American Economic Review 2008, 98:4, 1245–1268
http://www.aeaweb.org/articles.php?doi=10.1257/aer.98.4.1245

Competition and Price Variation when Consumers Are Loss Averse

By Paul Heidhues and Botond Kőszegi*

We modify the Salop (1979) model of price competition with differentiated products by assuming that consumers are loss averse relative to a reference point given by their recent expectations about the purchase. Consumers' sensitivity to losses in money increases the price responsiveness of demand—and hence the intensity of competition—at higher relative to lower market prices, reducing or eliminating price variation both within and between products. When firms face common stochastic costs, in any symmetric equilibrium the markup is strictly decreasing in cost. Even when firms face different cost distributions, we identify conditions under which a focal-price equilibrium (where firms always charge the same "focal" price) exists, and conditions under which any equilibrium is focal. (*JEL* D11, D43, D81, L13)

Menu costs, tacit collusion, search costs, kinked demand curves, and many other theories have been invoked to explain a widespread empirical fact: that prices in imperfectly competitive industries often do not change when costs or demand change.[1] Existing theories, however, cannot convincingly account for another way in which prices vary surprisingly little: they are often identical across differentiated products. As is part of industrial organization folklore starting from R. L. Hall and C. J. Hitch (1939) and Paul M. Sweezy (1939), and is documented by the Competition Commission of the United Kingdom (1994), Jonathan Beck (2004), and Heidhues and Kőszegi (2007), many nonidentical competitors charge identical, "focal," prices for their differentiated products. In addition, as documented for instance by Robert S. McMillan (2004) and Liran Einav and Barak Y. Orbach (2005), and familiar to anyone who has bought clothes, books, or movie tickets, many retailers selling multiple products with different cost and/or demand characteristics charge the same "uniform" price for them.

* Heidhues: Department of Economics, Bonn Graduate School of Economics, Adenauerallee 24–26, 53113 Bonn, Germany (e-mail: Heidhues@uni-bonn.de); Kőszegi: University of California, Berkeley, 508-1 Evans Hall #3880, Berkeley, CA 94720-3880 (e-mail: botond@econ.berkeley.edu). We are very grateful to Jeremy Bulow, Marc Dudey, Ulrike Malmendier, Jörg Oechssler, Matthew Rabin, Philipp Schmidt-Dengler, and three anonymous referees for great suggestions. We also thank seminar participants at University of California, Berkeley, University of Bonn, Central European University, University of Cologne, University of Edinburgh, ESMT Berlin, University of Frankfurt, University of Freiburg, University of Innsbruck, University of Iowa, London School of Economics, LMU Munich, University of California, San Diego, Wissenschaftszentrum Berlin, University of Zürich, as well as participants at the CEPR's Competition Policy Conference in Brussels (2005), EARIE (2006), Edinburgh Workshop on Behavioral and Experimental Economics (2007), the European Symposium in Psychology and Economics (2006), and the Stanford Institute in Theoretical Economics (2006). Kőszegi thanks the National Science Foundation for financial support under award 0648659. Heidhues gratefully acknowledges financial support from the Deutsche Forschungsgemeinschaft through SFB/TR-15.

[1] Indirect evidence for this fact—often referred to as price stickiness—is provided by Anil K. Kashyap (1995), Margaret E. Slade (1999), and Judith A. Chevalier, Kashyap, and Peter E. Rossi (2003), who document in various retail industries that regular prices typically do not change for months at a time. There is also direct evidence that marginal-cost changes are sometimes fully absorbed by retailers (Competition Commission of the United Kingdom 1994, 150, sect. 7.41). And evidence on countercyclical markups reviewed in Julio J. Rotemberg and Michael Woodford (1999) suggests that even when prices do adjust to circumstances, they move less than marginal costs.

To explain these tendencies toward reduced price variation, we develop a model of price competition between profit-maximizing firms selling to loss-averse consumers. Because consumers are especially averse to paying a price when it exceeds their expectation of the purchase price, the price responsiveness of demand—and hence the intensity of competition—is greater at higher than at lower market prices, reducing or eliminating price variation. Unlike in most previous theories, this logic applies both to different possible prices of the same product and to prices of different products, so that our theory not only explains the unresponsiveness of prices to changing circumstances, but also often predicts focal and uniform pricing as the *unique* outcome even for asymmetric firms and products. And because a change in the responsiveness of demand affects competition more when the value of an extra consumer is high, we predict that these tendencies are stronger in more concentrated industries.

Section I presents our model and illustrates our solution concepts and some key results using a two-firm example. Building on Steven C. Salop (1979), a consumer's "taste" is drawn uniformly from the circumference of a circle, and she is looking to buy exactly one of n products located equidistant from each other on the same circle. Her utility from or "satisfaction with" a product is decreasing in the product's distance from her taste, and she also suffers additive disutility from paying the product's price. But (applying Kőszegi and Matthew Rabin 2006, 2007) we posit that in addition to this *intrinsic utility*, a consumer derives *gain-loss utility* from comparing outcomes in money and product satisfaction to her lagged rational expectations about those outcomes, with losses being more painful than equal-sized gains are pleasant. For example, if she had been expecting to spend \$14.99 on a Britney Spears CD—her favorite music—she experiences a sensation of loss if she buys that CD for \$18.98, and also if she instead buys a—less agreeable—Madness CD for \$14.99.[2] And if she expected to pay either \$14.99 or \$19.99 for something, paying \$18.98 for it generates a mixture of two feelings, a loss of \$3.99 and a gain of \$1.01, with the weight on the loss equal to the probability with which she expected to pay \$14.99.

The firms, none of which owns two neighboring products, are standard: they face uncertain privately observed costs of production and simultaneously set prices to maximize expected profits given other firms' behavior and consumer expectations. We begin the analysis in Section II by showing that the necessary and sufficient condition for a focal-price equilibrium—an equilibrium in which all firms always charge the same focal price p^*—to exist is that any two cost realizations of any two firms are within a given constant. This condition allows for, say, one firm to have higher costs than another in all states of the world. If consumers had expected to pay p^* with probability one, they assess buying at a price greater than p^* as a loss in money and buying at a price lower than p^* merely as a gain in money, so that demand is more responsive to unilateral price increases than decreases from p^*. Due to this asymmetry, for a range of cost levels p^* is the optimal price to charge.

We next establish two properties of focal-price equilibria. First, a focal-price equilibrium is more likely to exist in more concentrated industries. Since the profits from a consumer are higher in this environment, the asymmetric demand responsiveness at p^* creates a greater difference in marginal profits from price increases versus price decreases from p^*, and hence yields a greater range of costs for which p^* is the optimal price to charge. Second, loss aversion increases prices. Since a consumer is more sensitive to a loss from getting surprisingly low product satisfaction than to a gain from paying surprisingly little, attracting her from a competitor is difficult, decreasing competition.

In Section III, we derive a sufficient condition that guarantees that *any* equilibrium is a focal-price equilibrium, even though it does not require firms to have the same cost distribution. The

[2] Actual prices were taken on September 4, 2005, from www.amazon.com; \$14.99 was the retail price of both CDs (and numerous others), while \$18.98 was a typical list price.

key is to argue that a firm sets a deterministic price in any equilibrium; then, if the supports of firms' cost distributions have even a single common point, firms cannot charge different deterministic prices, so any equilibrium is a focal-price one. If the consumer expected a firm's prices to be stochastic, the sense of loss from comparing the realized price to lower possible ones would make her demand more responsive at higher than at lower prices in the firm's anticipated distribution. Then, if the firm's costs do not vary much, in contradiction to equilibrium it could increase profits either by decreasing high prices (attracting a lot of extra demand) or by increasing low prices (not losing much demand).

In Section IV, we characterize all equilibria with common stochastic marginal costs and symmetric pricing strategies. This leads to a tractable model for studying price variation when conditions for focal pricing are not necessarily met, and for analyzing firms' responses to industry-wide cost shocks. As above, if a consumer had expected stochastic prices, her demand is more responsive at higher than at lower prices within the anticipated price distribution. Hence, competition is fiercer at higher prices, leading to markups that strictly decrease in cost. Using the empirical observation that costs are strongly procyclical, this means that markups are countercyclical. In addition, in some regions of cost it may be that competition at higher prices is tougher to an extent that is inconsistent with firms raising their price in response to cost increases at all. In such regions, the price must be constant in cost.

In Section V, we argue that our results are robust to a number of modifications of our model, including dynamics, demand asymmetries and shocks, heterogeneity in consumer preferences, and the endogenous determination of the number of firms. In Section VI, we discuss theories of pricing most closely related to our model. By dint of predicting equal and sticky prices even in a one-shot setting, our theory cautions against the common interpretation of these patterns as signs of collusion. In fact, we can go further: because playing an equilibrium with *equal* in addition to *sticky* prices does not help firms in detecting each other's deviations, models of collusion do not provide a compelling reason for ex ante asymmetric firms to set equal prices, whereas our theory does. In the same vein, other explanations of price stickiness are not intended to, and largely do not, predict focal and uniform pricing. For instance, menu costs may explain sticky pricing and can perhaps contribute to uniform pricing, but they cannot address equal pricing across firms. And if unexpected price increases trigger costly search by consumers but unexpected price decreases do not, sticky pricing can result, but there is no reason for differentiated products to have the same price.

All proofs are relegated to the Web Appendix, available at http://www.aeaweb.org/articles.php?doi=10.1257/aer.98.4.1245.

I. Setup and Illustration

This section introduces our theory and explains the solution concepts and some key results through simple examples and heuristic calculations. We incorporate consumer loss aversion into the model of Salop (1979) using a disciplined approach introduced by Kőszegi and Rabin (2006): we base the reference-dependent "gain-loss utility" on classical "intrinsic utility" taken straight from Salop (1979), and fully endogenize the reference point as lagged rational expectations. If there is no loss aversion, our theory reduces to Salop's.

A. *Reference-Dependent Utility*

A mass one of consumers have tastes $\chi \in [0, 1]$ distributed uniformly on the circumference of a circle with perimeter one. There are $n \geq 2$ products denoted $y_1, \ldots, y_n$ on the same circle, equidistant from each other. A consumer can buy at most one product, and to avoid unenlightening extra notation, we assume her utility from not consuming is negative infinity, so that she always does buy a product.[3] Letting $d(\chi, y)$ denote the distance of χ and y on the circle, the intrinsic utility of consumer χ from buying product y at price p is $v - t\,d(\chi, y) - p$, where $k_1 = v - t\,d(\chi, y)$ is her intrinsic utility from, or "satisfaction" with, the good, and $k_2 = -p$ is her intrinsic utility from paying its price. Like previous authors, we interpret χ as the consumer's "ideal variety," and $t\,d(\chi, y)$ as her disutility from consuming a product y different from her ideal. The constant t is a measure of the (intrinsic) differentiation between products.

For a riskless consumption outcome $k = (k_1, k_2)$ and riskless reference point $r = (r_1, r_2)$ defined over product satisfaction and money, total utility $u(k|r)$ is composed of two additive terms: intrinsic utility introduced above, and reference-dependent "gain-loss utility" equal to $\mu(k_1 - r_1) + \mu(k_2 - r_2)$. To capture loss aversion, we assume that μ is two-piece linear with a slope of 1 for gains and a slope of $\lambda > 1$ for losses.[4] This specification incorporates three key assumptions. First, the consumer assesses gains and losses in the two dimensions, satisfaction and money, separately. Hence, she evaluates a good that costs more but is closer to her ideal than the reference point as a loss in money and a gain in satisfaction—and not, for example, as a single gain or loss depending on total intrinsic utility relative to the reference point. This is consistent with much experimental evidence commonly interpreted in terms of loss aversion.[5] Second, while money is on a different psychological dimension from any of the n products, our model also says that the n products are on the *same* dimension. This assumption reflects our impression that goods that compete most strongly with each other are typically hedonically substitutable; indeed, that is partly why they compete. Third, since the gain-loss utility function μ is the same in both dimensions, the consumer's sense of gain or loss is directly related to the intrinsic value of the changes in question. While we find this assumption psychologically realistic, we point out in Section V that it is only necessary for one of our results, Proposition 2.

Since we assume below that the reference point is expectations, we extend the utility function above to allow for the reference point to be a probability measure Γ over $\mathbb{R}^2$:

$$U(k|\Gamma) = \int_r u(k|r)\,d\Gamma(r). \tag{1}$$

In evaluating k, the consumer compares it to each possibility in the reference lottery. For example, if she had been expecting to pay either \$15 or \$20 for her favorite CD, paying \$17 feels like a loss of \$2 relative the possibility of paying \$15, and like a gain of \$3 relative to the possibility of paying \$20. In addition, the higher the probability with which she expected to pay \$15, the more important is the loss in the overall experience.

We assume that consumers' prior on χ is identical to the population distribution, $U[0, 1]$. Since it gives rise to the same distribution of satisfaction from each product, an equivalent model is one in which consumers know their ideal variety, but are uncertain about the positioning of

[3] Our results would be identical if consumers had an option of not buying, but v below was sufficiently high (or costs and product differentiation sufficiently low) that no consumer took advantage of this option in equilibrium. And in Section V, we argue that our qualitative results on reduced price variation would survive even if consumers made a relevant decision of whether to buy.

[4] In order not to clutter our formulas with extra notation, we do not introduce a weight on gain-loss utility relative to intrinsic utility. This does not qualitatively affect any of our results, and we have confirmed that our calibration at the end of Section II also remains unchanged.

[5] Specifically, it is key to explaining the endowment effect and other observed regularities in riskless trades. The common and intuitive interpretation of the endowment effect—that randomly assigned "owners" of an object value it more highly than "nonowners"—is that owners construe giving up an object as a painful loss that counts more than the money they receive in exchange, so they attach a high monetary value to the object. But if gains and losses were defined over the value of an entire transaction, owners would not be more sensitive to giving up the object than to receiving money in exchange, so no endowment effect would ensue.

products.[6] A situation where consumers have a very good idea about their ideal variety as well as the products offered corresponds to a narrow or even degenerate prior distribution on χ, and yields a different model. In Section V, we argue that our results in Sections II and IV carry over to this case unchanged, and that reasonable specifications also yield our results in Section III.

B. *Concepts and Results: Illustration*

While Section IA defined how the consumer's utility depends on her reference point, we must also specify what the reference point is and how firms behave when selling to loss-averse consumers. In this section we illustrate some of our definitions and results in a two-firm example and without the full formal detail of later sections.

To both motivate our model of reference-point determination and explain a key result, suppose the two firms in the market are expected to set deterministic prices p_1 and $p_2 > p_1$ for products 1 and 2. In the face of these prices, what is the consumer's reference point for evaluating her purchase? We posit that it is her *lagged rational expectations* about outcomes. But since these depend on her own behavior, our assumption requires elaboration. To illustrate, suppose that the consumer had planned to buy the cheaper product if her taste is within distance $\alpha \in (1/4, 1/2)$ of firm 1, and to buy the expensive product otherwise, as shown on the left-hand side of Figure 1. This plan induces an expected purchase-price distribution F with mass 2α on p_1 (the probability that χ falls within α of $y_1 = 0$) and mass $1 - 2\alpha$ on p_2, as well as an expected distribution of the purchased product's distance from ideal that is shown on the right-hand side of Figure 1. Hence, the consumer's reference point—and so her utility at the time of purchase—is affected by the plans she had formed earlier: if α is higher, she expected to pay less and get lower product satisfaction with higher probability, which makes paying a high price more painful and getting a less satisfying product less painful. To close the model, we follow Kőszegi and Rabin (2006) in requiring a consistency condition called personal equilibrium: the consumer can only make plans she knows she will follow through. In the current setting, this means that given the expectations above, when she has taste α the consumer must be indifferent between purchasing from firms 1 and 2.

While our model of consumer behavior is new, firm behavior is more or less standard: each firm maximizes expected profits given other firms' behavior and the consumer's reference point.[7] As in a standard model, a major factor in determining equilibrium is the consumer's reaction to price changes. To understand a key part of this reaction, we focus on the money dimension. If the consumer above (unexpectedly) pays a price p satisfying $p_1 < p < p_2$, her reference-dependent utility in money is

$$-p - \lambda(2\alpha)(p - p_1) + (1 - 2\alpha)(p_2 - p).$$

The first term is intrinsic utility. The second term represents a sense of loss from comparing p to the lower expected purchase price p_1—a loss of $p - p_1$ weighted by the probability with which she expected to pay p_1, 2α. And the third term represents a gain from comparing p to the higher expected purchase price p_2—a gain of $p_2 - p$ weighted by the probability with which she expected to pay p_2, $1 - 2\alpha$. Hence, a small price increase decreases the consumer's utility in the money dimension by $1 + \lambda(2\alpha) + (1 - 2\alpha)$. More generally, at any price p that is not a

[6] More precisely, our model is identical to one in which each consumer knows her ideal variety, and product 1's location is drawn from a uniform distribution on the circle, with the n products still equidistant from each other.

[7] We assume profit maximization to capture our impression that firms display reference-dependent preferences far less than consumers do, and to isolate the effect of *consumer* loss aversion on market outcomes.

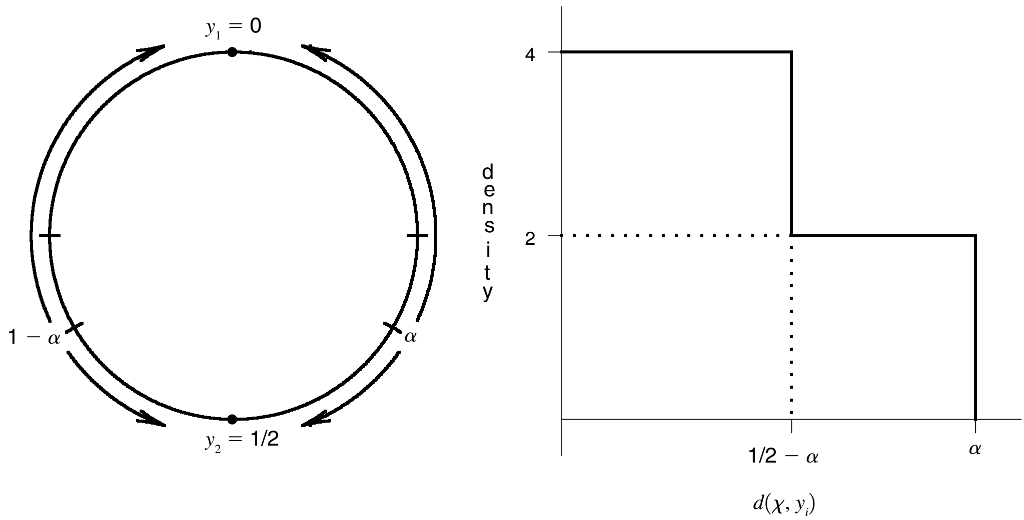

Figure 1. Illustration of a Consumer's Plans with Two Firms Charging Deterministic Prices

Notes: The figure on the left illustrates the consumer's strategy: if her taste is within distance $\alpha \in (1/4, 1/2)$ of firm 1, she buys the cheaper product 1; otherwise she buys the more expensive product 2. The figure on the right illustrates the density of the expected distribution of the purchased product's distance from ideal that is induced by this plan. If the consumer's taste is very close to a product, she buys that product, so the density is 4 for small distances. For larger distances, the consumer is willing to buy a product that is far from her taste only if it is the cheaper product, so the density shrinks to 2. Given her plans, the consumer does not expect to purchase a product that is farther than α from her ideal variety.

mass point of the expected purchase-price distribution F, the utility impact of a marginal price change is equal to $1 + \lambda F(p) + 1(1 - F(p))$. The intuition is simple. Paying p is experienced as a loss relative to lower prices in the expected purchase-price distribution, and as a gain relative to higher prices in that distribution. Due to this "comparison effect," a change in p is counted as a change in loss with weight $F(p)$—the probability with which the consumer expected to pay lower prices—and as a change in gain with weight $1 - F(p)$—the probability with which she expected to pay higher prices.

Based on the considerations above, at any price p_i that is not a mass point of F, the partial derivative of firm i's demand with respect to its price p_i is

$$-\frac{[1 + \lambda F(p_i) + (1 - F(p_i))]}{t\,z}, \tag{2}$$

where z reflects the consumer's gain-loss utility in product satisfaction. When p_i is a mass point of F, demand is continuous and left and right differentiable with the derivatives given by the left and right limits of expression (2). Although z depends on the consumer's reference point and the prices set by the firm's neighbors, for the purposes of this section we assume heuristically that it is an exogenous constant.

Now we can illustrate results about possible equilibria in the market. Proposition 3 says (more generally and precisely) that so long as there exist realizations of marginal costs c_1 and c_2 with $c_1 \geq c_2$, an equilibrium in which the firms charge deterministic prices p_1 and $p_2 > p_1$ does not exist. With these prices, when $c_1 \geq c_2$, firm 2 has a higher markup than firm 1, so it benefits more from one extra consumer than firm 1 suffers from one less consumer. Furthermore, since firm 2 has lower infra-marginal demand than firm 1, its infra-marginal losses from lowering its price are lower than firm 1's inframarginal gains from raising its price by the same amount. And because by expression (2) the responsiveness of demand at prices just below p_2 is the same as the responsiveness just above p_1, either firm 1 wants to raise its price or firm 2 wants to lower its price.

In contrast, Proposition 1 says (generally and precisely) that even if the firms have different marginal-cost distributions, an equilibrium in which they charge the *same* deterministic price p^* often exists. If consumers expect to pay p^* with certainty, expression (2) implies that the price responsiveness of a firm's demand when it unilaterally raises its price is $-(1+\lambda)/(tz)$, while the responsiveness when it lowers its price is only $-2/(tz)$. Intuitively, a price decrease of a given amount expands demand less than a price increase of the same amount reduces demand because consumers are not as attracted by a gain in money as they dislike a loss in money. Since the effect of these price changes on profits from inframarginal consumers is symmetric, for a range of cost levels neither deviation can increase profits.

To conclude this section, we illustrate the reasoning behind our trickiest result, Proposition 4, which provides conditions for all firms to set a deterministic price. Combined with conditions above ruling out different deterministic prices for different firms, this leads to conditions under which any equilibrium is a focal-price equilibrium. The essence of the argument can be seen most simply by assuming that in equilibrium firm 1's cost and price are continuously distributed on $(\underline{c}_1, \overline{c}_1)$ and $(\underline{p}_1, \overline{p}_1)$, respectively, and firm 2's price is also continuously distributed. We show a condition under which a contradiction results. If the marginal costs of the two firms are sufficiently similar and densely distributed, firm 1's expected demand is about one-half. Using this and expression (2), the fact that at cost $\overline{c}_1$ firm 1 does not want to set a price lower than $\overline{p}_1$ implies $\overline{p}_1 - \overline{c}_1 \lessapprox (tz)/[4 + 2(\lambda - 1)F(\overline{p}_1)]$. Similarly, that at cost $\underline{c}_1$ firm 1 does not want to set a price above $\underline{p}_1$ implies $\underline{p}_1 - \underline{c}_1 \gtrapprox (tz)/[4 + 2(\lambda - 1)F(\underline{p}_1)]$. Subtracting the latter inequality from the former one, and using that the consumer must have expected to buy from firm 1 with probability of about one-half, so that $F(\overline{p}_1) - F(\underline{p}_1) \gtrapprox 1/2$, we get

$$\overline{c}_1 - \underline{c}_1 \gtrapprox \frac{tz}{2}\left[\frac{(\lambda-1)(F(\overline{p}_1) - F(\underline{p}_1))}{[2+(\lambda-1)\;F(\underline{p}_1)][2+(\lambda-1)F(\overline{p}_1)]}\right] + \overline{p}_1 - \underline{p}_1 \gtrapprox \frac{tz}{2}\frac{(\lambda-1)\frac{1}{2}}{(1+\lambda)^2} = \frac{tz}{4}\frac{\lambda-1}{(\lambda+1)^2}.$$

Hence, if $\overline{c}_1 - \underline{c}_1$ is small, firm 1's incentives are incompatible with the purported equilibrium above. Intuitively, if firm i chooses a stochastic price, then—no matter how close its highest and lowest prices—expectations-based loss aversion dictates an amount by which the consumer is more price responsive at the firm's high price than at its low price. This, in turn, implies an amount by which the markup at the high price must be lower than at the low price. But if the firm's highest and lowest costs do not differ by that amount, such a situation is impossible.

C. *Personal Equilibrium and Market Equilibrium*

This section formally specifies consumer behavior and defines market equilibrium. We begin by defining the concept of personal equilibrium motivated above with two firms setting deterministic prices for more firms and arbitrary price distributions. Suppose that the consumer has a prior $H \in \Delta(\mathbb{R}^n_+)$ on the nonnegative price vectors she might face. Her decision of which good to buy is made after observing the realized χ and the realized price vector, and is described by the strategy $\sigma : [0, 1] \times \mathbb{R}^n_+ \to \{1, \ldots, n\}$. As emphasized above, her reference point for evaluating outcomes is her lagged rational expectations about outcomes. This is the distribution $\Gamma_{\sigma,H}$ induced by σ, H, and her uniform prior over χ, over vectors (k_1, k_2) of product satisfaction and expenditure. To deal with the resulting interdependence between behavior (σ) and expectations ($\Gamma_{\sigma,H}$), personal equilibrium (Kőszegi and Rabin 2006) requires the behavior generating expectations to be optimal given the expectations:

DEFINITION 1: *σ is a* personal equilibrium *for the price distribution H if*

$$\sigma(\chi, p) \in \operatorname{argmax}_{i \in \{1,\ldots,n\}} U(v - t \cdot d(\chi, y_i), p_i \mid \Gamma_{\sigma,H}) \text{ for all } \chi \in [0, 1] \text{ and } p \in \mathbb{R}^n.$$

We now integrate consumer behavior into a notion of market equilibrium. The timing of our full market model is illustrated in Figure 2. Consumers first form the expectations regarding consumption outcomes that later determine their reference point. Next, firms observe their cost realizations and simultaneously set prices. Finally, each consumer observes her ideal variety and the realized market prices, and purchases a good.

For expositional simplicity, we assume that the n available products are produced by n different firms, with firm i producing good i, y_i. In Section V, we argue that as long as each product is owned by exactly one firm and no firm owns neighboring products, our results on how products are priced extend unchanged to situations where some or all firms produce multiple products. Hence, results on focal pricing below also imply uniform pricing for multiproduct firms.

Firms' costs are jointly distributed according to Θ on the set $\Pi_{i=1}^{n} [\underline{c}_i, \overline{c}_i]$, where $[\underline{c}_i, \overline{c}_i]$ is the smallest closed interval containing the support of firm i's cost distribution. Let $\underline{c} = \min \{\underline{c}_i\}$ and $\overline{c} = \max \{\overline{c}_i\}$. Denote firm i's pricing function by $P_i : [\underline{c}_i, \overline{c}_i] \rightarrow \mathbb{R}_+$. Let $\mathbf{P} = (P_1, \ldots, P_n)$ be the vector of pricing strategies, H_P the market price distribution induced by $\mathbf{P}$ and Θ, and $\mathbf{P}_{-i}(c_{-i}) = (P_j(c_j))_{j \neq i}$ the price vector of firms other than i.

DEFINITION 2: *The strategy profile* $\{\mathbf{P}, \sigma\}$ *is a* market equilibrium *if*

(i) σ is a personal equilibrium for the price distribution H_P.

(ii) For each i and $c_i \in [\underline{c}_i, \overline{c}_i]$,

$$\mathbf{P}_i(c_i) \in \operatorname{argmax}_{p_i \in \mathbb{R}_+} (p_i - c_i) \Pr[\sigma(p_i, \mathbf{P}_{-i}(c_{-i}), \chi) = i \mid c_i].$$

Our definition of market equilibrium extends Bayesian Nash equilibrium to allow for reference effects in consumer behavior. A market equilibrium needs to satisfy two conditions. First, consumers play a personal equilibrium given the correctly forecasted price distribution.[8] Second, each firm at each cost realization plays a best response to other firms' pricing strategies, taking consumers' reference points as given.[9]

[8] For notational simplicity, our definition implicitly imposes that there is a single representative consumer, or all consumers play the same personal equilibrium. We argue in Section V that this does not affect our results.

[9] Hence, at the stage when firms' prices are chosen, these prices do not influence a consumer's reference point. This specification reflects our assumption that the reference point is *lagged* rather than contemporaneous expectations. If, for instance, a consumer had been confidently expecting to pay \$14.99 for a CD that she now finds costs \$18.98, she presumably adjusts her beliefs about the price distribution immediately, but she would still experience paying \$18.98 as a loss. While the expectations that are relevant for specifying the reference point are clearly lagged, the fact that we do not specify when exactly these expectations are formed is a weakness of our approach. In addition, we exclude any direct influence of firm behavior on what expectations consumers form. As we show below, for instance, there is typically a continuum of focal prices, so firms have a strong incentive to manage consumers' price expectations. While firms might be able to do so through public commitments to prices, advertising, or other marketing activities, analyzing these motives is beyond the scope of this paper.

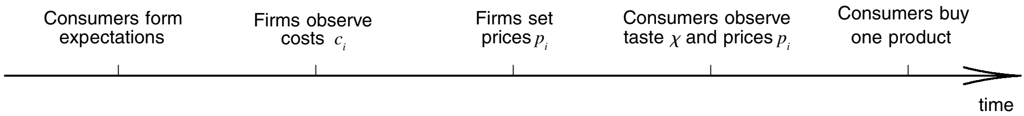

FIGURE 2. TIMING OF FULL MARKET MODEL

One of our goals in this paper is to investigate circumstances under which all firms charge the same price irrespective of their cost positions. We introduce a term for such a situation.

DEFINITION 3: *A market equilibrium is a* focal-price equilibrium *if there is a price p^* that all firms charge with probability one. In a focal-price equilibrium, p^* is the* focal price.

To address what we believe is mostly a technical issue that arises in our model as well as the standard Salop model with cost asymmetry, we introduce a restricted class of market equilibria.

DEFINITION 4: *A market equilibrium is an* interior equilibrium *if, for any equilibrium price vector, all firms sell to a positive measure of consumers on each side of their location.*

In principle, it is possible that a firm prices so low relative to a neighbor that it attracts all consumers between them. In the standard Salop model, at such a price level there is a discontinuity in the firm's demand, as it suddenly captures *all* consumers on the other side of its neighbor. In our model, both convex kinks and discontinuities in demand are possible. Since we do not know how to handle equilibria with such situations, we follow previous models with cost uncertainty (e.g., Philippe Aghion and Mark Schankerman 2004) and focus on (what we call) interior equilibria.

In the following three sections, we analyze the model above. As a foundation for this analysis, Lemma 1 derives a key part of firms' incentives, the price responsiveness of demand. Suppose the consumer had expected the distribution of her purchase price and the distribution of her acquired product's distance from ideal to be $F(\cdot)$ and $G(\cdot)$, respectively. These distributions are determined as part of a personal equilibrium by the consumer's strategy σ and the distribution of market prices H in the following way:

$$F(p') = \Pr_{p\sim H, \chi\sim U[0,1]}[p_{\sigma(p,\chi)} \le p'] \text{ for any } p' \ge 0, \text{ and}$$

$$G(s) = \Pr_{p\sim H, \chi\sim U[0,1]}[d(\chi, y_{\sigma(p,\chi)}) \le s] \text{ for any } s \ge 0.$$

Denoting right and left limits by subscripts $\downarrow$ and $\uparrow$, $F(\cdot)$ and $G(\cdot)$ can be used to express how demand depends on prices:

LEMMA 1: *Suppose that given the consumer's lagged expectations $F(\cdot)$ and $G(\cdot)$ and the realized prices, the realized tastes that make the consumer indifferent between purchasing y_i and the respective neighboring products are at distances $x^+, x^- \in (0, 1/n)$ from y_i on its two sides. Then, the right derivative of firm i's residual demand, $D_i(p_i, p_{-i})$, with respect to its price p_i is*

$$(3) \qquad D_{i\downarrow}(p_i, p_{-i}) = -\frac{1}{2t}\left[\frac{2 + (\lambda - 1)F(p_i)}{2 + ((\lambda - 1)/2)(G(x^+) + G(1/n - x^+))} + \frac{2 + (\lambda - 1)F(p_i)}{2 + ((\lambda - 1)/2)(G(x^-) + G(1/n - x^-))}\right],$$

and the left derivative $D_{i\uparrow}(p_i, p_{-i})$ is given by the expression in which $F_{\uparrow}(p_i)$ replaces $F(p_i)$ above.

Whenever firm i's two neighbors set the same price p, we will denote the indifferent consumer on each side by $x = x^+ = x^-$, and firm i's demand as a function of its price by $D_i(\cdot, p)$.

The price responsiveness of a firm's residual demand D_i derives partly from the comparison effect explained above: as reflected in the numerators in equation (3), the higher the probability with which the consumer expected lower prices, the more she experiences paying a given price as a loss, and hence the more responsive she is to price changes. The comparison effect has two important implications we will use repeatedly in the paper. First, the residual demand curve is kinked at p_i if $F(\cdot)$ has an atom at p_i, and it is differentiable at p_i if $F(\cdot)$ has no atom at p_i. Second, the price responsiveness of demand is greater at higher prices in the purchase-price distribution.

More subtle than the effect of utility from money itself is the effect of product satisfaction on the price responsiveness of demand. A small price change can affect a consumer's choice if she is approximately indifferent between firm i's product at a distance x^+ from ideal and the neighbor's product at a distance $1/n - x^+$ from ideal. If any of these options is evaluated as a loss to a greater extent—that is, if the expected probability of a better product, $G(x^+)$ or $G(1/n - x^+)$, is higher—then a change in the consumer's realized taste has a greater effect on which option she prefers. This means that a given price change reverses the consumer's decision for a smaller range of taste realizations, lowering the price responsiveness of firm i's demand.

II. Existence and Properties of Focal-Price Equilibria

As our first step in analyzing the model, we establish a necessary and sufficient condition for a focal-price equilibrium to exist, and explore the condition's implications for the price level and the effect of industry concentration on pricing. The condition allows for stochastic costs, and even for commonly known *differences in* (stochastic or deterministic) costs.

To derive our condition, we solve for the cost levels c_1 for which firm 1 does not want to deviate from a focal price of p^*. Market equilibrium requires that the consumer anticipated all prices to be p^*, so that she expected to spend p^* with probability 1. In addition, since (having expected equal prices) she expected to buy the product closest to her taste, $G(\cdot)$ is the uniform distribution on $[0, 1/(2n)]$. In addition, the consumers who are indifferent between a firm and its neighbor are at distances $x^+ = x^- = x = 1/(2n)$ from the firm, so $G(x) = G(1/n - x) = 1$.

Given these considerations, equation (3) implies that $D_{1\downarrow}(p^*, p^*) = -1/t$. Using that $D_1(p^*, p^*) = 1/n$, so long as $(p^* - c_1)/t \geq 1/n$, firm 1 cannot benefit from locally raising its price. Similarly, since $D_{1\uparrow}(p^*, p^*) = -2/(t(1 + \lambda))$, so long as $2(p^* - c_1)/(t(1 + \lambda)) \leq 1/n$, firm 1 cannot benefit from locally lowering its price. Combining and rearranging these conditions, charging p^* is locally optimal if and only if

$$(4) \qquad p^* - \frac{t}{n}\,\frac{1 + \lambda}{2} \leq c_1 \leq p^* - \frac{t}{n}.$$

In the Web Appendix, we show that when local deviations are unprofitable, nonlocal deviations are also unprofitable. Therefore:

PROPOSITION 1: *A focal-price equilibrium exists if and only if*

$$\overline{c} - \underline{c} \leq \frac{\lambda - 1}{2}\frac{t}{n}.$$

When there is no loss aversion ($\lambda = 1$), a focal-price equilibrium exists only if $\overline{c} - \underline{c} = 0$—if all firms have the same deterministic cost. As explained above, however, if consumers are loss averse and expect all firms to charge the same price p^*, there is a kink in residual demand at p^*, so for a range of cost levels p^* is the optimal price to charge. Hence, with loss aversion a focal-price equilibrium can exist despite cost differences and variation.

Proposition 1 has a number of important comparative-statics implications for when a focal-price equilibrium exists. Naturally, a focal price is easier to sustain when the range of marginal costs $\overline{c} - \underline{c}$ is smaller. Also, a focal-price equilibrium is more likely to exist when consumer loss aversion is greater. The greater is λ, the greater is the difference between a consumer's sensitivity to price increases from p^* and price decreases from p^*. Hence, the greater is the difference between the effects on profits of price increases and decreases, and the greater is the range of cost levels for which p^* is the optimal price. One implication of this comparative static and our model more generally may be that, *ceteris paribus*, prices are less variable in consumer markets than in transactions between (presumably less loss averse) firms. Evidence in Alan S. Blinder et al. (1998) is broadly consistent with this prediction.

Most interestingly, a focal-price equilibrium is more likely to exist when market power, as measured by product differentiation relative to the number of firms (t/n), is greater. For an intuition, consider the price p^* at which a firm with cost $\overline{c}$ is just indifferent to raising its price. Then, due to a kink in demand, for a range of cost decreases it strictly prefers not to decrease its price. This range—and hence the allowed cost variation for a focal-price equilibrium to exist—is increasing in the markup $p^* - \overline{c}$, so that it is larger in less competitive industries. With a higher markup, the value of a marginal consumer is higher, so a change in the responsiveness of demand has a greater effect on the firm's incentives to change its price. Hence, the low responsiveness of demand to price decreases makes the firm more reluctant to cut its price, and it will not want to do so for a greater range of cost decreases.

In addition to identifying conditions under which a focal-price equilibrium exists, inequality (4) determines what the focal price level can be:

PROPOSITION 2: *There is a focal-price equilibrium with focal price* p^* *if and only if*

$$\overline{c} + \frac{t}{n} \leq p^* \leq \underline{c} + \frac{t}{n}\frac{1 + \lambda}{2}.$$

In the corresponding Salop model without loss aversion, the support of a firm's interior-Bayesian-Nash-equilibrium prices is bounded above by $\overline{c} + t/n$*, and this bound can be attained only if the firm has realized cost* $\overline{c}$.

Proposition 2 says that in a focal-price equilibrium, consumer loss aversion leads to increased prices: even at the *lowest* possible cost, a firm charges a higher price than it would in the standard model at the *highest* possible cost. Intuitively, if a firm unilaterally lowers its price, it attracts some consumers of the neighboring firms, who (unexpectedly) choose a good that both costs and matches their taste less than expected. Since consumers are more sensitive to the loss in satisfaction than to the gain in money, the firm attracts fewer of them than without loss aversion. But if the firm unilaterally raises its price, its consumers must either pay a higher price or get a less satisfactory product than they expected was possible, so—as either choice involves a loss—the firm loses the same number of consumers as without loss aversion. Since loss aversion decreases a firm's incentive to lower its price and leaves a firm's incentive to raise its price unchanged, it increases equilibrium prices.

Proposition 2 implies that if there is a focal-price equilibrium, there are generically multiple ones, with the set of possible focal prices being a closed interval. If consumers' expectation of the price increases from p to $p' > p$, the difference between paying p' and paying p turns from a loss to a foregone gain. Because this makes demand less responsive, firms are more willing to increase prices, within limits exactly matching the increased expectations.

Beyond a theoretical possibility, our model predicts that focal-price equilibria can exist for calibrationally nontrivial amounts of cost variation. Assuming $\lambda = 3$, which corresponds to the conventional assumption of about two-to-one loss aversion in observable choices (Amos Tversky and Daniel Kahneman 1992, for example), a focal-price equilibrium exists for cost variation $\overline{c} - \underline{c}$ up to t/n. Since by Proposition 2 the equilibrium markup lies in the interval $[t/n, 2t/n]$, the allowed cost variation is between 50 percent and 100 percent of firms' markups.

III. Conditions for All Equilibria to be Focal

In this section, we identify sufficient conditions under which firms with possibly different cost distributions suppress cost shocks and adhere to focal pricing of differentiated products in *any* interior market equilibrium. To our knowledge, no price-setting model predicts focal prices so robustly. We first establish that if the intervals containing the supports of firms' cost distributions overlap, there cannot be an equilibrium with stable but different prices—if each firm sets a deterministic price, they set the same one. Then, we show that if the density of each firm's cost distribution is sufficiently large everywhere on its connected support, prices are stable—each firm sets a deterministic price. Then, when both conditions hold, any market equilibrium is a focal-price equilibrium. We also give examples illustrating that if firms' cost distributions do not overlap, equilibria with different deterministic prices can exist.

The following proposition is the first part of our argument:

PROPOSITION 3: *Suppose* $\bigcap_{i \in N} [\underline{c}_i, \overline{c}_i] \neq \emptyset$. *If all firms set a deterministic price and either*

$$\lambda \leq 1 + \frac{2}{n-1}\left(1 + \sqrt{1 + 2n(n-1)}\right) \tag{5}$$

or $n = 2$, *the market equilibrium is a focal-price equilibrium.*

The intuition for Proposition 3 is the same as in the two-firm example of Section I: if firms do not charge the same price, a highest-price firm has a higher markup and a lower inframarginal demand than a lowest-price firm, and because by the comparison effect it tends to face a greater responsiveness of demand, either it or the lowest-price firm wants to deviate. This intuition, however, ignores an effect that (for $n > 2$) makes it necessary to impose condition (5). A change in a firm's price changes the distribution of marginal consumers in its two markets. By Lemma 1, this typically changes the price responsiveness of its residual demand. If demand responsiveness changed too fast, the firm's profit-maximization problem might not be single-peaked, and this would generate many technical difficulties. To rule out such possibilities, Proposition 3 above and Proposition 4 below impose restrictions on λ.

But condition (5) is relatively weak. It applies only when $n > 2$, and it is satisfied for any number of firms whenever $\lambda \leq 1 + 2\sqrt{2} \approx 3.8$. Since the conventional assumption of two-to-one loss aversion is equivalent to $\lambda = 3$, the condition does not seem very problematic.

As a second ingredient for the main result of this section, we give conditions such that all firms charge a deterministic price. Because analyzing a more general model is technically very difficult, we restrict attention to independent (idiosyncratic) cost shocks, still allowing for asymmetries in firms' cost distributions.[10]

PROPOSITION 4: *Suppose costs are independently distributed with $c_i \sim \Theta_i\,[\underline{c}_i, \overline{c}_i]$ and corresponding densities θ_i. If $38 > \lambda > 1$ and $(\overline{c} - \underline{c}) < (t/n)(3 + \lambda)/(2(1 + \lambda))$, there is a real number $\rho(\lambda, t, n, \overline{c} - \underline{c}) > 0$ such that, if*

$$\theta_i(c) > \rho(\lambda, t, n, \overline{c} - \underline{c})$$

for all $c \in [\underline{c}_i, \overline{c}_i]$, then firm i sets a deterministic price in any interior equilibrium.

Combining Propositions 3 and 4:

COROLLARY 1: *If the conditions of Propositions 3 and 4 hold, any interior market equilibrium is a focal-price equilibrium.*

It is worth noting that the function $\rho(\lambda, t, n, \overline{c} - \underline{c})$ that naturally drops out of our approximations underlying the proof of Proposition 4 is decreasing in t and increasing in n, and approaches zero as $t \to \infty$ and infinity as $n \to \infty$. Our sufficient conditions for all equilibria to be focal are therefore more likely to be met in less competitive industries.

To conclude this section, we provide some examples where the conditions of Proposition 3 do not hold but those of Proposition 4 may, illustrating the logic of market equilibrium with unequal prices and discussing further issues.

Example 1. *Suppose $n = 2$, $\lambda = 5$, and $t = 1$. As we verify in the online Appendix, there is a market equilibrium in which firm 1 always charges price $p_1 = 2$, firm 2 always charges price $p_2 = 9/4$, and the consumer buys from firm 1 with probability 3/4, if and only if $c_1 \in [1/8, 5/4]$ and $c_2 \in [2, 49/24]$.*

The conditions above for the existence of a market equilibrium with prices $p_1 = 2$ and $p_2 = 9/4$ allow for several possibilities. If costs are deterministic with $c_1 = 9/8$ and $c_2 = 97/48$, for instance, there is a market equilibrium with deterministic prices $p_1 = 2$ and $p_2 = 9/4$, and by Proposition 1 a focal-price equilibrium also exists. The intuition for why both types of equilibria can exist is the following. If consumers had expected the two firms to charge the same price, demand will be very responsive to increases from this price and not very respon-

[10] If costs are not independent, a change in c_i changes the distribution of competitors' prices conditional on c_i and hence also the distribution of marginal consumers for a given p_i. By Lemma 1, this typically changes the price responsiveness of residual demand. While we believe this consideration would not substantially modify the comparison effect, the main force driving our result, we cannot formally analyze this more general case.

sive to decreases from this price, so that it is optimal for both firms to charge this price. But if consumers had expected different prices, the responsiveness of demand in-between the two expected prices is at an intermediate level, so that it is optimal for the low-cost firm to charge the lower of the prices and for the high-cost firm to charge the higher of the prices.

In contrast, if costs are deterministic with $c_1 = 1$ and $c_2 = 97/48$, charging deterministic prices $p_1 = 2$ and $p_2 = 9/4$ is still a market equilibrium, but in this case a focal-price equilibrium does not exist. More generally, if firm 1's and firm 2's marginal costs are independently and narrowly distributed around 1 and 97/48, respectively, Proposition 4 implies that each firm charges a deterministic price in any market equilibrium, and Proposition 1 implies that these prices are different. Hence, sticky pricing—the unresponsiveness of prices to cost circumstances—does not necessarily go hand in hand with focal pricing—equal pricing across firms. Intuitively, if each individual firm's cost distribution is sufficiently narrow, the firm's price will be invariant to its cost realization. But if one firm is at the same time much more efficient than the other, the deterministic prices of the two firms must be different.

While not generating focal pricing, in some ways the example above still illustrates how loss aversion can lead to reduced price variation and lower competition—two important themes in the paper. It is easy to check that in a standard Salop model, an equilibrium with a price difference of 1/4 requires a cost difference of 3/4. In our example, a cost difference of up to 23/12 can support the same price difference, showing that with loss aversion prices can be much closer to each other. Indeed, unlike in our setting, in the standard model any cost difference above 3/2 would lead the low-cost firm to price the high-cost firm out of the market. Loss aversion therefore reduces competition and allows both firms to make positive profits.

Although our example does not speak directly to situations with more than two goods, its logic also suggests that in many situations uniform pricing is more likely to happen than focal pricing. If a single firm's cost distributions for its different products are narrow and overlapping, the firm will often set the same deterministic price for all its products. But, again, if one firm has much lower costs overall than the other, the uniform prices of the two firms will have to differ.

IV. Industry-Wide Cost Shocks

In this section, we fully characterize symmetric equilibria when firms always have identical marginal costs—that is, when they are subject only to industry-wide cost shocks. This allows us to study, in a tractable model, the responsiveness of price to cost when conditions for a focal-price equilibrium are not necessarily met. We find that markups strictly decrease with cost in any market equilibrium, and that the price may be sticky (unchanging in cost) in some regions. Furthermore, markups decline faster with cost, and prices tend to be more sticky, in more concentrated industries.

Suppose firms' common marginal cost is continuously distributed according to Θ on $[\underline{c}, \overline{c}]$, with corresponding density θ. We first establish two basic properties of symmetric market equilibria:

LEMMA 2: *Suppose firms have identical, continuously distributed marginal costs. In a symmetric market equilibrium, price is a continuous and nondecreasing function of marginal cost.*

To understand the lemma, take costs c and c' and corresponding equilibrium prices p and $p' > p$, and suppose that residual demand is differentiable at both p and p'. Because firms use symmetric strategies, inframarginal demand is the same at the two prices (and equal to $1/n$). In addition, due to the comparison effect, demand is (weakly) more responsive to unilateral price changes at the high price p' than at the low price p. In order for firms' first-order conditions to be satisfied at both costs, therefore, c' must be greater than c and not arbitrarily close to it.[11]

We now fully characterize the set of symmetric-equilibrium pricing functions $P:[\underline{c}, \overline{c}] \rightarrow \mathbb{R}$, and then turn to a detailed discussion of the implications of this characterization. As a step toward a full analysis, we posit that for a cost c, $P(c)$ is not an atom of the market price distribution, and we derive $P(c)$. Since in a symmetric equilibrium firms set identical prices in all states of the world, consumers always choose the product closest to their taste. Hence, as in Section II, $G(\cdot)$ is the uniform distribution on the interval $[0, 1/(2n)]$. Furthermore, equation (3) implies that the derivative of firm 1's residual demand exists at $P(c)$ and is equal to

$$(6) \qquad -\frac{1}{t}\frac{2+(\lambda-1)F(P(c))}{1+\lambda} = -\frac{1}{t}\frac{2+(\lambda-1)\Theta(c)}{1+\lambda},$$

where $F(P(c)) = \Theta(c)$ because $P(\cdot)$ is nondecreasing and $P(c)$ is not a pricing atom. Substituting equation (6) into the firm's first-order condition, using that $D_1(P(c), P(c)) = 1/n$, and rearranging yields

$$(7) \qquad P(c) = c + \frac{t}{n}\frac{2+(\lambda-1)}{2+(\lambda-1)\Theta(c)} \equiv \Phi(c).$$

Expression (7) and Lemma 2 impose strong restrictions on a symmetric-market-equilibrium pricing function. For any $c \in [\underline{c}, \overline{c}]$ that is not on a flat part of $P(\cdot)$, $P(c)$ is not a pricing atom, so $P(c) = \Phi(c)$. In addition, arbitrarily close to an interior end of a flat part, there are costs c for which $P(c)$ is not a pricing atom, where again $P(c) = \Phi(c)$. Hence, at interior ends a flat part of $P(\cdot)$ connects continuously to $\Phi(\cdot)$. Finally, because for $c = \underline{c}$ equation (6) is the left derivative of demand whether or not $\underline{c}$ is a pricing atom, for price decreases from $\underline{c}$ to be unprofitable, we must have $P(\underline{c}) \leq \Phi(\underline{c})$; and by a similar argument, $P(\overline{c}) \geq \Phi(\overline{c})$.

The conditions above are in fact not only necessary, but also sufficient for $P(\cdot)$ to be a symmetric-market-equilibrium pricing function:

PROPOSITION 5: *Suppose firms have identical marginal costs distributed according to Θ on $[\underline{c}, \overline{c}]$. A pricing function $P:[\underline{c}, \overline{c}] \rightarrow \mathbb{R}$ is a symmetric-market-equilibrium pricing function if and only if all of the following are satisfied:*

(i) $P(\cdot)$ is continuous and nondecreasing;

(ii) There are disjoint intervals $[f_1, f_1'], [f_2, f_2'], \ldots \subset [\underline{c}, \overline{c}]$ such that $P(\cdot)$ is constant on all $[f_i, f_i']$ and not constant on any interval not contained in any $[f_i, f_i']$;

(iii) $P(c) = \Phi(c)$ for any $c \notin \cup_i [f_i, f_i']$;

(iv) $P(\underline{c}) \leq \Phi(\underline{c})$ and $P(\overline{c}) \geq \Phi(\overline{c})$.

To start identifying the implications of Proposition 5 in specific cases, suppose that $\Phi(\cdot)$ is strictly increasing. In that case, $P(\cdot)$ cannot have a flat part: because $P(\underline{c}) \leq \Phi(\underline{c})$ and $P(\overline{c}) \geq$

[11] If the price distribution has atoms at p or p', so that residual demand is not differentiable, the same argument still works by considering—instead of first-order conditions—incentives to lower one's price from p' as compared to incentives to raise one's price from p.

$\Phi(\bar{c})$, a flat part cannot start at either of these points and connect continuously to $\Phi(\cdot)$; and an interior flat part cannot connect continuously to $\Phi(\cdot)$ at both ends. Hence, there are no pricing atoms, and the unique symmetric market equilibrium has $P(c) = \Phi(c)$ everywhere:

COROLLARY 2: *Under the conditions of Proposition 5, if $\Phi(c)$ is strictly increasing, the unique symmetric market equilibrium has pricing strategies $P(c) = \Phi(c)$. Otherwise, a symmetric equilibrium with strictly increasing pricing strategies does not exist.*

But $\Phi(\cdot)$ is not necessarily strictly increasing. Differentiating equation (7) with respect to c,

$$\Phi'(c) = 1 - \frac{t}{n}\frac{(1+\lambda)(\lambda-1)\theta(c)}{[2+(\lambda-1)\Theta(c)]^2}, \tag{8}$$

which is negative if $\theta(c)$ is very high. If $\Phi(\cdot)$ is nonincreasing, then $P(\cdot)$ cannot have a strictly increasing part—where it would have to coincide with a nonincreasing $\Phi(\cdot)$—so that it is constant. Hence, in these situations, any symmetric market equilibrium is focal:

COROLLARY 3: *Under the conditions of Proposition 5, if $\Phi(c)$ is nonincreasing, any symmetric market equilibrium is a focal-price equilibrium. Otherwise, symmetric equilibria other than focal-price equilibria exist.*

As with Proposition 4, the intuition for this result is most easily seen by first assuming that consumers expected firms' prices to be strictly increasing in cost. If the density of the cost distribution is high, a small increase in c implies a large increase in $F(P(c))$ and hence a large increase in the comparison effect and the corresponding price responsiveness of demand. This is inconsistent with equilibrium: a firm can increase profits either by decreasing prices at higher costs and attracting substantial extra demand, or by increasing prices at lower costs without losing many consumers. Since this is true for any strictly increasing pricing strategy, the equilibrium price must be constant.

If $\Phi(\cdot)$ is neither strictly increasing nor nonincreasing, Proposition 5 implies that market-equilibrium pricing functions will generally consist of flat parts pasted together continuously with strictly increasing parts that coincide with $\Phi(\cdot)$. Figure 3 illustrates a nonmonotonic $\Phi(\cdot)$ and possible market equilibria. For $c \in [\underline{c}, c']$ and $c \in [c'', \bar{c}]$, the pricing function cannot have a flat part, because that could not be pasted continuously with $\Phi(\cdot)$. Hence, in these regions $P(\cdot)$ is strictly increasing and therefore equal to $\Phi(\cdot)$. The nondecreasing $P(\cdot)$, however, must be "ironed out" over the range where $\Phi(\cdot)$ is decreasing. Furthermore, because at the ends of a flat interval $P(\cdot)$ connects continuously to increasing parts of $\Phi(\cdot)$, it has exactly one flat part. $P^1(\cdot)$ and $P^2(\cdot)$ are two possible market-equilibrium pricing functions.

In combination with equation (7), Proposition 5 has a number of important implications for symmetric equilibria. Two implications are about the level and variation in markups in our model relative to the standard one (identical to $\lambda = 1$ here). In the standard Salop model, the markup is constant in cost and equal to t/n. As in focal-price equilibria (Proposition 2), one effect of loss aversion is to increase the price level: the markup is strictly greater than t/n for $c < \bar{c}$, and greater than or equal to t/n for $c = \bar{c}$. The consumers a firm attracts by lowering its price experience a pure loss in product satisfaction (from choosing a product unexpectedly far from ideal), and unless $c = \bar{c}$, only some combination of gain and avoided loss in money. Hence, they are more difficult to attract than in the standard setting, decreasing competition and increasing prices.

The other effect of loss aversion is to decrease price variation by making markups strictly decreasing in c:

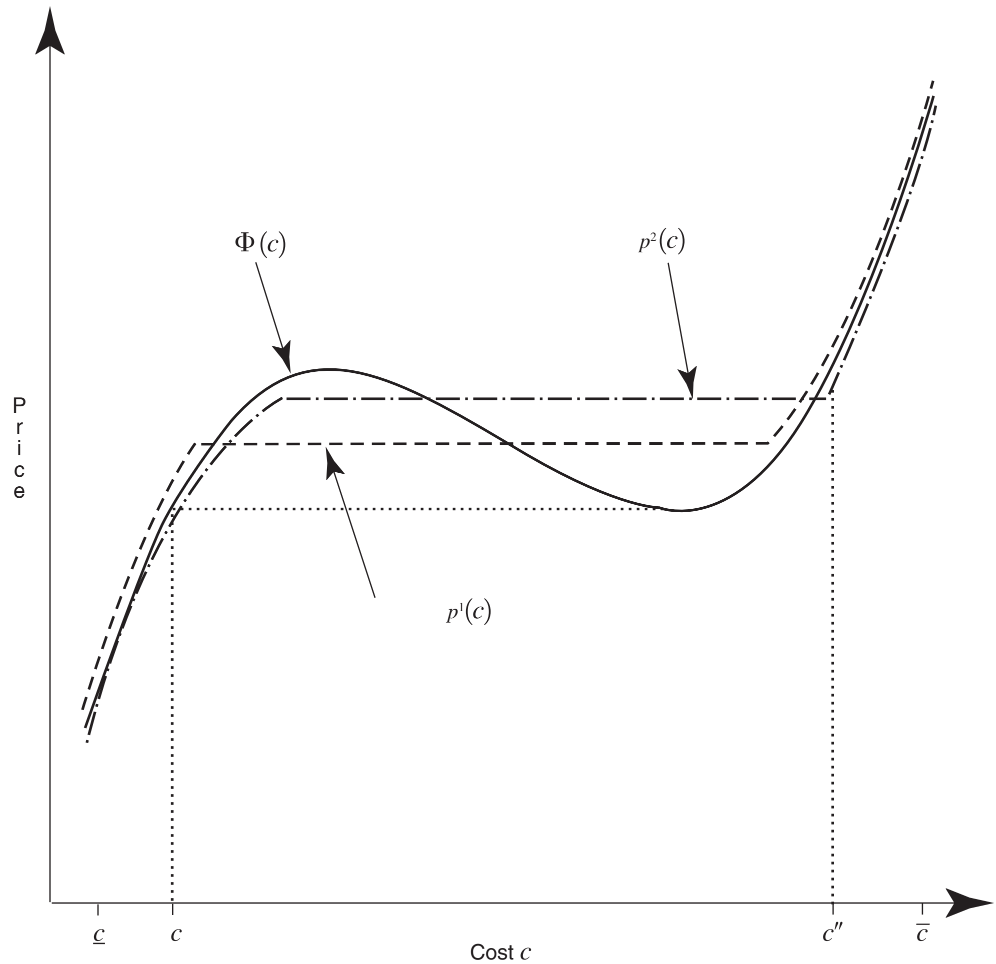

FIGURE 3. DETERMINATION OF EQUILIBRIUM PRICING FUNCTIONS WITH INDUSTRY-WIDE SHOCKS

Notes: For better visibility, overlapping curves are drawn as close parallel curves. The solid curve $\Phi(c)$ is the solution to the first-order condition for optimal pricing assuming that $P(c)$ is not an atom of the price distribution. Since a market-equilibrium pricing function is nondecreasing and continuous, it consists of constant parts pasted together with strictly increasing parts that coincide with $\Phi(\cdot)$. Two market-equilibrium pricing functions are $P^1(c)$ and $P^2(c)$.

COROLLARY 4: *Under the conditions of Proposition 5, in any symmetric market equilibrium* $P(c) - c$ *is strictly decreasing in c on the support of* Θ.

This prediction of our theory is potentially relevant for understanding macroeconomic fluctuations. Extensive evidence reviewed by Rotemberg and Woodford (1999) indicates that costs are strongly procyclical. Hence, our model implies markups are countercyclical.[12] Intuitively, recall that due to the comparison effect, consumers are more responsive to price changes at higher than at lower prices within the price distribution. Since inframarginal demand is constant across the price distribution, firms compete more fiercely at higher prices, reducing markups.

[12] Of course, if one measures countercyclicality using the Lerner index $(p - c)/p$, the Salop model without loss aversion also features countercyclical markups.

Proposition 5 implies not only that price variation is lower than in the standard model, but also that it is systematically related to the competitiveness of the market. The more concentrated is the industry (the lower is n) and the greater is product differentiation (the greater is t), the lower is $\Phi'(c)$ at any c (equation (8)). As a result, the more countercyclical are markups, the faster $P(c) - c$ decreases with c when price is strictly increasing in cost, and the more likely it is that any symmetric equilibrium is a focal-price one. Intuitively, with the higher average markups firms enjoy in a less competitive industry, the increased ability to attract consumers at higher prices has a greater impact on firms' incentive to cut prices, generating markups that decrease faster in cost. If markups are very high, the impact of an increase in demand responsiveness on firms' incentive to cut prices is so great that firms are unwilling to raise their price at all—they charge a sticky (and focal) price.

In fact, Proposition 5 allows us to more fully describe pricing patterns for industries ranging from very competitive ($t/n \approx 0$) to very uncompetitive ($t/n \to \infty$). If competition is sufficiently strong, the unique symmetric market equilibrium features a strictly increasing pricing function, which is close to marginal-cost pricing if competition is very strong. At lower levels of competition, markups are higher and more countercyclical. At even lower levels of competition, the price is constant in cost near regions where the cost distribution is relatively dense, but may remain strictly increasing in cost in other regions. At very low levels of competition, the price is sticky and focal.

It is important to note that in this section identical pricing across firms was assumed, not derived. The question arises whether such identical pricing would be the outcome in an environment where idiosyncratic cost shocks also exist. For cases in which the cost variation is sufficiently small, Section II has shown that a focal-price equilibrium exists even when there is both industry-wide and idiosyncratic cost uncertainty. More strongly, although (for reasons mentioned above) we cannot fully analyze a general model with both types of cost uncertainty, intuition developed in the last two sections suggests that in regions where both components vary little, the price will be focal in any equilibrium.[13] But in regions where the common cost shocks are not absorbed, residual demand will be smooth, so idiosyncratic cost shocks will also not be absorbed.

V. Robustness

In this section, we argue that our results are largely robust to natural variations of our model. In short, most of our results rely on the simple intuition that—due to loss aversion in money—a consumer's sensitivity to price changes is increasing in the probability with which she expected to pay lower prices, and this force is not eliminated by reasonable modifications of the model.

Because in many situations consumers are unsure either about what they want or about what is available, we have assumed a dispersed prior on χ. But most of our results do not depend on this assumption. Even if χ is known perfectly, our results in Sections II and IV remain unchanged, with the same logic and essentially the same proofs.[14] Under reasonable refinements, such as that consumers play the ex ante optimal personal equilibrium, a version of Proposition 3 also remains

[13] Again, the intuition can be seen by assuming that cost shocks are not absorbed. If costs vary little, the price distribution will be dense, so that the comparison effect implies that a small increase in the price leads to a large increase in the price responsiveness of demand. Then, a firm can increase profits either by increasing lower prices (where demand is relatively inelastic) or by decreasing higher prices (where demand is relatively elastic).

[14] Technically, if different consumers have different information about their preferred locations, there cannot be a representative consumer. The definition of market equilibrium can be modified easily to account for such heterogeneity.

true—although for a completely different reason than above.[15] To illustrate, suppose $n = 2$ and firms 1 and 2 charge prices p_1 and $p_2 > p_1$. If a consumer prefers product 1 ex ante and plans to buy it, to avoid a loss in money she prefers it even more ex post. If she prefers product 2 ex ante and plans to buy it, to avoid a loss in satisfaction she prefers it even more ex post. With all consumers "locked in," both firms want to raise their price, contradicting equilibrium.

Our methods in Proposition 4 do not extend to the case when χ is known with certainty. The logic of our proof, however, seems to rely only on sufficiently many *marginal* consumers being sufficiently uncertain about their relative preference for at least two neighboring products that they are unsure as to which one they will buy. These consumers exhibit a similar pattern of behavior to our consumers above, so they give a firm similar incentives.

As do most applications of the Salop model, our model assumes that firms' prices affect only the allocation of demand, not its level. One way to model a market-size effect is to assume that consumers have an outside option with a randomly determined level of utility. In this case, the comparison effect makes consumers on the margin between two firms, as well as on the margin between a firm and the outside option, more responsive to price changes at higher than at lower prices in the purchase-price distribution. Hence, our qualitative results on reduced price variation (but not necessarily our result that loss aversion increases prices) are likely to survive.

We assume above that each firm sells exactly one product. As long as no firm owns neighboring products, our results carry over unchanged to multiproduct firms, so that results on focal pricing translate directly into uniform pricing.[16] In an interior equilibrium, the incentive for locally changing one product's price is unaffected by how many nonneighboring products a firm owns. But global deviations are weakly less profitable for a multiproduct firm because such a firm might be cannibalizing its own market.[17]

In our model, consumer demand is stable and symmetric across firms, but firms have possibly different and uncertain marginal costs. In most industries, firms do differ in the features and popularity of their products, leading to different elasticities of residual demand even if all of them set the same price. In our setting, for instance, the intrinsic valuations of firms' products could be different and random instead of always taking the same value v. Because a change in its product's intrinsic value has similar implications for a firm's pricing incentives as a change in its marginal cost, results closely analogous to those above would likely hold in this alternative model.

While we have assumed that industry structure is exogenous, our model can be extended to allow for endogenous entry. Suppose industry concentration is determined by a fixed cost that firms must pay to enter the market, and post-entry products are located equidistant from each other. Since the fixed cost determines the number of firms but has no impact on market equilibria given the number of firms, our qualitative results on the effect of industry concentration on market equilibria survive.

Our results on focal pricing and reduced price variation more generally hold in a model in which consumers are loss averse only in the money dimension. This assumption would, in fact,

[15] In fact, the assumption that consumers play the ex ante optimal personal equilibrium is unnecessarily strong. It suffices to assume, for example, that consumers have an arbitrarily small amount of self-discipline in the sense that they can select ex ante whether to impose an arbitrarily small ex post cost on themselves if they choose a certain action deemed undesirable ex ante.

[16] In this case, the relevant measure of the competitiveness of the industry depends (in addition to t) on the number of products rather than on the number of firms.

[17] Even if firms can own two neighboring products, all the forces behind our results are still present, so that focal pricing will often be an equilibrium, and often the only type of equilibrium. Our conditions and proofs, however, would have to account for the decreased competition between firms and for the fact that products may differ in how many neighboring products they compete with. When a firm can own three neighboring products, the middle one faces no immediate external competition, so the firm always sets a higher price for it.

substantially simplify some of our formal statements and proofs (especially those of Propositions 3 and 4). We make the assumption that consumers are also loss averse in the product dimension both because it is far more realistic and experimentally and theoretically well-motivated, and because it shows the robustness of results to including loss aversion in things other than money.

Our definition of market equilibrium assumes that all consumers play the same personal equilibrium. Relaxing this assumption does not affect our results. In all situations in Sections II and IV, selection is a nonissue simply because the personal equilibrium is unique. Our proofs in Section III work by investigating how the responsiveness of a firm's residual demand changes across the price distribution. Since our bounds hold for any personal equilibrium a person might be playing, they also hold if consumers play different equilibria.

The results in this paper are also robust to heterogeneity in loss aversion among consumers. Our estimation methods would have to account for such heterogeneity, but as long as there is some loss aversion in the population, the results would survive in some form.

A more fundamentally different extension of our model than all the considerations above is the incorporation of dynamics, and we conclude this section by intuitively discussing how this may affect our results. Suppose the firms play the pricing game T times with costs independently redrawn in every period, and consumers' reference points depend on lagged rational expectations regarding the distribution of outcomes they are going to get in each period. Clearly, since for a range of costs firms cannot increase even current profits by deviating from focal pricing in any single period, our existence results easily extend to this dynamic setting. But whether focal and sticky pricing emerges as the *unique* possible outcome is far trickier. A major complication is that past prices can, in general, affect a consumer's expectations and hence also her reference point for future outcomes. While a full and realistic analysis of this issue seems important—and could potentially lead to novel models of advertising and price leadership—it is beyond the scope of this paper.[18] Hence, to abstract from this consideration, we assume that the expectations determining a consumer's reference point are formed before she observes any prices.

With this assumption, the same logic that underlies our results in this paper seems to imply that if firms' cost distributions overlap and have sufficiently high density, then in each period prices will be focal. Because there can be a continuum of static focal-price equilibria, however, it is not necessarily the case that firms set the same price from period to period. Whether this is guaranteed depends on how exactly lagged expectations determine consumers' reference points. At one extreme, suppose that consumers compare their outcomes in a period only to their lagged expectations specific to that period, possibly because new consumers arrive in each period. Then, any sequence of static focal prices is an equilibrium; it could be that consumers expect the price of a CD to be \$15 one week and \$20 the next, and firms comply with both these expectations. At the other extreme, suppose that a consumer forms expectations regarding all T periods, and her reference point for outcomes in each period is an average of these expectations. Then, the logic of our results seems to indicate that the unique equilibrium is to charge the same focal price in each period: just like it cannot be an equilibrium for firms to charge different deterministic prices, it cannot be an equilibrium for them to charge different prices in different periods. More generally, if consumers' average reference point changes slowly—for example because a small fraction of consumers is replaced in each period—there is a tendency for firms not to move away from prices set in previous periods.

[18] Nevertheless, there is a way in which observing past prices might increase the tendency for sticky pricing: if consumers take past prices as salient indicators of the future play of firms, firms may have a strong incentive to comply with these expectations.

VI. Related Literature

Loss aversion features prominently in at least two somewhat separate literatures. In experimental and behavioral economics, loss aversion and reference dependence explain a number of robust phenomena, including the endowment effect and small-scale risk aversion.[19] More closely related to our topic, empirical evidence in marketing indicates loss aversion in consumer behavior that is broadly consistent with the consumer model of this paper. Consumers seem to compare actual market prices to "reference prices" determined at least partly by "price beliefs" or expectations, and purchases are more sensitive to losses from the reference price than to gains relative to it (Gary M. Erickson and Johny K. Johansson 1985; Manohar U. Kalwani and Chi Kin Yim 1992; Russell S. Winer 1986). Bruce G. S. Hardie, Eric J. Johnson, and Peter S. Fader (1993) find loss aversion in evaluations of quality as well. Our paper develops a model of reference prices based on insights from behavioral economics, and—asking a question not formally addressed in either literature—examines ways this affects the strategic interactions between firms. As such, it belongs to a growing literature sometimes called "behavioral industrial organization," which examines the impact of bounded rationality and other psychological factors on the competition between firms. See Glenn Ellison (2006) for a review.

There are several prominent theories of price stickiness, some of which also feature focal-price equilibria. The stickiness of prices is a robust feature of these theories, but we will argue that their aim is not to explain the *equality* of prices: if they are extended to allow for asymmetric firms and differentiated products, they either become inconsistent with focal pricing, or predict a large number of equilibria with no compelling reason to select the focal one. Furthermore, these theories do not ask how the competitiveness of the industry affects price variation, and do not address the issue of uniform pricing.

Perhaps the most commonly invoked theory of price stickiness is that of menu costs. Menu costs generate a disincentive to change prices, but not an incentive to set identical prices. Furthermore, in some situations prices tend to be sticky even though menu costs seem to be zero.[20]

Formalizing the casual view of many researchers and observers that price stickiness and focal pricing are due to collusive behavior,[21] Susan Athey, Kyle Bagwell, and Chris Sanchirico (2004) show that in a repeated price-setting game, the optimal symmetric equilibrium is often a focal-price equilibrium. This equilibrium is enforced by the threat of a price war in case of a price change, and is efficient because the price war is never triggered. In asymmetric environments, however, there is no reason for firms' sticky prices to coincide. Similarly, if each colluding oligopolist sells multiple differentiated products, there is no reason to set the same price for all those products.

Rotemberg (2004) develops a monopoly model in which consumers both dislike price changes and are willing to punish the firm if they perceive it to be insufficiently altruistic. Even if selfish, the firm pretends to be sufficiently altruistic by setting the highest acceptable price. The model predicts that observable increases in input costs lead to price increases, but increases in

[19] Kahneman, Jack L. Knetsch, and Richard H. Thaler (1990, 1991), for instance, find that the minimal acceptable selling price for an object is higher than the maximum price at which people are willing to buy the same object, presumably because subjects construe giving up the object as a loss. And as argued by Rabin (2000), Rabin and Thaler (2001), Nicholas Barberis, Ming Huang, and Thaler (2006), and other researchers, the most significant source of aversion to risk over modest stakes is loss aversion.

[20] For example, Kashyap (1995) finds sticky pricing in retail catalogues, even when new catalogues are printed anyhow. David Genesove (2003) documents substantial rigidity in apartment rents, even though a new lease is filled out and signed every year for most apartments in his sample.

[21] This view is expressed, for instance, in Denis W. Carlton (1989, 914–15) and Christopher R. Knittel and Victor Stango (2003, 1704–05). In addition, focal prices and reduced price variability seem to have raised suspicions of collusion in other cases, such as the recent Sony-BMG merger case in Europe.

demand may not. While this captures an important aspect of price dynamics our model misses, Rotemberg's single-product monopoly setup cannot address focal or uniform pricing.

An important class of models with implications for price variation assumes that consumers must pay search costs to sample firms' products and prices. These models, however, often generate excess rather than reduced-price variation. If search costs are bounded away from zero and the first search is costly, there is no focal-price equilibrium, even with deterministic identical costs: if a consumer expecting price p shows up at a firm, the firm knows she values the good above p, and can raise the price. If search costs are not bounded away from zero, the situation is more complicated. If consumers observe the price distribution, a price increase by a firm triggers search by some consumers arriving at the firm, and a price decrease triggers search by some consumers arriving at other firms. Joseph E. Stiglitz (1987) shows that as a result of these opposing forces, price stickiness obtains if search costs are convex in the number of searches, but excess price variation results if search costs are concave. Finally, if consumers do not observe the price distribution, a price increase triggers search by consumers arriving at the firm, but (because it is not observed) a price decrease does not trigger search by consumers arriving at other firms. This can lead to price stickiness.[22] But even in this case, equilibria with nonequal prices exist, and with asymmetry there is no compelling reason to select a focal-price one.

Our model is related to an older literature on kinked demand curves (Hall and Hitch 1939; Sweezy 1939). In these models, each firm believes that rivals will follow price decreases but not price increases—leading to a kinked demand curve. Eric Maskin and Jean Tirole (1988) provide a game-theoretic foundation for these beliefs, but do not investigate the impact of cost shocks on pricing behavior. In addition, once we drop their assumption that all consumers buy from a lowest-price firm, there is no reason to presume that equilibria would necessarily be focal. More distantly, our paper is also related to the literature on switching costs, whereby consumers face an exogenously given cost when buying a product different from the one they purchased previously.[23] Consumers in our model face a kind of "psychological" cost when switching away from their expected outcomes: they dislike trading an unanticipated loss in one dimension for an unanticipated gain in the other dimension. Both kinds of switching costs predict increased prices. But whereas in typical switching-cost models the profits from increased prices are competed away in earlier periods when firms fight for unattached consumers, the same is not true in our model. More importantly, in our model the size of the switching cost is endogenous and situation-dependent. In particular, the key feature of our model is that consumers are less reluctant to switch in response to a price change if they construe it as more of a change in a loss rather than a gain. Since this feature is not present in classical switching-cost models, these models do not predict reduced price variation.

VII. Conclusion

A basic premise of our model is that consumers have accurate expectations about prices. Marketing studies disagree whether consumers can even recall prices for recently purchased products accurately, with estimates ranging from 5 percent to 50 percent of consumers.[24] The surveys upon which these estimates are based typically focus on the knowledge of average

[22] Nevertheless, for reasons similar to the logic in Hal R. Varian (1980), if there is a mass of informed consumers—who find out all prices for some reason—a focal-price equilibrium can once again not exist. With search costs, any equilibrium must have positive expected profits. Then, if all firms were to charge the same price, undercutting other firms slightly would attract all informed consumers, increasing profits.

[23] For a recent survey, see Joseph Farrell and Paul Klemperer (2007).

[24] For a recent meta-study, which includes an extensive literature survey, see Hooman Estelami and David R. Lehmann (2001).

consumers and do not study the accuracy of price expectations. While for simplicity we have assumed that all consumers have correct price expectations, our effects are driven by marginal consumers—consumers who will switch in response to some relevant price changes—and so require only (some of) these consumers to have correct expectations. Unfortunately, we are not aware of any empirical work on whether they do.

In contrast to our results, which predict reduced price variation in a number of senses, there often seems to be excess price variation between even identical products.[25] Such price variation seems to occur primarily in industries where consumers cannot or do not compare prices across different sellers, partly because firms deliberately make comparisons difficult.[26] Our theory is more suitable to environments in which prices and relevant features of products are transparent.

REFERENCES

Aghion, Philippe, and Mark Schankerman. 2004. "On the Welfare Effects and Political Economy of Competition-Enhancing Policies." *Economic Journal*, 114(498): 800–24.

Athey, Susan, Kyle Bagwell, and Chris Sanchirico. 2004. "Collusion and Price Rigidity." *Review of Economic Studies*, 71(2): 317–49.

Barberis, Nicholas, Ming Huang, and Richard H. Thaler. 2006. "Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case of Narrow Framing." *American Economic Review*, 96(4): 1069–90.

Baye, Michael R., John Morgan, and Patrick Scholten. 2004. "Price Dispersion in the Small and in the Large: Evidence from an Internet Price Comparison Site." *Journal of Industrial Economics*, 52(4): 463–96.

Beck, Jonathan. 2004. "Fixed, Focal, Fair? Book Prices with Optional Resale Price Maintenance." WZB Markets and Political Economy Working Paper SP II 2004–15.

Blinder, Alan S., Elie R. D. Canetti, David E. Lebow, and Jeremy B. Rudd. 1998. *Asking About Prices: A New Approach to Understanding Price Stickiness*. New York: Russell Sage Foundation.

Carlton, Dennis W. 1989. "The Theory and the Facts of How Markets Clear: Is Industrial Organization Valuable for Understanding Macroeconomics?" In *Handbook of Industrial Organization*, Vol. 1, ed. Richard Schmalensee and Robert Willig, 909–46. Amsterdam: North-Holland.

Chevalier, Judith A., Anil K. Kashyap, and Peter E. Rossi. 2003. "Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data." *American Economic Review*, 93(1): 15–37.

Competition Commission of the United Kingdom. 1994. "The Supply of Recorded Music: A Report on the Supply in the UK of Prerecorded Compact Discs, Vinyl Discs and Tapes Containing Music." http://www.competition-commission.org.uk/rep_pub/reports/1994/356recordedmusic.htm#full 1994.

Einav, Liran, and Barak Y. Orbach. 2005. "Uniform Prices for Differentiated Goods: The Case of the Movie-Theater Industry." Unpublished.

Ellison, Glenn. 2005. "A Model of Add-on Pricing." *Quarterly Journal of Economics*, 120(2): 585–637.

Ellison, Glenn. 2006. "Bounded Rationality in Industrial Organization." In *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress,* ed. Richard Blundell, Whitney Newey, and Torsten Persson. Cambridge: Cambridge University Press.

Erickson, Gary M., and Johny K. Johansson. 1985. "The Role of Price in Multi-Attribute Product Evaluations." *Journal of Consumer Research*, 12(2): 195–99.

Estelami, Hooman, and Donald R. Lehmann. 2001. "The Impact of Research Design on Consumer Price Recall Accuracy: An Integrative Review." *Journal of the Academy of Marketing Science*, 29(1): 36–49.

[25] Baye, Morgan, and Scholten (2004), for example, document that for many products, different Internet-based retailers charge very different prices.

[26] Ellison (2005) and Xavier Gabaix and David Laibson (2006) identify competitive advantages of making price comparisons difficult. Gabaix and Laibson (2006) show that in the presence of some consumers who ignore add-on costs, firms have no incentive to make add-on prices transparent, even when it is very cheap to do so. Ellison (2005) gives natural conditions under which rational consumers who are unresponsive to an ex ante hidden add-on price also decrease competition in transparent aspects of the product.

Farrell, Joseph, and Paul Klemperer. 2007. "Coordination and Lock-In: Competition with Switching Costs and Network Effects." In *Handbook of Industrial Organization*, Vol. 3, ed. Mark Armstrong and Robert H. Porter, 1967–2071. Amsterdam: North-Holland.

Gabaix, Xavier, and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *Quarterly Journal of Economics*, 121(2): 505–40.

Genesove, David. 2003. "The Nominal Rigidity of Apartment Rents." *Review of Economics and Statistics*, 85(4): 844–53.

Hall, R. L., and C. J. Hitch. 1939. "Price Theory and Business Behavior." *Oxford Economic Papers*, 2(1): 12–45.

Hardie, Bruce G. S., Eric J. Johnson, and Peter S. Fader. 1993. "Modeling Loss Aversion and Reference Dependence Effects on Brand Choice." *Marketing Science*, 12(4): 378–94.

Heidhues, Paul, and Botond Kőszegi. 2007. "New Evidence on Focal Pricing." In Preparation.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1990. "Experimental Tests of the Endowment Effect and the Coase Theorem." *Journal of Political Economy*, 98(6): 1325–48.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1991. "Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias." *Journal of Economic Perspectives*, 5(1): 193–206.

Kalwani, Manohar U., and Chi Kin Yim. 1992. "Consumer Price and Promotion Expectations: An Experimental Study." *Journal of Marketing Research*, 29(1): 90–100.

Kashyap, Anil K. 1995. "Sticky Prices: New Evidence from Retail Catalogs." *Quarterly Journal of Economics*, 110(1): 245–74.

Knittel, Christopher R., and Victor Stango. 2003. "Price Ceilings as Focal Points for Tacit Collusion: Evidence from Credit Cards." *American Economic Review*, 93(5): 1703–29.

Kőszegi, Botond, and Matthew Rabin. 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics*, 121(4): 1133–65.

Kőszegi, Botond, and Matthew Rabin. 2007. "Reference-Dependent Risk Attitudes." *American Economic Review*, 97(4): 1047–73.

Maskin, Eric, and Jean Tirole. 1988. "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles." *Econometrica*, 56(3): 571–99.

McMillan, Robert S. 2004. "Estimating Demand for Differentiated Products with Continuous Choice and Variety-Seeking: An Application to the Puzzle of Uniform Pricing." PhD diss. Stanford University.

Rabin, Matthew. 2000. "Diminishing Marginal Utility of Wealth Cannot Explain Risk Aversion." In *Choices, Values and Frames*, ed. D. Kahneman and Amos Tversky, 202–08. New York: Russell Sage Foundation.

Rabin, Matthew, and Richard H. Thaler. 2001. "Anomalies: Risk Aversion." *Journal of Economic Perspectives*, 15(1): 219–32.

Rotemberg, Julio J. 2004. "Fair Pricing." National Bureau of Economic Research Working Paper 10915.

Rotemberg, Julio J., and Michael Woodford. 1999. "The Cyclical Behavior of Prices and Costs." In *Handbook of Macroeconomics*, Vol. 1B, ed. John B. Taylor and Michael Woodford, 1051–1135. Amsterdam: North-Holland.

Salop, Steven C. 1979. "Monopolistic Competition with Outside Goods." *Bell Journal of Economics*, 10(1): 141–56.

Slade, Margaret E. 1999. "Sticky Prices in a Dynamic Oligopoly: An Investigation of (s,S) Thresholds." *International Journal of Industrial Organization*, 17(4): 477–511.

Stiglitz, Joseph E. 1987. "Competition and the Number of Firms in a Market: Are Duopolies More Competitive Than Atomistic Markets?" *Journal of Political Economy*, 95(5): 1041–61.

Sweezy, Paul M. 1939. "Demand Conditions under Oligopoly." *Journal of Political Economy*, 47(4): 568–73.

Tversky, Amos, and Daniel Kahneman. 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty*, 5(4): 297–323.

Varian, Hal R. 1980. "A Model of Sales." *American Economic Review*, 70(4): 651–59.

Winer, Russell S. 1986. "A Reference Price Model of Brand Choice for Frequently Purchased Products." *Journal of Consumer Research*, 13(2): 250–56.