TPD scenario using 201 (different) undergraduate subjects, drawn from the same subject pool. This also provides a robustness check for Loyalty, Betrayal, Switching, and Jealousy effects in the TPD scenario, as the resulting estimates can be compared, at least ordinally, to those obtained by Feinberg et al. (2002). We replicated their design, estimating the same four effects $\{s_D, b_D, l_D, j_D\}$, which allowed us to assess parametric differences relative to an identical underlying model structure. All four effects are significant at $p = 0.0001$ as they, too, found: $s_D = 0.260$; $b_D = 0.112$; $l_D = 0.325$; $j_D = 0.162$. Moreover, the effect sizes (parameter values) are roughly concordant, and a Jonckheere-Terpstra test does not reject that the parameters are in the same order (cf. their Table 6, p. 288) in TPI, TPD, and the TPD scenario run in Feinberg et al. (2002).

While it is tempting to compare the TPI and TPD effects, we caution the reader that the relative size of these effects clearly depends on the exact framing of the price scenarios. Nonetheless, eyeballing the parameter values for the various parameters in TPI and TPD suggests that some effects are stronger in the increase condition. Parametric equality across models can be assessed by restricting the parameters in question to be the same, jointly re-estimating the models, and computing likelihood ratios. Such tests reveal that the Loyalty and Betrayal effects are each significantly different (in fact, stronger) between the TPI and TPD environments ($p < 0.01$), but that the Switching ($p > 0.05$) and Jealousy effects are not ($p > 0.2$).

One reason for the somewhat larger effects in the TPI environment may be that all deals in the TPI scenario occur in a loss context (to the segment not offered the deal) given the reference point of status quo (no price change), whereas all TPD effects occur in a gain context (to the segment offered the deal). Consequently, effects generated by your firm like Loyalty and Betrayal (a “negative” effect, in particular) might be expected to be stronger in TPI than in TPD. Further enhancing the degree of asymmetry is that, in the TPI scenario, the loss is to you; by contrast, in

the TPD scenario, the gain is to others. Simply put, this suggests, for example, that consumers feel more “betrayed” by their firm when Switchers are offered a constant price (when they themselves are faced with a price increase) versus when Switchers get a price cut (when they themselves are offered a constant price).

Next, we consider implications of the effects generated by TPI for a competitive market.

5. Competitive Targeting Implications

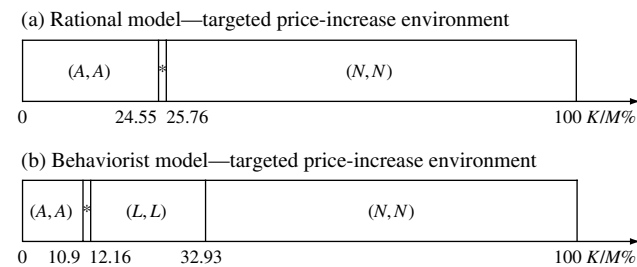
The strategic implications of the Loyalty, Switching, Betrayal, and Jealousy effects in the price-increase environments can be surmised by developing a competitive game in which two competing firms enact promotional decisions to maximize their steady-state payoffs. A firm can target its own Loyal customers with a price special by raising its price only to Switchers, so that its profit margin from Switchers is M and that from Loyals is $M - K$, where $M > K$. We denote this strategy as L , since Loyals are favored. Alternatively, a firm can target only Switchers by raising its price only to Loyals, and this is denoted as S . In this case, the firm’s profit margin is M from Loyals and $M - K$ from Switchers. A firm may also opt to raise its price to all consumers in the market such that no targeted promotion is instituted. In that case, it obtains the same profit margin M from all customers. We denote this strategy as N . Finally, a firm may decide to favor all of its customers, so that its profit margin from all consumers is $M - K$, and we denote this strategy by A .

Let $\tilde{\Pi}^{hk}$ be firm i ’s steady-state payoffs, given that firm i chooses promotion strategy h and the rival firm chooses strategy k , where $h, k = \{L, S, N, A\}$. We derive $\tilde{\Pi}^{hk}$ in the appendix, using the parameter estimates from Table 3, and conduct three sets of equilibrium analysis. First, we derive the competitive equilibria under TPI to examine firms’ optimal strategies in this environment. Second, we conduct the same analysis with asymmetric market shares for the two firms to investigate how market share may moderate the firm’s promotion strategy. Third, we alter the proportion of people in the market who are aware of prices to the other segment, i.e., we vary the proportion of consumers who follow the behaviorist (as opposed to the “strongly rational) model, and derive the associated competitive equilibria.

5.1. Equal Market Shares

The competitive equilibria, given that the baseline market shares are the same for both firms, are illustrated in Figure 1. Equal baseline market shares are reflected in the fact that all consumers are equally likely to switch away from either Firm A or Firm B on the next purchase occasion once having purchased

Figure 1 Symmetric Competitive Equilibria



Note. The figures are drawn with $\alpha = \beta = 0.5$; * indicates equilibria (L, A) and (A, L) .

from the firm ($\alpha = \beta$). Figure 1(a) reflects the “strong rationality” model and Figure 1(b) the “behaviorist” model.

Looking at Figure 1(a) (the strong rationality model), one can see that, except for a small range of values ($24.55\% < K/M < 25.76\%$), the competing firms choose not to target their price increase to a single segment alone: They either offer price specials (i.e., no price increase) to all (equilibrium (A, A)) or increase their prices across the board (equilibrium (N, N)). In the small range where they do practice targeted pricing, they favor Loyals. Intuitively, when offering price specials indiscriminately (i.e., to all), a firm weakens its customer loyalty slightly (i.e., $l' < l$), thus reducing the likelihood of its own customers’ repeated purchases. At the same time, the firm becomes much more attractive to the rival’s customers. Thus, for small price specials, each firm looks to increase its sales by making the price special available to all. As the cost of the special increases, one of the firms will find it more profitable to retain its customers and cut down the cost of the special through favoring only its own repeat customers. When the cost of the special is too high, none of the firms offers any special.

When one includes behaviorist effects (Figure 1(b)), the range where a firm deploys a targeted promotion strategy increases considerably ($10.9\% < K/M < 32.93\%$), while the ranges respectively for no specials and across-the-board specials shrink. In fact, whenever a firm targets its customers with a special, the targeted customers are always loyal customers. This occurs because targeting own loyal customers can help to retain and attract more customers through the sizable Jealousy effect, which also increases the profitability of the special. A firm would face the double jeopardy of strong Betrayal and Jealousy effects if it were to target only Switchers.

5.2. Asymmetric Market Shares

One can argue that if a firm has a larger market share, then it cannot afford not to increase price to its Loyals. Therefore, it cannot afford to favor Loyals in a

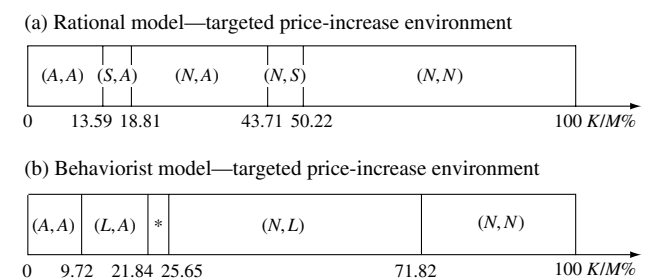
TPI strategy. However, the counterargument is that it cannot afford to upset its loyal consumer base either. If the Loyals become aware of a better price to Switchers and are upset by this, detrimental effects may ensue. How should a firm strike a balance between these two forces? Our experimental and competitive analysis results suggest that a high market share firm should offer fewer across-the-board price specials, i.e., it should enact more across-the-board price increases than should a low market share firm. However, when offering price specials to just one of the segments, both firms decisively favor Loyals.

To illustrate how differences in market share can impact equilibrium pricing strategy, we consider a concrete example in which consumers who have purchased from Firm A have a probability of switching (away from Firm A) of 0.2 (i.e., $\alpha = 0.8$), but the corresponding switching probability for Firm B is 0.6 (i.e., $\beta = 0.4$). In such a market, Firm A has the larger market share in steady state (absent any targeted promotion), commanding a 75% share. When targeted promotions are considered, the equilibrium strategies are illustrated for this market in Figure 2.

Figure 2(b) shows that, as in the symmetric behaviorist case, whenever a firm implements a targeted price increase, it always favors its Loyals regardless of whether a firm has a larger or smaller share. The larger market share firm (Firm A) always targets its Loyal customers whenever it is profitable to employ targeted promotions ($9.72\% < K/M < 21.84\%$) because the larger share firm, with its large loyal base, has much to lose via Betrayal and/or Jealousy effects. The small share firm, by contrast, targets its own Loyals to take advantage of the Jealousy effect and takes the aim at the loyal base of the larger share firm. However, the larger share firm (Firm A) has a far narrower promotion rate range ($K/M < 21.84\%$) compared to its smaller rival (Firm B) ($K/M < 71.82\%$). This reflects the fact that the larger share firm must forgo too much profit if it does not increase price to its large loyal base. In other words, it is more costly for the larger share firm to offer TPI.

A comparison of Figures 2(b) and 2(a) indicates that equilibrium targeted pricing strategies incorporating

Figure 2 Asymmetric Competitive Equilibria



Note. The figures are drawn with $\alpha = 0.8$ and $\beta = 0.4$; * indicates (N/A) .

behaviorist effects are clearly distinct from those assuming merely “strong”ly rational (and no behaviorist) effects on the part of consumers. As in the symmetric market share case, in the asymmetric case, we too find that the consideration of behavioral effects generates a larger range of parameter values where a targeted price increase is used by one of the firms. Furthermore, the incidence of indiscriminate price specials (A, A) or indiscriminate price increases (N, N) is reduced. This means that behavioral effects do not necessarily discourage firms from deploying targeted (to one segment) promotions and, indeed, they can incentivize firms to do more of it, so long as the firms favor the correct consumer segment. Figure 2 shows that while targeting Switchers (via a price deal) is among the equilibrium strategies in the strong rationality model, it is not in the behaviorist model (for the parameter values empirically obtained).

As market share for a firm increases (Figure 1(b) versus Figure 2(b)), our results suggest that promotions are supported for a smaller range (of promotion proportion) for the larger share firm, but a wider range for the smaller share firm; this difference reflects the fact that less profit is lost by not increasing price to Loyals if the consumer base is small. In an environment of increasing prices, large share firms may implement price increases (though small share firms may not), and could well use this state of affairs as an opportunity to build up their customer base. Further, even for the smaller firm, it may be better to favor its Loyals (and not Switchers) if it implements targeted price increases. Thus, counter to intuition, even when the larger market share firm is not promoting to its Loyals—and, hence, there is greater opportunity for the other firm to attract these customers—it may still be better for the other firm to focus on retaining its own Loyals rather than seeking to attract the larger share firm’s customers.

5.2.1. Decrement in Firm Profits by Adopting the Incorrect Pricing Strategy. Our competitive analysis indicates that firms could sacrifice a great deal of profit by choosing a suboptimal targeting strategy, either by adopting a strong rationality model instead of allowing for behaviorist effects, or by implementing an across-the-board price increase instead of TPI. This is especially so when the two competing firms have asymmetric market shares.

For instance (in Figure 2(b)), when $25.65\% < K/M < 43.71\%$, the small share firm should raise its price to all but its own Loyal customers, so that the Loyals are thereby targeted (see Figure 2(b)), given that the rival firm raises its price to all consumers. However, in this range, the small share firm would choose the incorrect policy of indiscriminately offering specials if it ignores behaviorist effects (see Figure 2(a)). In that case, the firm’s profitability can decrease by as much


as 13%. Similarly, when $43.71\% < K/M < 50.22\%$, the small share firm may use the wrong strategy of targeting Switchers when, in fact, it should target its own loyal customers. In that case, the firm’s profit can decrease by as much as 24%. An inappropriate pricing strategy can thereby result in substantial foregone profit, arguing that firms need to well understand the nuances of targeted pricing in a competitive market.

5.3. Awareness of Prices to the Other Segment

Our analyses thus far have presumed that all consumers are aware of prices to the other segment. In reality, only a proportion of consumers may be aware of these prices. Further, some consumers may be aware of (lower) prices to the other segment but it may not affect the way they behave, i.e., they may not care that the other segment is offered a better price than they are. We use “aware” to refer to those consumers who both know and care about prices to the other segment. We now consider the case where only a fraction of consumers may be “aware” and, hence, susceptible to the influence of Betrayal and Jealousy effects; the remainder may only know about or care about deals that they themselves receive, consistent with the strong rationality model. To examine this case, let γ represent the fraction of consumers in the market who know or care only about the promotions they themselves receive; these can be termed the “rational” consumers. The remainder ($1 - \gamma$ proportion) represent the “aware” (i.e., know and care) segment, who are susceptible to Betrayal and Jealousy effects. Figure 3 illustrates the equilibrium strategies for the asymmetric market share case, holding $K/M = 15\%$.

Figure 3 shows that when a vast majority of the customers in the market are neither aware of nor care about prices to others ($\gamma > 91.62\%$), the larger share firm should target Switchers (for this 15% price special) and the smaller share firm should offer a price special to all (as also indicated in Figure 2(a)). However, as more consumers become aware of (and care about) the deals offered to the others, the larger share firm first matches the strategy of the smaller share firm, offering price specials to all ($48.52\% < \gamma < 91.62\%$), and then targets its own Loyals with a special ($\gamma < 48.52\%$). Intuitively, as more customers become “aware and care,” the Betrayal effect plays an ever more prominent role such that any benefit

Figure 3 Competitive Equilibria—Mixed Consumers



Note. The figure is drawn with $\alpha = 0.8$ and $\beta = 0.4$.

the larger share firm obtains from incremental sales will be offset by the negative effect on its own Loyal customers. Consequently, the way to attract Switchers and generate incremental sales is for the larger share firm to offer deals to all, minimizing the Betrayal effect. However, the effect of using indiscriminate (i.e., to all) specials to attract Switchers diminishes as more and more consumers become “aware and care” (s' is larger under the strong rationality versus the behaviorist model), whereas the Loyalty effect increases (l is larger than l' under the behaviorist model). This explains why the larger share firm begins to target its own Loyals with the price special. Compounding this is the Jealousy effect, offering additional inducement for a firm to favor its Loyals, since doing so shores up loyal customers and tempts those of the rival.

This analysis thus suggests that the extent to which price information is disseminated in a market, or the extent to which customers care to pass along their pricing information, can have a substantial impact on a firm's targeting strategy. When many customers in a market are the “aware and care” type, targeting own loyal customers is, once again, not only a safe strategy but also a profitable one.

6. Conclusions and Potential Extensions

There is presently neither extant theory nor empirical evidence to guide managers in setting promotions when price increases are part of the feasible “strategy space,” that is, when there is pricing power. We show in this paper that when a rival firm enacts a price increase, offering a constant price can indeed be a viable strategic option for a firm. Our analysis cautions firms against presuming that “behaviorist” effects (Betrayal, Jealousy) are negligible. Indeed, firms can forgo a substantial proportion of profit by allowing only for the “strong(ly) rational” and widely studied effects of Switching and Loyalty, choosing a suboptimal pricing strategy as a result. Perhaps the most salient implication arising from our analyses is this: firms will not necessarily fare best by imposing an across-the-board price increase on all consumers but rather by imposing it selectively on the appropriate segment.

The paper began with a list of questions we hoped to answer. In response, we find in summary that

- Being offered a constant price by a firm when others are faced with a price increase can be broadly construed as a price promotion: there can be a positive impact to implementing a price increase selectively on one segment alone.
- When the firm offers a deal to all consumers, the Loyalty effect is smaller than when it offers a deal to its Loyals only, indicating that Loyals may “feel

more special” and thus be more inclined to repurchase the brand when the price special is not an across-the-board strategy. On the other hand, we find that the Switching effect is larger when the other firm offers a deal to all compared with only to Switchers. This may come about because consumers are even more strongly inclined to switch to a firm that does not increase prices to anyone (in an environment of increasing prices).

- Effects are not identical in the TPI and TPD environments. Loyalty and Betrayal effects are significantly larger in the TPI environment. Also, the optimal pricing strategies in the two scenarios are different in ways difficult to anticipate in the absence of a game-theoretic analysis.

- Firms can achieve greater profits in many cases by imposing price increases selectively instead of uniformly.

- Many different equilibrium pricing strategies can come about depending on variety of factors: promotion magnitude, firm market shares, the proportion of “aware and care” consumers, and whether one allows for behaviorist effects over and above the traditional rationalist ones. An unusually wide spectrum of possible equilibrium types can occur, including deals to Loyals by both firms, deals to neither segment by both firms, or one firm promoting to Loyals and the other to Switchers.

- As market share for a firm increases, our results suggest that promotions are supported for a smaller promotion range for the larger share firm and a wider range for the smaller share firm. Further, even for the smaller share firm, it may be better to favor its Loyals (and not Switchers) if it implements targeted price increases. Thus, counter to intuition, even when the larger market share firm is not promoting to its Loyals—and, hence, there is greater opportunity for the other firm to attract these customers—it may still be better for the other firm to focus on retaining its own Loyals rather than seeking to attract the larger share firm's customers.

- In an environment of increasing prices, large share firms may implement price increases whereas small share firms may not. Small share firms may use this as an opportunity to try and build their customer base.

- The equilibrium strategy depends critically on proportion of “aware and care” consumers in the market. Presuming that all consumers have full knowledge of various promotional offers can lead to serious errors in targeting.

A strength of the general approach used here is the fusion of several widely applied methodologies. The stochastic model allows one to encode six different effects and specify how they drive repurchase probabilities away from their baseline levels. The

Markovian framework is particularly amenable to estimation using readily collected experimental data, and indicates strongly significant values for all six effects parameters (Betrayal, Jealousy, and two distinct types each of Switching and Loyalty). And, finally, game theory addresses competitive equilibria and how they differ with varying promotion magnitude, market share, proportion of “aware and care” consumers, and across the rationalist and behaviorist models.

There are, of course, limitations to our approach. We have used a steady-state analysis to derive competitive strategy implications; this assumes either that firms have a relatively low discount rate for future payoffs or that consumers respond quickly to changes in promotional strategy. Without these assumptions, a highly complex differential games framework would be necessary to derive equilibria and then only numerically. We also adopted a first-order model, so that consumers’ purchase probabilities are dependent only on what they did in the directly previous period. Such a modeling framework overlooks that, for some consumers at least, the longer they stay with a brand, the less likely they may be to switch. Extending the model to a higher-order framework could incorporate such behavior, though at a cost of some complexity.

For reasons of tractability, we also used a two-firm, two-segment (Loyal, Switcher) market in which firms can promote to Loyals, promote to Switchers, or not promote. In “real world” markets, there are typically more than two relevant firms. However, even if there are not two but three (or more) firms, we would still expect the types of effects we hypothesize and find in the two-firm case—thus, if your firm increases price to you but another firm does not, then you would be more likely to switch to it (Switching or substitution effect); if your firm does not increase price to you, but other firms do, then you would be more likely to stay with your firm (Loyalty effect); if your firm increases price to you but not to others, then you would feel “betrayed” and would be more likely to switch out (Betrayal effect); and, if your firm increases price to you but another firm does not increase its price to its Loyals, then you would be jealous and would be more likely to switch to it (Jealousy effect).

In the multifirm setting, however, the size of the effects may depend on how many firms are following each type of policy—effects would be exacerbated if fewer firms are following a certain policy (whether the effects are negative or positive) and vice versa. Thus, if your firm is the only one not increasing price to Loyals, whereas all the others are, then the Loyalty effect could be even stronger. Contrarily, if all firms are increasing prices but your firm is the only one not

increasing price to Switchers, then the Betrayal effect could be very strong. Another issue relevant in the case of multiple firms is that consumers may consider and learn about prices from only a few (two or three) firms, so that policies of all firms may not matter and only those of a few do.

Loyalty, Switching, Betrayal, and Jealousy effects were assumed symmetric for the two firms, by the model and the experiment. In our experimental data, with highly similar firms, relaxing this symmetry produced negligible differences in parameter values (and, thus, in implications following from them). However, this certainly should not be assumed for all data sets or promotional scenarios. We also assumed the price increase K to be exogenous. This assumption may be valid in situations where firms try to pass the cost increase to customers, so that the price always increases by the amount of the cost increase. Since the cost increase is exogenous, K may be as well. However, future research should relax this assumption and extend the class of results established here to scenarios in which K is fully endogenous. We suspect that endogenizing K should not qualitatively change the nature of our conclusions but would help instead to pin down factors driving the magnitude of K . Of course, to endogenize K , we need to further specify the reason for the price increase about which our analyses and, in particular, experiments were deliberately silent. For instance, prices could increase because of material costs, channel issues, reduced competitive pressure, or simply to generate profit. Future experimentation should repeat scenarios similar to those used here but should present respondents with specific, alternate reasons behind the price increase; parameters could then be measured for each price-increase driver and corresponding equilibria compared.

We further anticipate that field experiments could clarify a number of important empirical issues, chief among them whether the model’s posited effects tend to be over- or underestimated in a laboratory setting. They could also assess a force we attempted to strenuously control: inferences about why firms are raising prices in the first place. In some cases, such as when fuel or transportation costs go up, consumers may view price increases not only as fair but equally so across all firms. In others, as when a labor strike affects just one firm increases may be seen as fair for that firm but not if another firm takes advantage of the situation. And, in still other cases, for example, as a result of a natural disaster or deliberate cartel action, any attempt to raise prices may be seen as unwarranted.

In terms of implementing targeted price increases in practice, some environments may be more conducive than others. If many consumers are aware of

prices to other segments and are upset if these prices are lower than those they are being offered, manufacturers may not wish to practice targeted pricing. The proportion of “aware” consumers can also vary by industry. For example, in business-to-business industries, prices to various customers may be more private than in business-to-consumer industries. Hence, the proportion of “aware” consumers may be intrinsically smaller in the former. Further, Web-based business may allow for easier price comparisons, allowing for a larger proportion of “aware” consumers.

Our analysis highlights a finding of potential importance to practicing managers: that it may be prudent to impose price increases selectively rather than to all one’s potential customers. Managers would also do well to realize that the mere act of failing to increase prices to one segment can, depending on how one treats another segment, be viewed as a bona fide price promotion, much as being given a day off from work (when others are not) can be viewed as a paid vacation, as opposed to a scheduled holiday. Given that managers may, in accordance with our findings, choose to enact a “nonincrease” type promotion, they should realize that they are different from standard “decrease” type promotions: effects, strengths, and resulting competitive equilibria systematically differ. Given that targeted price increases are not currently common, we hope that our findings serve to expand the practice of this type of policy. From the consumer’s perspective, this would allow some consumers to continue with a lower price while enabling firms to enhance profitability. As such, this paper hopes to introduce a valuable strategic option to the manager’s toolbox as well as to encourage further research on targeted promotional programs in general.

7. Electronic Companion

An electronic companion to this paper is available as part of the online version that can found at <http://mansci.journal.informs.org/>.

Acknowledgments

The authors thank Elea Feit, Andy Gershoff, Michel Wedel, David Wooten, Carolyn Yoon, and Jie Zhang for their assistance in conducting the experiments. All authors contributed equally to this paper.

Appendix

To obtain the steady-state profit function, we need first to derive a firm’s total sales and sales on promotion in steady state. We illustrate here how to derive each firm’s steady-state profits when Firm A targets Switchers and Firm B does not promote. All other cases are analogous and omitted for brevity.

Let $Sales_A^{SN}$ and $Sales_B^{SN}$ be Brand A’s and B’s total sales in steady state, respectively, when Firm A offers deals to Switchers and Firm B does not promote. From Table 2, by normalizing the “size” of the market to one, we have

$$\begin{aligned} Sales_A^{SN} &= \frac{1 - \beta(1 - s)}{[1 - \beta(1 - s)] + [1 - \alpha(1 - b)]} \\ Sales_B^{SN} &= 1 - Sales_A^{SN}. \end{aligned} \quad (1)$$

This expression represents steady-state sales in each period and comprises both promotional and nonpromotional sales.

Because Firm A’s promotions are targeted at Switchers, a fraction of its sales are made on deal. In steady state, sales on promotion are given by

$$PromSales_A^{SN} = [1 - \beta(1 - s)]Sales_B^{SN}. \quad (2)$$

In other words, Firm A’s promotional sales are equal to the fraction of Brand B’s buyers who switch to Brand A due to its promotional incentives. Firm B’s promotional sales, $PromSales_B^{SN}$ are zero in this case, as it offers no promotion.

Then, given M (the normal margin for a brand) and K (the unit cost of redemption inclusive of any costs of targeting a consumer, handling, and administration), a firm’s profit in steady state can be written, in general, as

$$\tilde{\Pi}_i^{hk} = M[Sales_i^{hk}] - K[PromSales_i^{hk}], \quad (3)$$

where $i = A, B$ and $h, k = \{N, S, L\}$. Thus, in the specific case examined here, one obtains

$$\begin{aligned} \tilde{\Pi}_A^{SN} &= M[Sales_A^{SN}] - K[PromSales_A^{SN}], \\ \tilde{\Pi}_B^{SN} &= M[Sales_B^{SN}]. \end{aligned}$$

Once all profit functions are derived, a straightforward way to utilize steady-state analysis in a game-theoretic setting is to construct an infinitely repeated game in which each of the competing firms chooses its promotional strategy, that is, which segment to target: Switchers (S), Loyals (L), Neither (N), or All (A). For the purposes of formal analysis, the firms’ payoffs can be taken to be their steady-state profit values $\tilde{\Pi}_i^{hk}$, assuming that all consumers in the market are susceptible to the influence of both Betrayal and Jealousy effects.

References

- Adams, J. S. 1965. Inequity in social exchange. L. Berkowitz, ed. *Advances in Experimental Social Psychology*, Vol. 2. Academic Press, New York, 267–299.
- Bolton, L. E., L. Warlop, J. W. Alba. 2003. Consumer perceptions of price (un)fairness. *J. Consumer Res.* 29(4) 474–491.
- Campbell, M. 1999. Perceptions of price unfairness: Antecedents and consequences. *J. Marketing Res.* 36(2) 187–199.
- Chen, Y. 1997. Paying customers to switch. *J. Econom. Management Strategy* 6(4) 877–897.
- Feinberg, F. M., A. Krishna, Z. J. Zhang. 2002. Do we care what others get? A behaviorist approach to targeted promotions. *J. Marketing Res.* 39(3) 277–291.
- Garbarino, E., O. F. Lee. 2003. Dynamic pricing in internet retail: Effects on consumer trust. *Psych. Marketing* 20(6) 495–513.
- Givon, M. 1984. Variety-seeking through brand switching. *Marketing Sci.* 3(1) 1–22.

- Greenberg, J. 1986. Determinants of perceived fairness of performance evaluations. *J. Appl. Psych.* **71**(2) 340–342.
- Jeuland, A. P. 1979. Brand choice inertia as one aspect of the notion of brand loyalty. *Management Sci.* **25**(7) 671–682.
- Kahn, B. E., J. S. Raju. 1991. Effects of price promotions on variety-seeking and reinforcement behavior. *Marketing Sci.* **10**(4) 316–347.
- Kahneman, D., J. L. Knetsch, R. Thaler. 1986a. Fairness as a constraint on profit seeking: Entitlements in the market. *Amer. Econom. Rev.* **76**(4) 728–741.
- Kahneman, D., J. L. Knetsch, R. Thaler. 1986b. Fairness and the assumption of economics. *J. Bus.* **59**(4) S285–S300.
- Klemperer, P. 1987. Markets with consumer switching costs. *Quart. J. Econom.* **102**(2) 375–394.
- Klemperer, P. 1995. Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade. *Rev. Econom. Stud.* **62**(4) 515–539.
- Lal, R. 1990. Price promotions: Limiting competitive encroachment. *Marketing Sci.* **9**(3) 247–262.
- Martins, M., K. Monroe. 1994. Perceived price fairness: A new look at an old contract. *Adv. Consumer Res.* **21** 75–78.
- Naert, P. A., A. V. Bultez. 1973. Logically consistent market share models. *J. Marketing Res.* **10**(3) 334–340.
- Nunes, J. C., C. W. Park. 2003. Incommensurate resources: Not just more of the same. *J. Marketing Res.* **40**(1) 26–38.
- Prasad, A., V. Mahajan, B. Bronnenberg. 2003. Advertising versus pay-per-view in electronic media. *Internat. J. Res. Marketing* **20**(1) 13–30.
- Shaffer, G., J. Zhang. 1995. Competitive coupon targeting. *Marketing Sci.* **14**(4) 395–416.
- Shaffer, G., Z. J. Zhang. 2000. Pay to switch or pay to stay: Preference-based price discrimination in markets with switching costs. *J. Econom. Management Strategy* **9**(3) 397–424.
- Shaffer, G., Z. J. Zhang. 2002. Competitive one-to-one promotions. *Management Sci.* **48**(9) 1143–1160.
- Stark, O., J. E. Taylor. 1989. Relative deprivation and international migration. *Demography* **26**(1) 1–14.
- Taylor, C. R. 2003. Supplier surfing: Competition and consumer behavior in subscription markets. *RAND J. Econom.* **34**(2) 223–246.
- Thisse, J. F., X. Vives. 1988. On the strategic choice of spatial price policy. *Amer. Econom. Rev.* **78**(1) 122–137.
- Urbany, J. E., T. J. Madden, P. R. Dickson. 1989. All's not fair in pricing: An initial look at the dual entitlement principle. *Marketing Lett.* **1**(1) 17–25.