DAVID R. BELL, GANESH IYER, and V. PADMANABHAN*

Conventional wisdom suggests that the main effect of price promotion is on brand switching (i.e., secondary demand); however, some recent studies demonstrate that the primary demand expansion effect can be considerably larger than previously believed. A significant driver of this primary demand effect is consumer stockpiling in response to price promotions. Indeed, experimental studies have shown that additional inventory on hand can lead to an endogenous increase in consumption. The authors develop a model of price competition between firms in response to the stockpiling and subsequent consumption dynamics of consumers. In this setting, the flexible consumption effect causes more intense price competition, deeper promotions, and an increase in the frequency of promotions. The authors use two years of scanner panel data from eight product categories and 4313 stockkeeping units to test three implications of the theoretical model; they find strong support for each.

Price Competition Under Stockpiling and Flexible Consumption

The nature of consumer response to a price promotion is of substantial importance to managers and has received considerable attention in the literature (see Blattberg and Neslin 1990). Recent studies (e.g., Bell, Chiang, and Padmanabhan 1999; van Heerde, Leeflang, and Wittink 2000) demonstrate that the primary demand expansion effect of a promotion can be significantly larger than previously believed (e.g., Chiang 1991; Gupta 1988). Primary demand expansion is distinct from secondary demand shifting created by brand switching in a category: It is an increase in the overall category volume. Price promotions can induce consumers to stockpile by purchasing higher-than-usual quantities. However, to ascertain the true primary demand effect of a promotion, it is important to distinguish the case of *pure stockpiling* from that of *flexible consumption*: that is, additional consumption induced by the presence of additional inventory on hand.

When consumers stockpile without increasing their consumption, a temporary expansion in demand is observed, that is, an increase in purchase quantity in response to a price promotion, followed by a longer-than-normal elapsed time before consumers reenter the market for a subsequent purchase. This is different from a situation in which the promotion-induced increase in purchase quantities does not significantly extend the time until the next purchase in the category, implying that there has been an increase in consumption. Figure 1 illustrates both cases.

The inventory effect is depicted in the bottom half of Figure 1: Higher inventory levels create primary demand expansion by endogenously increasing the usage rate of the consumer.[1] Several experimental studies show that increased inventory on hand leads to increased consumption of the product (see, e.g., Moore and Winer 1978; Wansink and Deshpande 1994). Until recently, most empirical studies on secondary data assumed that individual consumption rates were fixed. Ailawadi and Neslin (1998) relax this assumption and calibrate an econometric model of flexible consumption using scanner panel data. They find evidence that price promotions can lead consumers to increase consumption to a greater extent for yogurt than for ketchup. Bell, Chiang, and Padmanabhan (1999) report cross-category differences in consumption effects: Categories such as bacon, salted snacks, soft drinks, and yogurt exhibit

*David R. Bell is an associate professor, Wharton School of Business, University of Pennsylvania (e-mail: davidb@wharton.upenn.edu). Ganesh Iyer is an assistant professor, Haas School of Business, University of California at Berkeley (e-mail: giyer@haas.berkeley.edu). V. Padmanabhan is John K. Wallace Jr. and Ellen A. Wallace Distinguished Professor, Olin School of Business, Washington University (e-mail: padmanabhan@olin.wustl.edu), and Visiting Professor of Marketing, INSEAD–Singapore (e-mail: paddy.padmanabhan@insead.edu). The authors are grateful to Scott Neslin, Seethu Seetharaman, Christophe Van den Bulte, and the three anonymous *JMR* reviewers for their valuable suggestions and comments.

[1]In keeping with existing literature, we use the terms “consumption effect’” and “inventory effort” interchangeably. Both terms refer to the situation in which an increase in inventory on hand leads to an increase in consumption.

Journal of Marketing Research
Vol. XXXIX (August 2002), 292–303

Figure 1
PURE STOCKPILING AND FLEXIBLE CONSUMPTION

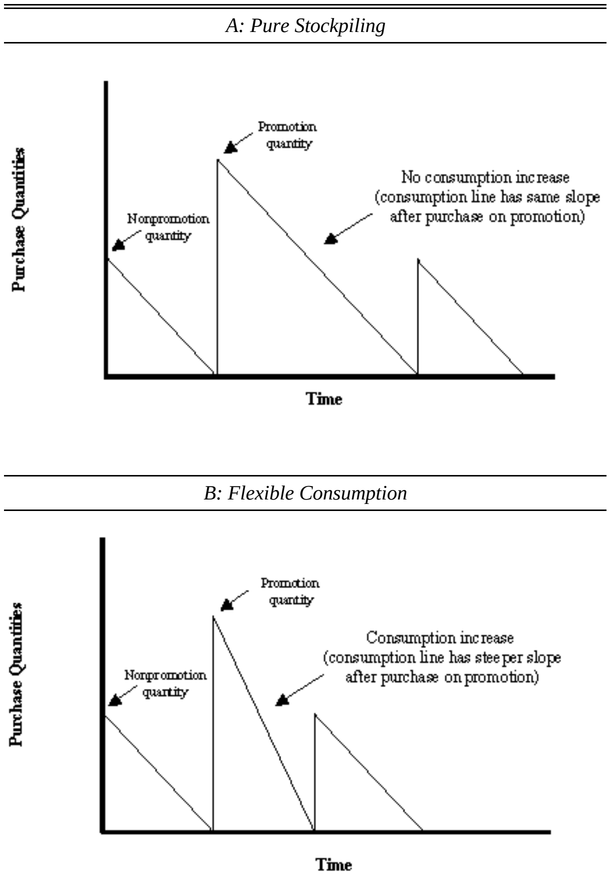

Notes: For ease of exposition, Figure 1 assumes that the consumption rate is approximately linear so the two cases can be distinguished by comparing the slopes of the consumption line.

flexible consumption and stockpiling; bathroom tissue, coffee, detergent, and paper towels exhibit stockpiling only.[2]

Collectively, these empirical findings raise the following research issue: How might the competitive pricing strategies of firms differ in categories that show only stockpiling, stockpiling plus consumption, or neither effect? We develop a game-theoretic model of price competition among firms to address the following questions:

- What should firms in retail markets do with regard to their price and promotional strategies for a product that exhibits stockpiling effects and consumption increases induced by stockpiling?
- How would the nature of price competition differ for a product that exhibits only stockpiling but no consumption effects?
- Is price competition more or less intense in markets with more pronounced consumption effects?

Given that previous research has empirically established the existence of both stockpiling and flexible consumption effects, we start with a game-theoretic model of price competition that captures consumer behavior on these dimensions. The model is based on the economics of consumer search in a market where consumers have imperfect price information. Salop and Stiglitz (1982) have shown that consumer stockpiling and imperfect price information lead to price dispersion even in a market with homogeneous consumers. The intuition is that in any given period, some firms might offer low prices (promotions) to generate additional sales from consumers who stockpile, whereas others choose to forgo these consumers to obtain the benefit of higher prices. Thus, the presence of consumer stockpiling can give rise to price promotions as an equilibrium outcome in a competitive market.

In this article, we go beyond the "pure" stockpiling phenomenon to integrate both stockpiling and consumption behaviors at the individual consumer level. That is, we allow the possibility that stockpiling can create the flexible consumption effect described previously and illustrated in Figure 1, Panel B. Consumers who decide to stockpile in response to promotion face the possibility that excess household stock might induce them to indulge in additional consumption. We establish how the equilibrium pricing strategies respond to consumers' stockpiling and flexible consumption behaviors.

The model delivers several insights pertaining to the impact of the consumption effect on price competition. The consumption effect motivates firms in retail markets to offer deeper promotions in equilibrium. The rationale is as follows: In the presence of this effect, consumers who decide to stockpile additional units anticipate that they might indulge in additional consumption. This implies that despite stockpiling, these consumers might need to reenter the market and face price uncertainty and the prospect of ending up at a high-priced store. This leads them to have a diminishing marginal utility for the additional units, forcing the firms to choose lower promotional prices in equilibrium. Thus, the consumption effect leads firms to compete by offering price promotions of greater depth. We also find that the consumption effect leads to more frequent promotions. Taken together, these findings imply that price competition is more intense in product categories (or markets) with stronger consumption effects, and this in turn leads to lower equilibrium profits and lower average market prices. In addition, we show that higher consumer inventory holding costs lead to smaller promotional depth. Higher holding costs also reduce the intensity of market competition and lead to an increase in the average market price and the equilibrium profits.

Our analysis not only shows why the consumption effect can intensify price competition but also produces specific predictions pertaining to the promotional price levels and the frequency with which firms will promote. To test the implications of the theory, we use store- and household-level scanner panel data on price, promotion, and purchase information. In particular, we test the following three model predictions:

1. The frequency of promotions is higher in categories that show the consumption effect. This implies that the proportion of stores offering a promotional price in any given period is higher in categories that show consumption effects.
2. The ratio of deal price to regular price is lower in these categories.
3. The ratio of average price to regular price is also lower.

In the empirical analysis, we use data on 4313 stockkeeping units (SKUs) from categories that show stockpiling

[2]Bell, Chiang, and Padmanabhan (1999) also report that other categories (e.g., butter, dryer softeners, ice cream, margarine, sugar) show neither effect.

and/or consumption effects. The differences between these category types (and a list of the categories used) are as follows. There are four pure stockpiling categories: bathroom tissue, coffee, detergents, and paper towels. In these categories, a promotional price leads to higher quantities purchased but not to greater consumption. There are four consumption effect categories in which a promotional price leads to both stockpiling and increased consumption: bacon, salted snacks, soft drinks, and yogurt. The no-stockpiling condition in the theoretical model occurs when consumers encounter the regular/high price and purchase sufficient quantities for current consumption only.

The remainder of the article is organized as follows: The next section provides a review of related research. We then present the model and structure of the game. We provide the theoretical results and compare and contrast the two cases of pure stockpiling and flexible consumption. We then give the data, empirical models, analysis, and results and conclude the article with a discussion and implications for further research.

BACKGROUND AND RELATED LITERATURE

There is a rich tradition of research on price promotions in both the marketing and industrial organization literature. The issues examined range from the economic and strategic rationales for price promotions to the empirical estimation of promotional effects and their managerial implications and finally to the psychological consequences of price promotions on consumer behavior and decision making. Although this article is most closely related to the first two branches of literature, it also draws on the latter. By way of summarizing the extant literature, we elaborate briefly on key findings from some of the relevant research.

Why Do Firms Offer Promotions?

The main economic rationales for why firms offer price promotions are (1) inherent demand uncertainty (Lazear 1986), (2) inventory-cost shifting (Blattberg, Eppen, and Lieberman 1981), (3) consumer heterogeneity with respect to information on market prices (Varian 1980), and (4) consumer stockpiling (Salop and Stiglitz 1982). Each rationale addresses a different aspect of consumer behavior or market dynamics.

Lazear (1986) analyzes the role of promotions when sellers are faced with demand uncertainty and provides a rationale for the clearance sales phenomenon. Promotion is viewed as a tool for capitalizing on the stochastic nature of demand. Blattberg, Eppen, and Lieberman (1981) argue that promotions help shift inventory holding costs from the retailer to the individual consumer because they encourage consumers to stockpile. The idea here is that it is more efficient for individual consumers to bear these costs separately rather than for the costs to be consolidated and borne in totality by the retailer.

Varian (1980) analyzes the idea of competitive promotions as a mixed strategy price equilibrium in response to the differences among consumers in whether they are informed about market prices. Similarly, Narasimhan (1988) and Raju, Srinivasan, and Lal (1990) examine price promotions as mixed strategies in response to differences in brand loyalty among consumers.[3] Salop and Stiglitz (1982) show that consumer stockpiling and imperfect information about prices can lead to price promotions as an equilibrium outcome in a competitive market. Price dispersion can arise even if firms and consumers are ex ante identical, because in any given period some firms can offer promotions to take advantage of consumer stockpiling. This rationale for promotions in competitive markets has received little attention and is most closely related to this article. Perhaps the smaller perceived occurrence of the primary demand effect of a promotion (and the early empirical findings that support the notion that such effects are dominated by brand switching) is part of the reason for this. Given the emerging stream of empirical and experimental studies reporting that the incidence of primary demand effects (and especially the consumption effect) is greater than previously believed, the research question of how these effects relate to retail competition and the use of price promotions becomes important.

Finally, there are models that offer an analysis of rational consumer behavior but treat firm pricing and promotional decisions as exogenous to the consumer problem. For example, Assuncao and Meyer (1993) examine the normative implications of price promotions for temporal buying and consumption behavior. Consumers anticipate prices through a first-order Markov process, and as a result, the rational shopping policy dictates that consumption rates should increase with household inventories, increase with holding costs, and decrease with temporal discounting. A related study by Ho, Tang, and Bell (1998) extends this idea and shows that a rational cost-minimizing consumer will increase consumption when he or she is faced with increased price variation. Krishna (1992) shows that the volume of product bought on deal decreases with promotion frequency, and this finding is empirically supported (see, e.g., Foekens, Leeflang, and Wittink 1999; Helsen and Schmittlein 1992). Although these articles do not address equilibrium pricing behavior, they nevertheless show a link between promotions and consumer behavior in the form of stockpiling and increased consumption.

How to Estimate Price Promotion Effects?

Gupta (1988) calibrates a model on scanner panel data and finds that the predominant effect of a price promotion is on brand switching and therefore secondary demand. This view is also supported by Chiang (1991) and Chintagunta (1993). With the advent of larger databases with more product categories, recent research (e.g., Bell, Chiang, and Padmanabhan 1999) has shown that primary demand effects of price promotions can be significant and that their importance varies considerably across categories. Similarly, van Heerde, Leeflang, and Wittink (2000), in their analysis of pre- and postpromotion dips, find some evidence of primary demand expansion. Nijs and colleagues (2001, p. 11) find that "price promotions significantly expand category demand in 58% of the cases over a dust-settling period that

[3]There are also several articles that investigate promotions in the context of a distribution channel. These include Lal and Villas-Boas (1998) and Jeuland, Kumar, and Rajiv (2001), which examine the impact of trade promotions on retail behavior.

lasts, on average, 10 weeks." In the long run, however, these effects appear to die out.

Thus, although some evidence for the primary demand effect has started to emerge, few empirical studies explicitly address the behavioral mechanisms that lead to primary demand expansion. That is, the econometric formulation typically assumes that consumption rates vary across consumers but are exogenous and fixed (i.e., consumers behave according to the pattern given in Figure 1, Panel A¾ even if additional units are purchased on promotion, consumers simply consume at the same rate and increase the elapsed time to the next purchase).

Ailawadi and Neslin (1998) address this issue using a reduced form approach that links the inventory and consumption component of the consumer utility function. In their model, the consumption rate can vary over time and in accordance with the level of inventory on hand. They find significant increases in promotion-induced consumption in yogurt and smaller effects for ketchup. Sun (2001) provides more evidence for flexible consumption in her study of the tuna category.

These findings, based on econometric models that are calibrated on secondary data, are convergent with results from prior experimental research. For example, Folkes, Martin, and Gupta (1993) and Wansink and Desphande (1994) demonstrate that consumption may be driven up by increases in household inventories. In summary, there is a good deal of empirical support for the idea that promotional tactics employed by firms stimulate stockpiling behavior and lead to increases in consumption.

MODEL

To relate promotions to stockpiling and consumption, we need to develop a model that is able to capture consumption dynamics at the individual consumer level. The framework must allow a consumer the option to not only buy the good for present consumption but also store it for future consumption. In addition, it must capture the possibility that stockpiling might result in additional consumption.[4] In this section, we first describe the assumptions of consumer and firm behavior and then outline the solution to the game.

Model Assumptions

Consumer market and consumer decisions. The market consists of T consumers who are homogeneous in their per-unit valuation, u, of the product. Consumers have a two-period consumption/planning horizon (this two-period assumption enables us to model intertemporal purchase dynamics). More precisely, the two periods should be thought of as pertaining to two possible purchase occasions a consumer faces within a planning cycle (we continue to use the term "periods" for expositional ease). On any given purchase occasion, consumers do not know the price charged by a particular store but know a priori only the distribution of prices, f(p), that they will encounter. Given price uncertainty, we assume that consumers randomly select a store. With a planning cycle of two periods, we assume that consumers choose from two alternative shopping strategies:

1. *No-stockpiling strategy*: If the price encountered is not low enough, a consumer might decide to buy only one unit for current consumption in Period 1 and then reenter the market to buy a second unit at the price encountered in Period 2. In this case, the consumer incurs a transaction cost, c, of entering the market again in the second period. We assume that if the consumer only bought one unit in Period 1, he or she will have a maximum demand of two units over the entire planning cycle (one unit in each period). In other words, if only one unit is bought in Period 1, there is no stockpiling and therefore no opportunity for additional consumption.
2. *Stockpiling strategy*: Alternatively, if the price encountered in Period 1 is attractive enough, the consumer can choose a stockpiling strategy, in which we assume that the consumer is able to stockpile an additional unit of the product.[5] Thus, a consumer's stockpiling strategy will involve his or her purchasing two units in Period 1: one unit for current consumption and a second unit that can be stored for future consumption. The consumer stockpiles the additional unit to potentially avoid the need to reenter the market in Period 2. As a result, the consumer avoids the possibility of encountering a high price in the second period and saves on the transaction cost, c, of reentering the market. However, the consumer will incur a holding/storage cost, h, if he or she ends up storing the second unit, where h captures the cost of physical storage as well as spoilage costs. The physical storage costs can be due to the bulkiness of the product or the extent of valuable freezer/refrigerator space used. Spoilage costs are relevant for perishable categories.

Introducing the consumption effect created by stockpiling. Allowing consumers to choose to buy two units in Period 1 and store one unit (at the holding cost h) for future consumption captures the possibility of consumer stockpiling. But the consumer model must also capture the consumption effect. This effect is relevant only for consumers who adopt the stockpiling strategy and purchase two units in Period 1. Upon stockpiling, these consumers face one of the following two states: With probability θ (0, 1) a consumer who has stockpiled consumes both the units (and enjoys consumption utility 2u) in the first period.[6] In this case, the consumer will reenter the market despite having bought the two units but will not incur the cost of storage. In the second state, which happens with the probability $(1 - \theta)$, a consumer who purchases two units in Period 1 consumes only one unit of the product (and enjoys a utility of u) in Period 1 and stockpiles the second for future consumption at the storage cost of h. In this event, the consumer does not need to reenter the market. Finally, note that in Period 2, the planning cycle ends and the consumer has utility for only up to one unit of consumption in this period.

The state-dependent utility described can be interpreted as follows: At the time of purchase in Period 1, consumers are uncertain as to whether they will indulge in additional

[4] As noted previously, Salop and Stiglitz (1982) examine the effect of consumer stockpiling on price promotions. The model developed here generalizes their approach to incorporate the flexible consumption effect that stockpiling can create.

[5] This implicitly means that the holding costs for stocking more than one unit are prohibitively high. Relaxing this assumption and allowing consumers to stockpile two units does not change the results of the article pertaining to the effects of consumption on price competition, but it complicates the algebra.

[6] An alternative interpretation of θ is that it is the proportion of consumers who end up consuming both units in Period 1, assuming that they stockpiled in Period 1.

consumption, but they have knowledge of the value of θ for the product. Consider the example of a consumer's purchase of soft drinks versus that of detergent. At the time of purchase, the consumer knows that, conditional on availability, he or she is more likely to consume additional quantities of soft drinks than detergent. However, the actual additional consumption of soft drinks is uncertain. For example, additional consumption of soft drinks for an individual consumer could be triggered by idiosyncratic complementary events, such as a favorite movie or an unexpected visit by a friend. Therefore, θ can be interpreted as a product category characteristic that represents the ex ante belief consumers have about the likelihood of additional consumption in that category, whereas the realization of θ can be perceived as related to idiosyncratic consumption events. It should be noted that in our empirical test, we explicitly recognize consumer heterogeneity when linking the theory to the data. In particular, we estimate flexible, hierarchical random effects models, which allow the possibility that specific SKUs and consumers within a category exhibit idiosyncratic effects that are distributed around a category-level mean.

Consumers are rational, and as such the rationality requirement implies that in Period 1 they make decisions by maximizing their full intertemporal expected utility. In other words, in deciding whether to stockpile in Period 1, consumers anticipate that stockpiling might result in additional consumption, thereby forcing them to reenter the market and face price uncertainty. Consumers maximize their intertemporal expected utility over the two periods on the basis of this anticipation. Finally, the consumer's Period 2 decision is simply to decide whether to buy one unit given the price encountered.

Firms. The retail market consists of n ex ante identical retailers. This assumption eliminates the possibility that individual retailer characteristics drive any price dispersion in the market. As a result, any observed price dispersion will be purely due to the strategic price choices of the retailers in equilibrium. Retailers have constant marginal cost of production for the good, which can be normalized to zero without any loss of generality.

Market Equilibrium

Each firm in the retail market chooses a pricing strategy to maximize profits for the planning cycle, given the prices of the other firms. We establish the symmetric mixed strategy equilibrium of the model. Furthermore, in equilibrium, each consumer searches optimally in response to $f^*(p)$ (the equilibrium price distribution he or she faces) and buys only if he or she receives nonnegative surplus. The mixed strategy equilibrium has the advantage of reflecting Hi–Lo promotional competition that characterizes the retail grocery industry we examine. Moreover, it implies that consumers who reenter the market after stocking out face price uncertainty.

Note that we can also consider a pure strategy equilibrium of a model in which consumers search with replacement (which implies price uncertainty for the consumers who reenter in period 2).[7] This model is fully equivalent to the model presented here, and in equilibrium, firms choose to be either high- or low-priced firms for the planning cycle. The results and insights of such a model are analogous to the mixed strategy results reported in this article, in which the cross-sectional distribution across retailers is similar to the within-retailer price distribution in the symmetric mixed strategy.[8]

THEORETICAL RESULTS

We are interested in understanding the impact of stockpiling and the consumption effect on retailers' strategic decisions, so it is necessary to establish the exact conditions under which stockpiling occurs. To do this, we must first calculate the threshold price that determines whether the consumer stockpiles at all. Subsequently, we describe the case in which consumers stockpile but stockpiling does not lead to an increase in consumption ($\theta = 0$), and then we describe the main case, in which stockpiling leads to additional consumption ($\theta > 0$).

General Condition

The first task is to identify the equilibrium price support in the mixed strategy profile. In Period 1, consumers must choose between buying one unit of the good for current consumption only or two units for current and future consumption. Consequently, the decision to buy two units versus one unit in Period 1 involves a trade-off between buy-and-hold versus shopping in the future period. This implies a "threshold" price, $\hat{p}$, which makes the consumer indifferent between the two choices.

Consider that consumers encounter a store charging a price $\hat{p}$ and decide to buy two units in Period 1. With probability θ, these consumers will use both units in Period 1 and will need to reenter the market in Period 2 for an additional unit. Consumers who reenter the market will face price uncertainty while purchasing the additional unit. Denoting the expectation of the price distribution $f(p)$ as $\bar{p}$, their total (expected) utility from this eventuality is $3u - 2\hat{p} - \bar{p} - c$. Alternatively, with probability $(1 - \theta)$, they will not consume the second unit in Period 1, and their total utility in this scenario is $2u - 2\hat{p} - h$. Consequently, their expected utility from stockpiling in the first period is $E(U_s) = \theta(3u - 2\hat{p} - \bar{p} - c) + (1 - \theta)(2u - 2\hat{p} - h)$.

In contrast, consider the utility obtained by the consumer who decides not to stockpile in Period 1. Such a consumer will buy one unit of the good at $\hat{p}$ and will reenter the market in Period 2 for the second unit. In doing so, the consumer incurs the transaction cost c for reentering and faces the uncertain price distribution $f(p)$. We have the expected utility of such a consumer as $E(U_{ns}) = u - \hat{p} + u - \bar{p} - c$.

A comparison of the utility for the two cases (stockpiling versus no stockpiling) generates the threshold price $\hat{p}$, which is computed as follows: The consumer will stockpile if $E(U_s) \geq E(U_{ns})$, which implies the threshold price

$$p \leq \hat{p} = \theta u + (1 - \theta)(\bar{p} + c - h). \quad (1)$$

In addition, we need to check that in equilibrium $E(U_s) \geq 0$; otherwise, buying two units and stockpiling will not be a feasible strategy for the consumer.

The next step in the analysis is to show that the price support in a symmetric equilibrium will involve at most two

[7]This can also be interpreted as the consumer market being memoryless. Therefore, consumers who reenter the market in Period 2 randomly select a store.

[8]We thank two anonymous reviewers for comments on this issue.

prices (denoted henceforth as p_l and p_h, with $p_l < p_h$). Because consumers shop randomly, it will be the case that all stores have the same number of consumers. Suppose there were two putative "low" prices p_l and p_m below $\hat{p}$ that a store was charging in equilibrium (where $p_m > p_l$). The p_l and p_m firms would have the same number of consumers who buy two units and stockpile; however, the profits for the store that charges p_m will be greater. Recall that the requirement of a mixed strategy equilibrium is that firms make their competitors indifferent between each of their pure strategies. This means that in equilibrium, the profits from choosing p_m and p_l cannot be different. Therefore, in the equilibrium chosen there will be only one price p_l at or below $\hat{p}$, and consumers who encounter this price in Period 1 will purchase two units of the product. Similarly, there cannot be two prices above $\hat{p}$, because the demand that a firm obtains from choosing either price would be the same. Thus, there will be a single price p_h above $\hat{p}$, and consumers who encounter this price will buy one unit of the product.

Given the preceding, an equilibrium must have the property that the low price is exactly $p_l = \hat{p}$. If $p_l > \hat{p}$, consumers would buy only one unit at the store, but because $p_h > p_l$, the p_l store will make less profit. In contrast, when a store charges $p_l < \hat{p}$, consumers at this store will buy for two periods. By raising the price slightly, such a store would not lose any customers but would increase profits. Therefore, the equilibrium low price must be $p_l^* = \hat{p}$. Similarly, the equilibrium high price charged by a store will be $p_h^* = u$.[9]

To solve for the symmetric equilibrium of this model, let λ denote the probability that a firm is charging the high price for the planning cycle (recall that the planning cycle consists of the two purchase occasion periods). Note that this implies that each firm in equilibrium will charge a price according to the probability λ for both purchase occasion periods. This helps reflect the reality in actual markets that retailers cannot discriminate among consumers on the basis of whether consumers are reentering the market because they stocked out. In other words, as consumers enter the store, there is no way for the retailer to distinguish between consumers who are reentering and those who are not.

As consumers shop randomly, each store receives T/n consumers in Period 1. Thus, λT consumers will enter a store that charges a high price and therefore buy only one unit. These consumers will need to reenter the market to purchase a second unit of the good. The remaining $(1 - \lambda)T$ consumers will encounter low-priced stores and buy two units in Period 1. Of these, a proportion, λ, will consume both units, reenter the market in Period 2, and buy an additional unit. The demand for a store that charges a high price S_h is

$$S_h = \frac{T}{n} + \frac{\lambda T}{n} + \frac{\theta(1-\lambda)T}{n}, \tag{2}$$

and the demand when charging the low price, S_l, is given by

$$S_l = \frac{2T}{n} + \frac{\lambda T}{n} + \frac{\theta(1-\lambda)T}{n}. \tag{3}$$

The first term in the demand functions comes from consumers who shop randomly in Period 1. The second term comes from consumers who faced a high-priced store in Period 1 and who therefore reenter the market. Finally, the third term comes from reentering consumers who faced a low-priced store in Period 1 and bought two units but consumed both the units.[10] From the definition of the mixed strategy, each firm will choose a distribution of the two possible prices to make other firms indifferent between their strategies, implying that in equilibrium,

$$p_h S_h = p_l S_l. \tag{4}$$

From now on, we set c = 0 for ease of exposition. The case of c > 0 yields qualitatively similar results (details of the c > 0 case are available from the authors). Substituting $p_h = u$, we have

$$[1 + \lambda + \theta(1 - \lambda)]u = [2 + \lambda + \theta(1 - \lambda)]p_l. \tag{5}$$

Recall that the average price is $\bar{p} = \lambda u + (1 - \lambda)p_l$. Finally, we have $p_l = \hat{p} = \theta u + (1 - \theta)(\bar{p} - h)$. Using these pricing identities and Equation 5, we have the following equilibrium low price (assuming $E(U_s) \geq 0$, which we show to be true in equilibrium):

$$p_l = \frac{u + h(1-\theta)}{2}. \tag{6}$$

The equilibrium can now be derived and is as shown in Proposition 1.

Proposition 1: If c = 0, a necessary and sufficient condition for the existence of an equilibrium is

$$\frac{\theta u}{2 - \theta(1+\theta)} < h < \frac{u}{3(1-\theta)}. \tag{7}$$

The equilibrium price profile consists of two prices with

$$p_l^* = \frac{u + h(1-\theta)}{2} \text{ and} \tag{8}$$

$$p_h^* = u, \tag{9}$$

and firms charge the high price with probability

$$\lambda^* = \frac{2h - \theta(u+h) - h\theta^2}{(1-\theta)[u - h(1-\theta)]}. \tag{10}$$

[9]Note that in models such as Varian's (1980), Narasimhan's (1988), and Raju, Srinivasan, and Lal's (1990), firms offer promotional prices in a mixed strategy equilibrium because of the presence of uninformed/switching consumers (or because of consumer heterogeneity in loyalty). The argument in this article is that promotional prices can occur even if there are no switching consumers. Even without any switching consumers, firms can offer promotions as long as consumers engage in stockpiling and flexible consumption.

[10]This model and the demand structure it generates can be viewed as analogous to the following infinite horizon overlapping generations model: Firms live in perpetuity, but consumers live for two purchase occasion periods and are then replenished. In each period, there is a mix of old consumers (who reenter because they stocked out) and new consumers, and firms cannot distinguish between these consumers. The new consumers buy either one or two units depending on whether they encounter a high- or low-priced store (similar to the first term in Equations 2 and 3), and the old consumers buy only one unit (similar to the second and third terms in Equations 2 and 3). For discount factors close to one, the insights derived from such an overlapping generations model would be similar to those presented here. We thank an anonymous reviewer for comments.

Proof: The necessary and sufficient conditions for a mixed strategy equilibrium with promotional pricing to exist are that $0 < \lambda^* < 1$ and that $p_l^* = \hat{p}$. Solving Equation 5 and the pricing identity $p_l = \hat{p} = \theta u + (1 - \theta)(\bar{p} - h)$ simultaneously yields the equilibrium values of p_l^* and λ^* shown previously. Some algebra reveals that the condition $0 < \lambda^*$ is equivalent to $\theta u/[2 - \theta(1 + \theta)] < h$, whereas $\lambda^* < 1$ implies $h < u/[3(1 - \theta)]$. Furthermore, $EU_s > 0$ after substitution of the equilibrium values implies that $h < u[\theta + \lambda(1 - \theta)]/(1 - \theta)[2\lambda + \theta(1 - \lambda)]$. Some algebraic manipulation shows that $u[\theta + \lambda(1 - \theta)]/(1 - \theta)[2\lambda + \theta(1 - \lambda)] > u/[3(1 - \theta)]$. This implies that the mixed strategy equilibrium exists when the condition shown in Equation 7 is satisfied.

Proposition 1 specifies the strategies of the firms and the condition for the existence of an equilibrium with price promotions. We now compare situations in which consumers stockpile with those in which they stockpile and consume more of the product. We first examine the special case of pure stockpiling ($\theta = 0$), because this provides us with a benchmark from which to evaluate the incremental effect of flexible consumption ($\theta > 0$).

Pure Stockpiling ($\theta = 0$)

Bell, Chiang, and Padmanabhan (1999) show that categories such as bathroom tissue, coffee, detergents, and paper towels exhibit significant stockpiling in response to a price promotion; however, the stockpiling in these categories is not accompanied by any change in the consumption rate of consumers. This case of pure stockpiling is a special case of our model where $\theta = 0$, which implies that consumers who stockpile two units in Period 1 do not indulge in excess consumption as a result of this stockpiling (and therefore do not reenter the market in Period 2). The two units are consumed one apiece in each of two periods.

This case of $\theta = 0$ is the analogue of Salop and Stiglitz's (1982) analysis. Thus, the within-firm price distribution in the mixed equilibrium here is analogous to the cross-sectional price distribution in Salop and Stiglitz's work. Note that without consumption effects, the only consumers in the market in Period 2 are those who encountered a high price in Period 1 and purchased a single unit. Consequently, the total sales when a store is charging the high price in the equilibrium is $S_h = [(1 + \lambda)T]/n$, and total sales when a store is charging the low price is $S_l = [(2 + \lambda)T]/n$.

Corollary 1: An equilibrium with promotional pricing for this case exists only if $h < u/3$. The equilibrium prices are

(11) $$p_l^* = \frac{u + h}{2} \text{ and}$$

(12) $$p_h^* = u,$$

and firms charge the high price with probability

(13) $$\lambda^* = \frac{2h}{(u - h)}.$$

The presence of consumer stockpiling can induce firms to offer price promotions. The two-price equilibrium occurs only if the holding costs are sufficiently low. If the cost of holding the additional inventory is large, consumers are less disposed to buy for storage, and as a result the low-priced stores will need to offer deeply discounted prices to motivate consumers to buy the additional unit. This makes promotional prices unprofitable and implies that the equilibrium frequency of charging the high price (λ^*) increases with consumer holding costs.

Table 1
EQUILIBRIUM RESULTS AND CONDITIONS

	General Results	
Demand	$S_l^* = \frac{[1 + \lambda + \theta(1 - \lambda)]T}{n}$	$S_h^* = \frac{[2 + \lambda + \theta(1 - \lambda)]T}{n}$
Prices	$p_l^* = \frac{u + h(1 - \theta)}{2}$	$p_h^* = u$
Equilibrium frequency	$\lambda^* = \frac{2h - \theta(u + h) - h\theta^2}{(1 - \theta)[u - h(1 - \theta)]}$	
Profits	$\pi_l^* = \pi_h^* = \frac{u[u - h(1 + \theta)]T}{[u - h(1 - \theta)]n}$	
Equilibrium condition ($\theta > 0$)	$\frac{\theta u}{2 - \theta(1 + \theta)} < h < \frac{u}{3(1 - \theta)}$	

Stockpiling and Consumption ($\theta > 0$)

Given the analysis so far, we begin with the question, How does the consumption effect relate to price competition? Proposition 1 indicates that, as in the pure stockpiling case, the holding costs of the consumers must be sufficiently small for promotions to occur in equilibrium (i.e., $h < u/[3(1 - \theta)]$. However, unlike the pure stockpiling case, we find that holding costs cannot be too small. Specifically, we require that $h > \theta u/[2 - \theta(1 + \theta)]$. The reason for this is as follows: Low holding costs imply that firms do not need to cut prices significantly to motivate consumers to buy the additional unit for storage. Furthermore, in the presence of the flexible consumption effect, firms also receive the benefit of increased demand in Period 2 (from consumers who stockpiled but indulged in excess consumption, using up both units in Period 1). This makes the low-priced position so attractive that all firms have the incentive to promote, with probability 1. We investigate the following aspects of the equilibrium: (1) the equilibrium prices and the frequency with which firms promote, (2) equilibrium profits, and (3) the impact of holding costs on the proportion of stores charging a low price. Table 1 provides the details of the equilibrium.

Equilibrium frequency of charging the high price. We begin with λ^*, the equilibrium frequency of charging a high price $p_h^* = u$. It can be shown that λ^* decreases with consumption effects (i.e., $\partial\lambda/\partial\theta < 0$). The presence of flexible consumption makes it more attractive for any given store to charge the promotional low price more frequently. This low price, in turn, encourages the consumers to stockpile. Having stockpiled, consumers have a probability θ of consuming the stockpiled unit and thereby increasing the rate of consumption. The effect of higher consumption (i.e., a larger value of θ) translates into higher overall market demand at the promotional price. Firms respond in equilibrium by increasing the frequency of promotions.

Equilibrium prices. The high or regular price ($p_h^* = u$) is the same in the theoretical model in all the scenarios; however, this is not true for the promotional price. Consider the situation of pure stockpiling. The choice of a promotional price, p_l^*, needs to be low enough to induce a consumer who visits the store to purchase two units and incur the holding cost for storage. Proposition 1 indicates that the equilibrium promotional price decreases with the consumption effect ($\partial p_l/\partial\theta < 0$ from Equation 8). In the presence of the consumption effect, consumers who decide to stockpile additional units anticipate that they might indulge in additional consumption. This implies that despite stockpiling, these consumers might need to reenter the market and face price uncertainty and the prospect of ending up at a high-priced store. As such, this leads them to have a diminishing marginal utility for the additional units, which forces the firms that decide to adopt a low price to choose lower promotional prices in equilibrium. Furthermore (as discussed previously), the probability and/or frequency with which firms promote also increases. Thus, the expected number of firms that offer the promotional price also increases with the consumption effect. All this also implies that the average market price, $\bar{p}$, goes down with the consumption effect. In other words, firms compete more intensely in product categories with a greater consumption effect.

Equilibrium profits. The symmetric equilibrium profits for each firm are

$$\pi^* = \frac{u[u + h(1 - \theta)]T}{[u - h(1 - \theta)]n}. \tag{14}$$

From Equation 14, it is straightforward to note that the equilibrium profits go down as the consumption effect becomes stronger. This also confirms the claim that flexible consumption intensifies price competition. This is interesting because the consumption effect produces some benefit (i.e., primary demand can be higher); however, firms must pay for this through more intense competition.

Holding costs. As h increases, (1) the frequency with which stores promote decreases, (2) the average price ($\bar{p}$) increases, and (3) the promotional price p_l^*) and firm profitability increase. When the consumer faces an increased cost of holding inventory, the firm needs to provide a bigger drop in price to induce consumers to stock up in the first period. The net effect of this is to reduce the attractiveness of promotions, and stores adopt the promotional price less frequently. This is evident from the result that λ, the probability of stores charging the high price $p_h^* = u$, increases as h increases.

In certain product categories, there might be a relationship between the holding costs and the likelihood of additional consumption. In particular, consumers might have a greater propensity to use up inventory in categories with higher holding costs. Note that the effects of h and θ on the promoted price, frequency, and profits, when they are independent, work counter to each other. In categories in which higher holding costs are positively related to θ, we expect price competition to be less intense.

EMPIRICAL ANALYSIS

We test key implications of the theoretical analysis. We focus on the equilibrium frequency distribution, or the likelihood of charging the high price (captured by λ), and the equilibrium pricing behavior. In particular, we consider the ratio of deal and average prices to regular prices in stockpiling and consumption effect categories.[11]

Data

To test the theory, we require categories that exhibit either flexible consumption effects or pure stockpiling. To aid in category selection, we refer back to previous work. Ailawadi and Neslin (1998) identify yogurt as a consumption effect category, and Bell, Chiang, and Padmanabhan (1999, p. 517, Exhibit 3) support this by reporting three other categories that show the consumption effect: bacon, potato chips, and soft drinks. The exhibit also lists four categories that show pure stockpiling effects: bathroom tissue, coffee, detergent, and paper towels.

We obtained two years of price and promotion information for each SKU in all eight categories listed in Bell, Chiang, and Padmanabhan's (1999) study. The market we examine covers five supermarkets located in close proximity, and the total number of SKUs in each consumption effect category is 98, 1189, 1305, and 341 for bacon, salted snacks, soft drinks, and yogurt, respectively. For stockpiling categories, we have 117, 713, 437, and 113 SKUs for bathroom tissue, coffee, detergents, and paper towels, respectively. Thus, with two years of price data for each of the 4313 SKUs in the data set, we have a potentially large number of observations for use in our analysis. However, not all SKUs are stocked by all stores for all 104 weeks.[12]

Before presenting the hypothesis tests and results, we first describe the testing strategy. Our goal is to develop an empirical model that not only reflects the basic elements of the theory but also takes into account important aspects of the data (e.g., unobserved nature of key constructs, heterogeneity).

Empirical Testing Approach

The parameter θ is a key construct in our theory. As noted previously, θ is a product category characteristic that represents an ex ante belief consumers have about the likelihood of additional consumption, whereas the particular realization of θ depends on idiosyncratic consumption events. The value of θ affects how firms set prices in equilibrium. This begs the question, How can we obtain reasonable estimates of θ from the data?

To generate category-specific estimates of θ that are consistent with the theory and the structure of the data, we return to the analysis presented by Bell, Chiang, and Padmanabhan (1999, p. 517, Exhibit 3). We reproduce the exhibit as Table 2, which contains information on the average quantities purchased ($\overline{Q}$), average elapsed time between purchases ($\overline{IP}$), and average rate of consumption ($\overline{Q/IP}$) in each of the eight categories. These measures were computed from panel data purchases and therefore reflect underlying buying patterns within the categories.

[11]An advantage of using these ratios is that they are unitless measures of price dispersion and natural candidates for across-category and across-store comparisons.

[12]The relatively large number of SKUs for two of the food categories¾salted snacks and soft drinks¾reflects the large degree of variety in flavors and sizes in those categories. Private labels are only carried by a single store, and many SKUs are dropped and added in the course of the two-year period. There is a fairly high stockout rate; approximately 7% of SKUs are stocked out of a given store in a given week.

Table 2

CATEGORY-SPECIFIC ESTIMATES OF θ

	Promotional Purchases			*Nonpromotional Purchases*			*Consumption Effects*[a]	
Category	$\overline{Q}$	$\overline{IP}$	$\overline{Q/IP}$	$\overline{Q}$	$\overline{IP}$	$\overline{Q/IP}$	$r_P > r_R$?	$\hat{\theta}$
Stockpiling only								
Bathroom tissue (roll)	6.76	27.65	.30	1.83	22.38	.28	No	0
Coffee (ounce)	35.71	58.81	.80	24.76	43.75	1.19	No	0
Detergent (ounce)	132.42	77.32	2.37	82.70	48.12	2.71	No	0
Paper towels (roll)	2.00	45.12	.08	1.40	34.77	.08	No	0
Increased consumption								
Bacon (ounce)	27.23	46.57	.92	19.92	47.49	.68	Yes	.26
Salted snacks (ounce)	10.25	44.63	.64	7.37	41.75	.41	Yes	.36
Soft drinks (ounce)	210.32	45.28	9.18	135.86	45.86	6.27	Yes	.32
Yogurt (ounce)	33.12	37.69	1.43	25.12	38.32	1.10	Yes	.23

[a]For regular purchases, $r_R = \overline{Q/IP}$, from column 6, and for promotional purchases, $r_P = \overline{Q/IP}$, from column 3; $\hat{\theta} = 1 - r_R/r_P$. When $\hat{\theta}$ is not statistically different from zero, we set it equal to zero.

Notes: The first six columns of this table have been reproduced from Bell, Chiang, and Padmanabhan (1999, p. 517, Exhibit 3).

Although Bell, Chiang, and Padmanabhan's (1999) exhibit does not provide a direct estimate of θ, it provides sufficient information from which to generate an estimate. In all eight product categories, price promotions lead to statistically significant increases in quantities purchased. In the case of the stockpiling-only categories, we find an accompanying increase in the elapsed time between purchases, which implies no significant increase in the rate of consumption. This is exactly the phenomenon presented in Figure 1, Panel A. For the increased consumption categories, the increased purchase quantity is accompanied by no significant increase in elapsed time, which suggests an increase in the slope of the consumption line (see Figure 1, Panel B).

An estimate of θ can be derived from the relationship between rates of consumption following purchases on promotion and purchases at regular prices. Let $\theta = 1 - r_R/r_P$, where r_R and r_P are the rates of consumption given purchases at regular and promoted prices, respectively. This means that $\theta = 0$ when a promotion-induced increase in purchase quantity is not accompanied by any increase in consumption (this definition is consistent with the assumptions of our theory). When the two consumption rates are different, $\theta > 0$.[13] In the consumption effect categories, the values of θ are strictly greater than zero (see Table 2), which is again consistent with our theory. In estimating the effect of θ on dependent measures that capture retail pricing behavior, we also control for idiosyncratic variation within categories (exact specifications follow).

Hypothesis Tests and Results

We examine the impact of the consumption effect (θ) on (1) the equilibrium frequency of charging the high price, (2) the equilibrium ratio of deal to regular prices, and (3) the equilibrium ratio of average to regular prices.

Equilibrium frequency of charging the high price. A key analytical result from the model involves the equilibrium frequency of charging the high price.

H_1: λ, the equilibrium frequency of charging the high price, should be lower for consumption categories than for stockpiling categories.

A simple test of H_1 involves separating the categories into consumption effect and stockpiling-only groups, as is shown in Table 2. Our objective is to show that the average frequency of charging the high price is lower for the consumption effect group. In the theoretical model, λ is the equilibrium frequency that retailers choose the high price. In the empirical model, we can interpret this as the probability that retailers choose the regular price at each point in time. This implies that at each time period t, for item i in category j, the dependent measure λ_{ijt} can be interpreted as the proportion of retailers in the market that charge the regular price. It is a time-dependent proxy frequency of a given store choosing the high price.[14]

To compute the dependent measure λ_{ijt}, we need to define a high price in a particular case. In the theoretical model, the ex ante identical retailers all charge the same high price equal to the consumer reservation price, u. In the data, store prices are also somewhat dictated by micro conditions and chain-level decisions. To appropriately control for such factors in testing H_1, we proceed as follows: For all SKUs i = 1, …, I; categories j = 1, …, J; and stores k = 1, …, K, we compute $mp_{ik}(j) = \max\{p_{ik1}(j), \ldots, p_{ikT}(j)\}$, where t indexes weeks t = 1, …, 104. The SKU price distributions within stores are overwhelmingly bimodal (regular price and deal price), so the maximum price charged by store k for a given SKU i in a particular category j, $mp_{ik}(j)$, is simply the regular price and is readily determined from the data. From these values, the within-group mean frequencies of charging the high price are .502 and .581 for the consumption effect and stockpiling-only groups, respectively, and the mean for the consumption effect category is significantly lower (F-value = 23.46, $p < .0001$).

[13]When the statistical test cannot reject the hypothesis that $r_R = r_P$, we assume that they are indeed equal. This means that for the category under consideration, $\theta = 1 - r_R/r_P = 0$.

[14]Note that our definition of maximum prices with respect to store-specific benchmarks enables us to use these maximum prices to create an across-store measure. That said, there are alternative ways to create a suitable dependent variable (we thank the anonymous reviewers for these suggestions). One natural candidate is an SKU-specific empirical frequency of the proportion of times an SKU is at regular price, or in our notation, λ_{ijt}. Although this measure also decreases in θ (details are available from the authors), we retain the formulation in Equation 15 that follows in the text. Analysis revealed nonstationarity in the time series of prices, and we wanted to determine whether our key theoretical results were robust to this phenomenon.

Although the simple test supports H_1, it would also be useful to conduct a test that better connects the theory to the particular characteristics of the data. In theory, θ is a parameter that captures differences across categories, and though it cannot be observed directly in the data, it can nevertheless be inferred at the category level (see Table 2). In the real data, there are two additional sources of variation in θ not captured explicitly by the theory: (1) variation across items, i = 1, …, I within a category j, and (2) variation over time t = 1, …, T for a particular item i in category j. Therefore, the dependent variable of interest in the statistical analysis is not simply θ but its empirical proxy, θ_{ijt}. That is, the data can be structured in a natural hierarchy: Within a particular category j, there are several items, i = 1, …, I, and for each item i, there are multiple observations over time, t = 1, …, T. In conducting the tests, we specify a hierarchical random effects model that takes this into account:

$$(15) \quad \lambda_{ijt} = \beta_{0i} + \beta_1 \times \theta_j + \beta_2 \times MS_{ij} + \beta_3 \times \theta_j \times MS_{ij} + \beta_4 \times TI_{ijt} + v_{ijt},$$

where

$\beta_{0i} = \gamma_0 + u_{0i}$, with $u_{0i} \sim N(0, \sigma^2)$;
θ_j = category-level estimate from Table 2;
MS_{ij} = average market share of SKU i in category j; and
TI_{ijt} = a time index for SKU i, category j, week t.

The item-level intercept, β_{0i}, enters as a random effect to reflect that within a category, pricing decisions are likely to differ across items for several reasons that are unobserved by the analyst (e.g., retailer knowledge of buyer preferences for particular items, frequency of trade deals from manufacturers). The parameter β_1 captures the effect of category-level differences in θ_j, and in accordance with H_1, we expect $\beta_1 < 0$.[15] The variable MS_{ij} controls for observable heterogeneity across SKUs. It is likely that, all else being equal, retailers prefer to promote high–market share brands (e.g., retailers would rather offer deals on Coke than on Dr. Pepper). Therefore, we expect $\beta_2 < 0$.

The interaction effect $\theta_j \times MS_{ij}$ is included in the model to account for the possibility that even within categories that display consumption effects, there is likely to be variation across items. In particular, large–market share brands in consumption-sensitive categories should have increased levels of promotion ($\beta_3 < 0$) after the category-level main effects are accounted for.[16] We also control for time-varying changes in promotional intensity through the inclusion of the item- and category-specific time-varying covariate, TI_{ijt}. As noted previously, exploratory analysis revealed that over the two years of the data set, on average all SKUs had increased levels of promotion. Therefore, we expect that the equilibrium frequency with which retailers charge the high price should decline over time so that β_4 should be negative.

To preserve the interpretation of the intercept γ_0, all covariates are appropriately centered on grand means or category-level means (Singer 1998); in the case of the time index, TI_{ijt}, the first week is indexed at zero. In estimating the model, we tested alternative specifications of the error term v_{ijt} (these are reported in Table 3). As might be expected for these types of pricing data, the weekly prices are serially dependent, so that the final model assumes that v_{ijt} follows a first-order autoregressive process. The estimated parameters of Equation 15 are given in Table 3, and a total of 34,141 observations are used in estimation.[17]

The category-level effect of flexible consumption is negative and significant, as predicted by the theory (β_1 = –.1130, $p < .0385$). This result holds even given the limited variation in θ_j and the provision for considerable flexibility through other aspects of the model. All else being equal, the equilibrium frequency of charging the high price declines for large-share brands (β_2 = –.2992, $p < .0001$) and over time (β_4 = –.1690, $p < .0001$). The interaction effect captured by β_3 is also negative and significant as expected. The total marginal effect ($\partial\lambda_{ijt}/\partial\theta_j$) is also negative given that $MS_{ij} > 0$ by definition.

The estimate for the variance component on the intercept suggests that there is random variation across items, even after we control for other effects (σ^2 = .0579, Pr Z = .0001). As expected, the data also show some serial dependence. In summary, we have strong support for H_1: The equilibrium

[15]We also estimated the model with β_1 entering as a random effect; however, given the limited number of categories, the inclusion of other relevant covariates, and random effects on the item-level intercept, there is no statistically significant variation in this parameter ($p < .148$).

[16]Note that the effects remain fully valid even without the interaction term and the estimates change little. In this case, the parameter estimates and *p*-values are (1) intercept γ_0 = .6724, $p < .0001$; (2) consumption effect β_1 = –.1404, p = .0284; (3) market share β_2 = –.4310, $p < .0001$; and (4) time β_3 = –.1694, $p < .0001$.

[17]For the consumption effect categories, there are 696, 5595, 11089, and 6616 observations for bacon, salted snacks, soft drinks, and yogurt, respectively. For the stockpiling categories, there are 948, 4913, 3444, and 840 observations for bathroom tissue, coffee, detergents, and paper towels, respectively. The correlation between these numbers and the number of SKUs reported in the "Data" section is .841.

Table 3

H_1: PARAMETER ESTIMATES FROM THE HIERARCHICALMODEL

Model Parameters	*Estimate*	*Standard Error*	*Pr < t*
Intercept (γ_0)	.6716	.0106	<.0001
Consumption effect θ_j (β_1)	–.1130	.0639	<.0385
Market share MS_{ij} (β_2)	–.2992	.0281	<.0001
Interaction effect $\theta_j \times MS_{ij}$ (β_3)	–1.3422	.1745	<.0001
Time TI_{ijt} (β_4)	–.1690	.0103	<.0001
Covariance Parameters	*Estimate*	*Standard Error*	*Pr Z*
Variance component (intercept, σ^2)	.0579	.0034	<.0001
Variance component (AR[1])	.7275	.0040	<.0001
Variance component (residual)	.0529	.0008	<.0001

frequency of charging the high price is lower for categories with consumption effects.

Promotional depth (ratio of deal to regular prices). H_2 complements H_1 by focusing on the relative depth of promotion for particular items in stockpiling and consumption effect categories.

H_2: The ratio of deal to regular prices decreases in categories that exhibit consumption effects.

As in H_1, we begin with a simple test. We identify deal and regular prices for each category and set R_D equal to the deal price divided by the regular price. The measure R_D is unitless and therefore suitable for cross-category comparisons. It is computed by analyzing the entire time series of prices $p_{ikt}(j)$, $t = 1, \ldots, T_{ik}(j)$ for each item i in store k in category j. With the regular price in hand (see the procedure for H_1), we determine each occasion in which a price discount is offered and calculate $R_D(ijt)$. Using these values of $R_D(ijt)$, we compute the overall average ratio of deal to regular prices for the stockpiling and consumption effect groups and find that the average is lower in consumption effect categories. The respective values are .642 and .741 (F-value = 144.22, $p < .0001$), so deals are significantly deeper in categories that have consumption effects.

Although this simple test supports H_2 for reasons given previously, we again use the random effects formulation of Equation 15 but this time with $R_D(ijt)$ as the dependent variable. As in Equation 15, the intercept enters as a random effect such that $\beta_{0i} = \gamma_0 + u_{0i}$, with $u_{0i} \sim N(0, \sigma^2)$. The other covariates are defined as before, and store-level fixed effects are added to the model to control for store-specific differences in promotional depth.[18] We expect to find $\beta_1 < 0$ as consumption effects lead to deeper promotions. Following the rationales given in the test of H_1, we expect greater deals on larger–market share brands ($\beta_2 < 0$) and a negative interaction effect ($\beta_3 < 0$); β_4 accounts for changes in deal depth over time.

[18]We also estimated the model on a store-by-store basis for both H_2 and H_3 and obtained qualitatively identical results. The consumption effect parameters, β_1, under H_2 are –.2499, –.3266, –.2808, –.3276, and –.3208 for Stores 1 through 5, respectively, with $p < .0001$ in all cases. For H_3, we have –.0707, –.1112, –.1363, –.0879, and –.0482, with $p < .0001$ in all cases. We thank an anonymous reviewer for suggesting this analysis.

As is shown in Table 4, the effect of θ on the ratio of deal to regular prices is negative and significant ($\beta_1 = -.0746$, $p < .0009$). The estimates for β_2 and β_3 are not critical to the theory; however, they have plausible magnitudes and expected signs. The negative estimate for β_4 suggests that, on average, promotions gradually became deeper over time (though the store-specific results show that this is true only at the first three stores). In summary, H_2 is supported: Promotional depth is greater in consumption effect categories.

Promotional depth (ratio of average to regular prices). H_3 complements H_2 by offering an alternative test of the promotional depth result.

H_3: The ratio of average to regular prices decreases in categories that exhibit the consumption effect.

We compute the ratios for the average prices to regular prices, $R_A(ij)$, and reestimate the model replacing $R_D(ij)$ with $R_A(ij)$.[19] As is shown in Table 5, the pattern of results is similar to that for H_2. β_1 has the expected negative sign and is significantly different from zero: H_3 is also supported by the data.

CONCLUSION

This article provides insights into price competition among retailers in categories in which consumers can be motivated to buy for storage and future consumption. These are categories that display the stockpiling effect (and exhibit temporary increases in primary demand) and/or experience long-term effects through increased consumption. Although several experimental studies and empirical research efforts have documented these effects, there exists no theory research to explain how retail price competition might be governed by the stockpiling and consumption behaviors of consumers.

We present an equilibrium analysis of price competition that traces the relationship between consumer tendencies to increase consumption and the behavior of firms that are pricing under such conditions. In this setting, the presence of consumer storage and flexible consumption effects is suf-

[19]The ratio for the average price to regular price pertains to the entire time series, and therefore there is no t subscript. The time covariate is excluded from the model.

Table 4

H_2: CONSUMPTION EFFECTS AND PROMOTION DEPTH (DEAL TO REGULAR PRICES)

Model Parameters	*Estimate*	*Standard Error*	*Pr < t*
Intercept (γ_0)	.9485	.0036	<.0001
Consumption effect θ_j (β_1)	–.0746	.0223	<.0009
Market share MS_{ij} (β_2)	–.1280	.0189	<.0001
Interaction effect $\theta_j \times MS_{ij}$ (β_3)	–.4107	.1313	<.0018
Time TI_{ijt} (β_4)	–.0266	.0027	<.0001

Table 5

H_3: CONSUMPTION EFFECTS AND PROMOTION DEPTH (AVERAGE TO REGULAR PRICES)

Model Parameters	*Estimate*	*Standard Error*	*Pr < t*
Intercept (γ_0)	.9326	.0012	<.0001
Consumption effect θ_j (β_1)	–.0855	.0230	<.0001
Market share MS_{ij} (β_2)	–.0972	.0962	<.3126
Interaction effect $\theta_j \times MS_{ij}$ (β_3)	–3.6309	.7114	<.0001

ficient for firms to offer price promotions. This general finding highlights a rationale for price promotions that is missing in the theoretical literature. As noted previously, this could result from many researchers either explicitly or implicitly assuming that price promotions do not increase consumption. Given recent empirical and experimental evidence for the consumption effect, it is useful to have a theoretical analysis of promotions that incorporates this behavioral phenomenon.

The theoretical model establishes that the frequency of choosing the high price decreases with consumption effects. We identify this new rationale for promotions in a situation in which both stores and consumers are ex ante identical in all respects. Our work also represents a departure from most studies in this area that either develop theory or undertake empirical analysis, but not both. In summary, we provide the following new insights:

1. The equilibrium frequency of charging a high price decreases in product categories subject to consumption effects.
2. The depth of promotion is also influenced by the presence of consumption effects, such that (a) the ratio of deal to regular prices decreases and (b) the ratio of average to regular prices decreases.
3. Firm profits go down in categories that exhibit consumption effects. Although retailers can benefit from the increase in primary demand, they must also compete more intensely for it.
4. The previous results are moderated by increases in consumer inventory holding costs, such that (a) the equilibrium frequency of charging the high price increases, (b) deal prices increase, and (c) firm profits increase.

We tested the theoretical predictions for promotion frequency and promotion depth using price and promotion data from eight product categories. The empirical tests offer strong support for the implications of the theory. In categories subject to the consumption effect, we find that (1) the equilibrium frequency of charging the high price decreases, (2) promotional depth increases as the ratio of deal to regular prices is lower, and (3) the ratio of average to regular prices is also lower. Retailers'pricing behavior suggests that they are cognizant of whether products are likely to exhibit flexible consumption patterns. In this sense, the findings support the notion of a market equilibrium in which consumers respond to retailer prices differentially according to category characteristics, and retailers endogenously set prices with this in mind (Villas-Boas and Winer 1999).

In conclusion, we offer a new rationale for price promotions: firm response to stockpiling and flexible consumption behavior of consumers. The effects of price promotions do not pertain exclusively to brand switching. Instead, price promotions can connect to fundamental consumer dynamics such as time shifting of purchase quantities and flexible consumption behavior. Our theory highlights a mechanism for this connection and provides support for its validity by carefully examining empirically observed patterns of price promotions in actual markets.

REFERENCES

Ailawadi, K. and Scott A. Neslin (1998), "The Effect of Promotion on Consumption: Buying More and Consuming It Faster," *Journal of Marketing Research*, 35 (August), 390–98.

Assuncao, Joao and Robert J. Meyer (1993), "The Rational Effect of Price Promotions on Sales and Consumption," *Management Science*, 39 (5), 517–35.

Bell, David R., Jeongwen Chiang, and V. Padmanabhan (1999), "The Decomposition of Promotional Response: An Empirical Generalization," *Marketing Science*, 18 (4), 504–26.

Blattberg, Robert C., Gary D. Eppen, and Joshua Lieberman (1981) "A Theoretical and Empirical Evaluation of Price Deals for Consumer Nondurables," *Journal of Marketing*, 45 (Winter), 116–29.

¾¾¾ and Scott A. Neslin (1990), *Sales Promotion Concepts, Methods and Strategies*. Englewood Cliffs, NJ: Prentice Hall.

Chiang, Jeongwen (1991), "A Simultaneous Approach to the Whether, What, and How Much to Buy Questions," *Marketing Science*, 10 (Fall), 297–315.

Chintagunta, Pradeep (1993), "Investigating Purchase Incidence, Brand Choice and Purchase Quantity Decisions of Households," *Marketing Science*, 12 (Spring), 184–208.

Foekens, Eijte W., Peter S.H. Leeflang, and Dick R. Wittink (1999), "Varying Parameter Models to Accommodate Dynamic Promotion Effects," *Journal of Econometrics*, 89 (March/April), 249–68.

Folkes, Valerie S., Ingrid M. Martin, and Kamal Gupta (1993), "When to Say When: Effects of Supply on Usage," *Journal of Consumer Research*, 20 (December), 467–77.

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

Helsen, Kristian and David C. Schmittlein (1992), "Some Characterizations of Stockpiling Behavior Under Uncertainty," *Marketing Letters*, 3 (1), 5–17.

Ho, Teck-Hua, Christopher S. Tang, and David R. Bell (1998), "Rational Shopping Behavior and the Option Value of Variable Pricing," *Management Science*, 44 (12), S145–S160.

Jeuland, Abel P., Nanda Kumar, and Surendra Rajiv (2001), "Effectiveness of Trade Promotions: Analyzing the Determinants of Retail Pass-Through," *Marketing Science*, 20 (Fall), 382–404.

Krishna, Aradhna (1992), "The Normative Impact of Consumer Price Expectations for Multiple Brands on Consumer Purchase Behavior," *Marketing Science*, 11 (Summer), 266–87.

Lal, Rajiv and J. Miguel Villas-Boas (1998), "Price Promotions and Trade Deals with Multi-product Retailers," *Management Science*, 44 (July), 935–49.

Lazear, Edward P. (1986), "Retail Pricing and Clearance Sales," *American Economic Review*, 76 (March), 14–32.

Moore, William L. and Russell S. Winer (1978), "An Experiment to Determine the Effects of Package Size on Consumption," working paper, Graduate School of Business, Columbia University.

Narasimhan, Chakravarthi (1988), "Competitive Promotional Strategies," *The Journal of Business*, 61 (October), 427–50.

Nijs, Vincent R., Marnik G. DeKimpe, Jan-Benedict E.M. Steenkamp, and Dominique M. Hanssens (2001), "The Category-Demand Effects of Price Promotions," *Marketing Science*, 20 (Winter), 1–22.

Raju, Jagmohan S., V. Srinivasan, and Rajiv Lal (1990), "The Effects of Brand Loyalty on Competitive Price Promotional Strategies," *Management Science*, 36 (March), 276–304.

Salop, S. and J.E. Stiglitz (1982), "The Theory of Sales: A Simple Model of Equilibrium Price Dispersion with Identical Agents," *American Economic Review*, 72 (December), 1121–30.

Singer, Judith D. (1998), "Using SAS PROC MIXED to Fit Multilevel Models, Hierarchical Models and Individual Growth Models," *Journal of Educational and Behavioral Statistics*, 24 (4), 323–55.

Sun, Baohong (2001), "Promotion Effects on Category Sales with Endogenized Consumption and Promotion Uncertainty," working paper, Graduate School of Industrial Administration, Carnegie Mellon University.

van Heerde, Harald, Peter S.H. Leeflang, and Dick R. Wittink (2000), "The Estimation of Pre- and Postpromotion Dips with Store-Level Scanner Data," *Journal of Marketing Research*, 37 (August) 383–95.

Varian, Hal R. (1980), "A Model of Sales," *American Economic Review*, 70 (September) 651–59.

Villas-Boas, J. Miguel and Russell S. Winer (1999), "Endogeneity in Brand Choice Models," *Management Science*, 45 (10), 324–38.

Wansink, Brian and Rohit Deshpande (1994), "Out of Sight, Out of Mind: Pantry Stockpiling and Brand Usage Frequency," *Marketing Letters*, 5 (1), 91–100.