

The Impact of Consumer Loss Aversion on Pricing*

Paul Heidhues
WZB (Social Science Research
Center Berlin)
10785 Berlin, Germany
heidhues@wz-berlin.de

Botond Kőszegi
Department of Economics
University of California, Berkeley
Berkeley, CA 94720-3880
botond@econ.berkeley.edu

October 2004

Abstract We develop a model in which a profit-maximizing monopolist with uncertain cost of production sells to loss-averse, yet rational, consumers. We first introduce (portable) techniques for analyzing the demand of such consumers, and then investigate the monopolist's pricing strategy. We provide conditions under which a firm with continuously distributed marginal costs chooses a discrete price distribution; that is, prices are "sticky". Price stickiness is more likely to obtain when the cost distribution has high density, the price responsiveness of demand is low, or consumers are likely to purchase. Whether or not the monopolist's prices are sticky, markups follow a countercyclical pattern. Despite the tendency toward price stability, there are also circumstances in which a firm with unchanging cost offers random "sales" to attract more demand at higher prices.

Keywords: Reference-dependent utility, price stickiness, monopoly pricing, kinked demand curve, countercyclical markups, sales and promotions, (seemingly) predatory pricing.

* We thank Stefano DellaVigna, Chris Shannon,  Szeidl, and especially Wei Li and Matt Rabin for numerous helpful discussions, and have benefitted from comments of seminar participants at CalTech, LSE, Royal Holloway, UC Berkeley, the University of Bonn, and the University of Vienna.

1 Introduction

Experimental investigations of small-scale trading decisions, as well as monetary risk taking, indicate that preferences are reference dependent—people compare economic outcomes to relevant “reference points”, not only evaluate them according to absolute measures. One of the most robust and quantitatively significant regularities related to reference-dependent preferences is the phenomenon of *loss aversion*—people are more sensitive to losses relative to their reference point than to gains relative to it.¹ Evidence from the economics and marketing literatures indicates that reference dependence and loss aversion influences individuals’ behavior in the marketplace,² and firms seem to be aware of this fact (Hall and Hitch 1939, Marketing News 1985).³

Inspired by these findings, we analyze the strategic pricing behavior of a profit-maximizing monopolist facing loss averse consumers. Our results can help explain, under a single umbrella, three stylized facts about the distribution and time pattern of prices for consumer goods. First, regular prices are often “sticky” in the sense that adjustments are infrequent (Carlton 1986, Kashyap 1995, Blinder 1998, for example). Second, markups are “countercyclical”: They are lower in booms than in recessions (Bils 1987, Chevalier and Scharfstein 1996, Warner and Barsky 1995, and others). Third, temporary sales and promotions are common (Chevalier, Kashyap, and Rossi 2000).

We begin in Section 2 by outlining a fully specified model of consumer behavior with reference-dependent preferences. A major challenge in developing such a theory lies in the specification of

¹ In selling and purchasing decisions, subjects ask for a higher price when selling a good than they are willing to pay when offered the opportunity to buy the same good, even though the two roles are randomly allocated between subjects. This probably occurs because subjects construe giving up the object they have just acquired as a loss, and are very sensitive to this loss (Kahneman, Knetsch, and Thaler 1990, 1991). In monetary gambles, loss aversion is reflected in people’s reluctance to accept small favorable lotteries, presumably because they are more afraid of the possibility of loss than they are thrilled about a somewhat larger gain (Kahneman and Tversky 1979, Rabin 2000).

² Genesove and Mayer (2001) document that sellers in the Boston housing market set higher selling prices if they have suffered a loss relative to their purchase price. Odean (1998) finds that small investors are more reluctant to sell losing stocks than winning stocks, significantly decreasing their returns in the process. And as a possible explanation for wage stability, Bewley (1998) states that “[d]eclines in living standards disrupt people’s lives and cause much more damage to well-being than corresponding increases improve it. Such declines are especially disruptive for people used to stable incomes.”

³ Hall and Hitch (1939) write “[c]hanges in prices [...] are disliked by merchants and customers. Several entrepreneurs referred explicitly to the fact that there are conventional prices to which customers are attached, and that these have to be charged, which means that in these cases only large changes in price which are clearly unprofitable are possible.”

the reference point.⁴ In particular, one could specify a consumer’s reference point in the money to be paid for an item as either her status quo (so that paying is always assessed as a loss), her customary spending on the good (so that only paying more than at other times is seen as a loss), or the regular price of the good (so that only a price increase is considered a loss). To sort out these possibilities without having to make arbitrary exogenous assumptions in each situation, a parsimonious theory of pricing and loss aversion should ideally build on a sufficiently general and precise specification of the reference point.

Our model draws on the framework developed in Kőszegi and Rabin (2004), in which a decisionmaker’s reference point is determined by her recent expectations (i.e. probabilistic beliefs) about the outcomes she is going to get. For example, a consumer who expects to eat a nice dinner experiences a sensation of loss if she finds that her favorite restaurant is temporarily closed. But she also experiences a loss from paying more for her dinner than she had previously anticipated. Based on this perspective, a person’s reference point in money depends on market conditions and her own anticipated behavior, endogenously determining which of the above candidates (if any) applies. The reference point is close to the consumer’s status quo if she expects to buy the good relatively rarely, or to be able to acquire it very cheaply. If the firm always charges the same price and she expects to buy at that price, it is this “regular price” to which she compares money outlays. And if she expects to pay a stochastic price, she has no single reference price; instead, in our framework she compares money outcomes to each possibility in the reference lottery, and takes the average of these gain-loss sensations.

Having assumed that expectations are the reference point, our model’s predictions rely on a theory of how expectations are determined. To complete our model, we adapt Kőszegi and Rabin’s (2004) framework: We assume that the reference point is determined endogenously, in a *personal equilibrium*, by the requirement that the stochastic outcome implied by optimal behavior conditional

⁴ Indeed, the lack of a precise theory of reference point determination may have limited the application of previous models (Kahneman and Tversky 1979, Tversky and Kahneman 1991). Most applications of loss aversion are in the area of finance. In one of the first prominent financial applications, Benartzi and Thaler (1995) argue that loss aversion can explain the equity premium if investors check their portfolios (and suffer gains and losses) about once a year. Barberis, Huang, and Thaler (2003) formally show that within a large class of preferences, observed risk attitudes can only be explained by a combination of loss aversion and “narrow bracketing,” the idea that investors do not integrate current risky choices with risks they already face. Barberis, Huang, and Santos (2001) study asset prices in an economy populated by loss averse consumers.

on expectations be consistent with expectations.

In our full market models, a risk-neutral monopolist interacts with loss averse consumers of the above type. The monopolist faces an uncertain cost of production, and has to decide how to map different cost realizations into prices.⁵ To capture, in a reduced-form model, the idea that a patient firm facing rational consumers would realize that these consumers learn the price distribution over time, we begin by assuming that the monopolist commits to a cost-contingent pricing distribution before observing the cost realization.⁶ The consumer observes the pricing distribution while forming expectations about her own price-contingent behavior (but does not necessarily know the cost distribution or observe the realization). Next, the firm's cost is drawn, the price is realized, and the consumer decides whether or not to buy a single item of the good.

For most of the paper, we assume that the consumer's willingness to buy the good is subject to a shock, which she observes just before making her purchase decision. This assumption is both psychologically reasonable (many unpredictable factors may affect the consumer's decision) and technically useful, as it turns a discrete decision into a smooth demand curve. To maintain discipline, we impose a condition on the shock such that personal equilibrium is unique for any price distribution, doing so with the impression that—as long as jumps between equilibria in response to price changes are not a concern—this does not affect the intuition for any of our main results.⁷ We believe that this specification is a generally applicable model of demand with loss averse consumers.

Our first major result is that even if marginal costs are continuously distributed, the firm might charge only finitely many prices. We interpret this as price stickiness. Intuitively, random prices induce uncertainty for the consumer as to how much she has to pay for the good. If she purchases at a relatively high price, she compares it to lower prices she could have paid, and experiences a

⁵ While in our formal model we focus exclusively on cost shocks, we argue that our results would be similar if the underlying source of uncertainty was in inframarginal demand.

⁶ Although we work with a static model in which the firm faces a random cost, like previous authors (Shilony 1977, Varian 1980) we interpret our model more broadly, and translate the static pricing results into pricing over time. While we discuss in detail below how our static model, and especially our commitment assumption, relates to pricing over time, the natural way to think about it is that realizations of random price draws over time translate into a price series. Thus, a stochastic price corresponds to price variation that is not predictable very far in advance, such as fluctuations due to unexpected sales.

⁷ In particular, the key result in our no-commitment model does not take advantage of the uniqueness of personal equilibrium. Since in that case the firm chooses its price after the consumer has formed expectations, changes in the price cannot affect the personal equilibrium played by her.

monetary loss. The anticipation of this loss reduces her willingness to pay for the good. By insuring the consumer against small price shocks, therefore, the firm increases her overall willingness to pay and thus its sales revenues. If the cost distribution is sufficiently thick, this gain dominates the loss from being unable to differentiate production levels according to marginal cost.

In addition to formalizing the connection between loss aversion and price stickiness, our theory makes predictions about the circumstances in which price stickiness is likely to be observed. We show that if the cost distribution has sufficiently low density, the monopolist chooses a nonatomic price distribution. The firm is also more likely to do so if its consumers buy the product less often, and if demand is more elastic. And looking within a price distribution, the price tends to fall less often in a region where the distribution is nonatomic than where it is discrete. That is, paradoxically, the price is more likely to fall where *fewer* price realizations are observed.

Even when the monopolist's price distribution is nonatomic, it is observably different from that of a firm facing standard consumers. Namely, an economist who studies the firm's behavior and the realized market demand curve, but ignores that demand depends on the consumer's expectations, would conclude that the firm systematically deviates from profit maximization in a number of ways. Most importantly, she would find that the firm's markups are (too) countercyclical. When the firm decides how to price in low-cost states, it needs to take into consideration that the low price it would like to charge would form a basis for an (unfavorable) comparison should prices be higher, decreasing the consumer's demand in these latter states of the world. Therefore, the monopolist is reluctant to aggressively cut its price in response to low cost realizations.

The firm's ability to benefit from a sticky price seems to rely strongly on its patience in setting the consumer's price expectations, modeled as commitment to a strategy in our setting. Suppose instead that the firm cannot commit to a pricing strategy, and the consumer expects it to charge a nonatomic price distribution. Then, the demand curve is smooth, and the firm's price indeed varies continuously with cost, seemingly indicating that a market equilibrium with a continuous pricing distribution should always exist. In contrast, Section 4 shows that our price stickiness result survives: For sufficiently dense cost distributions the firm charges a deterministic price in any equilibrium. If the consumer expected the firm to charge stochastic prices, its demand would

be more responsive at prices at the upper end of the anticipated distribution. This would lead the firm to charge less dispersed prices than the consumer expects. As a result, the equilibrium collapses to a single price. At this price, the demand curve is kinked. Importantly, we obtain a kinked demand curve endogenously—it is an *equilibrium* phenomenon, not an exogenously assumed feature of the consumer’s preferences. Strikingly, we obtain this result despite a great indeterminacy of the equilibrium price: A continuum of *deterministic* prices are possible.

In Section 5, we consider the impact of the multiplicity of personal equilibria on market outcomes. We show that when the firm can commit to a pricing distribution, multiplicity can give rise to random pricing even for deterministic costs. We interpret this as strategic sales by the monopolist. Indeterminacy of personal equilibrium arises because if the consumer expects to consume the good, she assesses not buying it as a loss, increasing her willingness to pay for it; but if she expects not to buy, it is spending money that she views as unpleasant, decreasing her willingness to pay. At first, this may seem to imply that consumer loss aversion greatly benefits the firm if she plays the former equilibrium, and greatly hurts it if she plays the latter one. Surprisingly, the situation is not quite this symmetric, and sales play a crucial role in generating the asymmetry. If the customer tends to play an unfavorable equilibrium, the monopolist sometimes offers a sufficiently attractive sale that would induce her to buy even if she expected not to do so. This increases her anticipated probability of purchase, and leads her to buy even at higher prices. Loosely speaking, therefore, sales manipulate consumers into playing a personal equilibrium with a high probability of purchase. We show that this effect is possibly so powerful that in a market equilibrium, consumer loss aversion cannot hurt the firm very much. While we feel this is an important reason for sales in some situations, this variant of our model has the drawback that it can support a wide range of pricing policies and profits.

To summarize, our model offers an intuitive explanation for both price stickiness and countercyclical markups, and may also yield a natural reason for sales. In Section 6, we discuss some existing evidence for these seemingly contradictory predictions. We also review the related industrial organization literature, which has tended to explain these different phenomena with completely different models. We point out possible extensions of the model, and conclude, in Section 7.




Figure 1: Timing

2 Model

The timing of our basic model is illustrated on Figure 1. A monopolist, who interacts with a single (representative) consumer, first commits to a pricing distribution, which the consumer observes. We will think of the pricing distribution as a map from the cost realization to a price. For our model, however, it is only important that the consumer observes the resulting price distribution. Then, the consumer forms expectations about her own price-contingent behavior. Next, the firm’s marginal cost is drawn, and the price is realized. Finally, a shock to the consumer’s willingness to buy is realized, and she decides whether or not to buy a single item of the good.

Three comments about this basic setup are in order. First, although all of our concepts and definitions would make sense without this, we include in our model the psychologically reasonable element that there is randomness in the consumer’s willingness to buy the good. One can think of many factors at the time of purchase that impact the consumer’s decision—other demands on her time, her mood, the weather, and so forth—that are neither under the control of the firm, nor perfectly predictable by the consumer. Adding a shock to the consumer’s decision does not affect the key properties of loss aversion that our model captures, and it turns a binary decision into a (better-behaved) continuous demand curve. In addition, we will show that if the shock is sufficiently variable, personal equilibrium is unique for any price distribution, while this is not the case in general. Basing our entire analysis on (ad hoc) assumptions about equilibrium selection would raise the concern that our results are largely driven by these assumptions. To maintain discipline, therefore, for most of the paper we will restrict attention to settings where demand is

uniquely defined. But the intuition for our results holds more generally as long as relevant changes to the firm’s price distribution do not induce jumps in the personal equilibrium played by the consumer. Section 5 examines some market outcomes when personal equilibrium is not in general unique, and further motivates the inclusion of the shock in the model.

Second, although our model is static in that there is only a single purchase decision by the consumer, its logic is applicable to pricing decisions over time. A static price distribution generates a dynamic price series through different realizations of the distribution. Even if cost (and therefore price) realizations are correlated over time, if consumer preferences are based on sufficiently lagged expectations, the logic of the dynamic model would be very similar to our static one.

Third, more importantly, our commitment assumption captures, in a static reduced form, a patient firm’s dynamic incentives in setting the optimal long-run price distribution. A patient firm realizes that over time, consumers would learn the empirical distribution of prices and incorporate it into their expectations.⁸ As a result, such a firm internalizes the effect of its pricing decisions on consumer expectations. There are also firms and circumstances where pricing choices are more myopic due to discounting, financial distress, etc. Therefore, we later also analyze our model without commitment.

2.1 Individual Decisionmaking Problem

We first model the consumer’s preferences and behavior, and based on this specification of demand, we go on to define a full market equilibrium. The model of the consumer’s behavior is derived from Kőszegi and Rabin (2004). Specifically, the consumer makes the simple decision of whether to purchase a single item of a single good. We normalize her initial wealth to zero. We start by defining the consumer’s utility when both the outcome and the reference point are riskless, and without the shock to her willingness to buy. Denote by $k_1, r_1 \in \{0, 1\}$ her consumption and reference point in the good, and by $k_2, r_2 \in \mathbb{R}$ her consumption and reference point in money, and let $k = (k_1, k_2), r = (r_1, r_2)$. Total utility is composed of two additive terms: intrinsic consumption utility $m(k)$ and gain-loss utility $n(k|r)$. Intrinsic consumption utility is analogous to the standard

⁸ Indeed, in a standard setting, Fudenberg and Levine (1989) prove that outcomes in a repeated game with a long-run player are equivalent to the Stackelberg outcome in which the long-run player can commit to a strategy.

notion of utility in economics, according to which preferences depend only on absolute outcomes. We suppose that $m(k) = vk_1 + k_2$, so that the intrinsic value of the consumption good is v . Gain-loss utility captures reference dependence and loss aversion. Following Kőszegi and Rabin, its specification incorporates two key assumptions. First, the consumer evaluates gains and losses in the two dimensions separately. For example, if she unexpectedly receives the good and has to pay money for it, this is assessed as a gain in the good dimension and a loss in the money dimension (instead of, for example, a single gain or loss depending on the net intrinsic value of the transaction).⁹ Second, the consumer’s sense of gain or loss is directly related to the usefulness of the goods in question—it is more painful to lose something we value (e.g. \$100) than to lose something we do not (e.g. a paper clip). To capture these properties, we assume that $n(k|r) = \mu(vk_1 - vr_1) + \mu(k_2 - r_2)$. That is, gain-loss utility in a dimension depends in a universal way on the difference between the intrinsic utility associated with consumption in that dimension, and the intrinsic utility that would have been derived from the reference point. We formalize loss aversion in the simplest possible way: μ is two-piece linear, and has a slope of 1 for gains, and a slope of $\lambda > 1$ for losses.¹⁰ Hence, the consumer is more sensitive to losses relative to her reference point than she is to equal-sized gains over it.

In addition to the above sources of utility, the consumer experiences a shock \tilde{w} to her willingness to buy the good. We make two reasonable psychological hypotheses about how exactly \tilde{w} affects utility. First, since we believe that comparisons between the firm’s possible prices are crucial in determining the consumer’s sense of gain or loss, we suppose that \tilde{w} is *not* a shock to her perceived price of the good (e.g. due to a random transportation cost).¹¹ A methodological advantage of this assumption is that it avoids adding a shock to the very variable whose distribution is the central question of the paper.

⁹ This phenomenon is key to the endowment effect and other observed regularities in riskless trades. If gains and losses were defined over the value of an entire transaction, loss aversion would have no implications for such trades.

¹⁰ In general, we would want μ to satisfy all the properties of Kahneman and Tversky’s (1979) “value function”: continuity, convexity for losses, and concavity for gains, in addition to being kinked at zero. We focus on the case in which μ is two-piece linear as it suffices to capture the effects of loss aversion, our primary interest in this paper.

¹¹ If that was the case, her gain-loss utility in money would not depend solely on the firm’s chosen prices, but on the total (shock-inclusive) “effective price” distribution, creating a very different model. In particular, the firm would then have much less control over the total distribution. As a result, its incentive to absorb small cost shocks would be significantly mitigated, because doing so would have a negligible effect on the overall effective price distribution.

Second, we assume that the consumer does not experience gain-loss utility in the shock \tilde{w} . While many random factors can influence whether the consumer buys in the end, we feel that gain-loss comparisons in these variables are likely to be much less salient and powerful than those about the good and its price, at least as far as the decision to purchase the good are concerned. Incidentally, this assumption is also convenient for technical reasons—our analysis below is made much simpler if we do not have to worry about gains and losses in this extra aspect of utility.¹²

These assumptions yield the following simple expression for the consumer’s utility when her reference point is r , her consumption is k , and she experiences the shock \tilde{w} : $u(k, \tilde{w}|r) \equiv m(k) + n(k|r) + k_1 \cdot \tilde{w}$. We assume that \tilde{w} is continuously distributed on the bounded interval (a, b) , with positive density everywhere. The cumulative distribution function H is twice continuously differentiable and H'' is bounded. Denote the probability distribution function by h .

Since we will suppose that the consumer’s reference point is her lagged probabilistic beliefs about what she is going to get, it is crucial to extend the above specification to allow for the reference point to be stochastic. We assume that if the consumer’s reference point is the probability measure Γ over $\{0, 1\} \times \mathbb{R}$, her utility is

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

This formulation captures, in a tractable way, the idea that in evaluating an outcome k , the consumer compares it to each possible consumption level in the reference lottery. For example, if she expects to receive \$100 with probability one-half, coming up empty-handed feels partly neutral (relative to the possibility of not getting \$100) and partly like a loss (relative to the possibility of getting \$100).

Having specified the consumer’s utility function, we turn to modeling her behavior. Suppose she faces a probability distribution F of non-negative prices, and can decide whether to buy at each price. Let $\sigma : \mathbb{R}_+ \times (a, b) \rightarrow [0, 1]$ be her strategy, which assigns a probability of buying to each price-shock pair. Our model of behavior is based on the premise that the consumer’s preferences,

¹² Nevertheless, we have derived some properties of demand when the noise is also subject to gain-loss comparisons, and the effect of this on our results would depend on the exact psychological source of the uncertainty. If the unpredictability was in the consumer’s utility from *not* buying the good, the effects on which our results rely would seem to strengthen. If, on the other hand, the shock was to the consumer’s valuation of the good, the same effects would seem to be weakened.

and therefore also the strategy σ , depend on expectations she forms before she finds out the realized price. We also believe that in most situations of interest for this paper, consumers would have at least some ability to predict their own behavior. We capture this notion by assuming rational expectations, positing that the consumer's reference point is the distribution $\Gamma_{\sigma,F}$ induced by σ and F over consumption good-money pairs. To deal with the resulting interdependence between behavior (σ) and expectations ($\Gamma_{\sigma,F}$), we use the personal equilibrium concept from Kőszegi and Rabin (2004), which requires the strategy that generates expectations to be optimal conditional on these expectations:

Definition 1. $\sigma : \mathbb{R}_+ \times (a, b) \rightarrow [0, 1]$ is a personal equilibrium for the price distribution F if for all $p \in \mathbb{R}_+$, $\tilde{w} \in (a, b)$,

$$\sigma(p, \tilde{w}) \in \operatorname{argmax}_{s \in [0,1]} s \cdot U(1, -p, \tilde{w} | \Gamma_{\sigma,F}) + (1 - s) \cdot U(0, 0, \tilde{w} | \Gamma_{\sigma,F}).$$

Consider, for example, a consumer's decision of whether to buy a nice stereo. Suppose for a moment that she expects to always buy the stereo and, with these expectations determining her preferences, there are values of \tilde{w} for which she prefers not to buy. Then, if she can predict her own behavior, she should not have expected to buy with probability one in the first place. Definition 1 captures the notion that her expectations must be consistent with future behavior. See Kőszegi and Rabin (2004) and Kőszegi (2004) for further details on this solution concept.

2.2 Market Equilibrium

Now consider the entire market game illustrated in Figure 1. Suppose the firm has to set non-negative prices. To make our results in the next sections most striking, we suppose the firm's cost is continuously distributed on the interval $[\underline{c}, \bar{c}]$, with continuous and positive density everywhere. Denote the cumulative and probability distribution functions of c by $\Theta(\cdot)$ and $\theta(\cdot)$, respectively.

In this environment, we define the firm's pricing strategy as a function from marginal cost to price. Clearly, any pricing strategy $P : [\underline{c}, \bar{c}] \rightarrow \mathbb{R}_+$ chosen by the firm induces a distribution of prices F_P faced by the consumer. For a strategy σ , let $\bar{\sigma}(p) \equiv \int_a^b \sigma(p, \tilde{w}) dH(\tilde{w})$; $\bar{\sigma}(p)$ is the consumer's probability of purchasing the good at price p .

Definition 2. The pricing strategy $P(\cdot)$ and the strategies $\sigma_F : \mathbb{R}_+ \times (a, b) \rightarrow [0, 1]$ for all price distributions F constitute a pricing equilibrium if

1. For any price distribution F , σ_F is a personal equilibrium for F .
2. $P(\cdot)$ maximizes the firm's expected profits:

$$P \in \operatorname{argmax}_{\hat{P}(\cdot)} \int_0^\infty (\hat{P}(c) - c) \cdot \bar{\sigma}_{F_{\hat{P}}}(\hat{P}(c)) \, d\Theta(c).$$

A pricing equilibrium is a situation in which (1) for any pricing strategy chosen by the firm, the consumer plays a personal equilibrium in the “continuation game”; and (2) the firm chooses an optimal pricing distribution correctly anticipating the continuation play of the consumer.

The sole source of uncertainty that drives price changes in this model and in our analysis below is uncertainty about marginal cost. An alternative specification is a model where the firm has unchanging marginal cost, but there are observable shocks to inframarginal demand. Specifically, suppose the firm's profits at price p drawn from the price distribution F are given by $p \cdot (\tilde{a} + \bar{\sigma}_F(p))$, where $\tilde{a} \leq 0$ is a random variable that shifts demand, and that is observed by the firm (such as the state of the economy). We will argue below that our results are robust to such a change in the source of underlying uncertainty.

2.3 The Demand Side—An Example and General Techniques

We first consider consumer behavior in our model. To illustrate the use of our personal equilibrium concept, and to demonstrate that a shock to the consumer's decision is necessary for it to be unique, we start with a simple multiple equilibrium example taken from Kőszegi and Rabin (2004). Let $F(\cdot)$ be the deterministic price p , take $v = 1$, and suppose \tilde{w} is degenerate and equal to zero. Then, if the consumer expects to buy the good $((r_1, r_2) = (1, -p))$, her utility from buying is $1 - p$, and her utility from not buying it is $p - \lambda$. Thus, she will buy if $p \leq \frac{1+\lambda}{2} \equiv p_{max}$. If she expects not to buy $((r_1, r_2) = (0, 0))$, her utility from buying is $1 - p - \lambda p + 1$, and her utility from not buying is zero. Thus, she does not buy if $p \geq \frac{2}{1+\lambda} \equiv p_{min}$. For $p \in [p_{min}, p_{max}]$, there are thus multiple personal equilibria.

The intuitive reason for multiple personal equilibria is simple. If the consumer expects to buy the good, not buying is assessed as a gain in the money dimension and a loss in the good dimension. Since the loss from not being able to consume is more painful than the gain from saving money is pleasant, it is optimal for her to buy. But in case she expects not to buy, it is spending money that she considers unpleasant, so it is optimal for her not to buy.

Solving for personal equilibria with non-degenerate price and noise distributions is considerably more complex. To identify some of the main forces that will be driving our results, suppose the consumer expects to buy the good with probability q , and she expects the distribution of buying price conditional on buying to be G . Then, her utility from buying at price p is

$$v + \tilde{w} - p - \lambda q \int_0^p (p - p') dG(p') + q \int_p^\infty (p' - p) dG(p') + (1 - q)(v - \lambda p). \quad (2)$$

The utility from not buying is

$$q \left[\int_0^\infty p' dG(p') - \lambda v \right], \quad (3)$$

which can be rewritten as

$$qp + q \int_0^p (p' - p) dG(p') + q \int_p^\infty (p' - p) dG(p') - q\lambda v. \quad (4)$$

Hence, the difference in utility between buying and not buying is

$$v + \tilde{w} - p + (1 - q)(v - \lambda p) + q(\lambda v - p) - (\lambda - 1)q \int_0^p (p - p') dG(p'). \quad (5)$$

Expression 5 will be a key starting point to many of our results below. The term $v + \tilde{w} - p$ is the intrinsic consumption utility associated with purchasing the good relative to not purchasing it. The other terms derive from gain-loss utility. The consumer expected to pass on the good with probability $1 - q$, relative to which buying results in a gain of the value v in the good dimension, and a loss of the price p in the money dimension. This is captured in the term $(1 - q)(v - \lambda p)$. On the other hand, the consumer expected to buy with probability q , relative to which buying avoids a loss of v and leads to a foregone gain of p . Thus the term $q(\lambda v - p)$. These last two terms imply that an increase in the consumer's anticipated probability of buying (q) increases her "attachment" to (net utility from buying) the good and thus her endogenous willingness to pay.

The last, and perhaps most interesting, term in Expression 5 captures the consumer’s sense of loss from comparing p with lower prices she may have received under the firm’s pricing distribution. Contrasts with these “low-price” states induce a sensation of loss in the money dimension if she buys at price p , whereas they merely reduce the gains she feels if she does not buy. As a result, these comparisons decrease her net utility from buying. At first sight, it may seem that Expression 5 should also include a term reflecting the consumer’s comparison between buying at price p and buying at *higher* prices in the distribution. While this comparison affects the consumer’s utility, it does not affect her decision of whether to buy, because it appears as a gain in both her utility from buying (Expression 2) and her utility from not buying (Expression 3). Intuitively, whether or not she buys, she saves money relative to states in which the good is more expensive than p .¹³

Expression 5 implies that, unlike in the standard model of consumer behavior, demand at a price p (and thus the demand curve) depends on the entire distribution F of prices. We now show how to derive the demand curve, and give a sufficient condition for it to be unambiguously determined. Based on this analysis, it will be possible to translate intuitions derived from examining Expression 5 into formal proofs of results about pricing equilibrium. Therefore, the role of the rest of this section is to provide techniques that are necessary to follow our proofs. We also feel that these techniques are useful for future market models with loss averse agents, and for that reason view them as an important contribution of this paper. But they are not necessary for an intuitive understanding of our main results.

Suppose the price distribution is F , the consumer expects to buy with probability q , and the distribution of conditional buying prices is G . Then, her strategy is described by a function $w_{F,q,G} : \mathbb{R}_+ \rightarrow \mathbb{R}$, where $w_{F,q,G}(p)$ is defined as the value of \tilde{w} at which Expression 5 is zero. Namely, she buys the good ($\sigma_F(p, \tilde{w}) = 1$) if $\tilde{w} > w_{F,q,G}(p)$, and does not buy ($\sigma_F(p, \tilde{w}) = 0$) if $\tilde{w} < w_{F,q,G}(p)$. Note that (given F) q and G fully determine $w_{F,q,G}$. For this to describe a personal equilibrium strategy, it must generate expectations consistent with q and G . Thus, a necessary and

¹³ This feature of our model depends on the linearity of μ in gains. With diminishing sensitivity, we would get the plausible prediction that expecting to buy at higher prices increases the consumer’s willingness to pay at a given price. Buying at a lower price than expected makes the consumer happy due to the gain she makes relative to the high price, and because of diminishing sensitivity, she would care less about the extra gain she could make by not buying. This is probably a real and important effect, and may have some interesting consequences for pricing. But it is unlikely to alter the results in this paper.

sufficient condition for personal equilibrium is that

$$q = \int_0^\infty (1 - H(w_{F,q,G}(p)))dF(p), \text{ and} \quad (6)$$

$$G(p) = \frac{1}{q} \int_0^p (1 - H(w_{F,q,G}(p')))dF(p') \quad \forall p. \quad (7)$$

In principle, solving for $w_{F,q,G}(p)$ requires us to look for a fixed point in a function space. Instead of solving this more complicated problem directly, we proceed by simplifying it to a fixed point problem in the consumer's probability of buying (q), a fixed point in the reals. More specifically, we take G out of the above fixed-point problem by showing that q in itself fully determines $w_{F,q,G}$.

Indeed, notice that through Expression 5, q determines $w_{F,q,G}(0)$ independently of G . Economically, at price $p = 0$, the consumer's monetary outcome, including comparisons to other possible prices, is the same whether or not she buys. To proceed, we exploit the property of our model that the comparison with prices greater than p does not affect the consumer's purchase decision at price p . Obviously, $w_{F,q,G}(p)$ is increasing in p . Thus, it is differentiable almost everywhere. Differentiating 5 with respect to p and using 7 gives

$$\frac{\partial w_{F,q,G}}{\partial p} = 2+(\lambda-1)(1-q)+(\lambda-1)qG(p) = 2+(\lambda-1)(1-q)+(\lambda-1) \int_0^p (1-H(w_{F,q,G}(p')))dF(p'). \quad (8)$$

Now, using that $w_{F,q,G}(p)$ is continuous in p , we can conclude (by the Fundamental Theorem of Calculus) that q determines $w_{F,q,G}$ entirely. Henceforth, we will thus denote the function by $w_{F,q}(\cdot)$. Intuitively, we can “build up” the function $w_{F,q}(\cdot)$ using only the consumer's overall probability of purchase. Since only comparisons with lower prices matter for the purchasing decision, q determines $w_{F,q}$ at zero and at prices “very close” to zero. Since again only comparisons with lower prices matter, this determines $w_{F,q}$ at slightly higher prices also, and so on.

Given the above simplifications, $w_{F,q}$ characterizes a personal equilibrium if and only if

$$q = \int_0^\infty (1 - H(w_{F,q}(p)))dF(p). \quad (9)$$

Since the right-hand side of Equation 9 defines a continuous map from $[0, 1]$ to itself, it has a fixed point, and therefore personal equilibrium exists. The following assumption and proposition establish that with “sufficient noise”, personal equilibrium is unique for any price distribution.

Assumption 1. For all $\tilde{w} \geq -(1 + \lambda)v$,

$$h(\tilde{w}) < \frac{1}{(\lambda - 1) \left(v + \frac{\tilde{w} + (1 + \lambda)v}{2} \right)}. \quad (10)$$

Proposition 1. (Unique Demand) *If Assumption 1 holds, then for any pricing distribution F , the personal equilibrium is unique.*

Noise in the consumer’s purchase decision weakens the feedback on which the existence of multiple equilibria hinges. Since the decision to buy depends on the realization of \tilde{w} in addition to expectations, a change in the anticipated probability of buying leads to a relatively small change in the actual probability of buying. Thus, the change in expectations cannot be self-fulfilling.¹⁴

3 Price Stickiness and Countercyclical Pricing

We are finally ready to derive our pricing results. If $q^*(F)$ is the equilibrium probability of buying, call $w_{F, q^*(F)}(p)$ —the \tilde{w} at which the consumer is indifferent between buying and not buying at price p —the “critical shock” at p . To guarantee that the firm can profitably sell to the consumer with positive probability even at the highest cost, we suppose that $\bar{c} < \frac{b+2v}{1+\lambda}$. And to simplify many of our proofs, from now on we impose the following condition on H :

Assumption 2. (Regularity Condition) H satisfies the monotone hazard rate condition: $\frac{h(\tilde{w})}{1-H(\tilde{w})}$ is increasing in \tilde{w} on (a, b) , and approaches ∞ as \tilde{w} approaches b .

Assumption 2 is a technical assumption; if the consumer had no gain-loss utility, it would be sufficient for the firm’s profit maximization problem to be well-behaved.¹⁵

We first establish that the firm’s maximization problem has a solution, and show that for any optimal pricing distribution, the probability of buying is uniformly bounded away from zero.

¹⁴ Notice that if λ increases, the right-hand side of Condition 10 decreases. Thus, we require “more noise” to guarantee a unique equilibrium. This is natural given the above logic that multiple equilibria rely on loss aversion. And as $\lambda \rightarrow 1$, equilibrium becomes unique even without noise.

¹⁵ That is, it guarantees that there is a solution to the firm’s pricing problem, and that the firm’s profits are concave in quantity. With no gain-loss utility, the price at which the firm sells quantity q is $p(q) = H^{-1}(1 - q) + v$. Thus, $p'(q) = \frac{-1}{h(p(q))}$ and $p''(q) = \frac{h'(p(q))p'(q)}{h(p(q))^2}$. As $q < 1$, $-\frac{h'(\tilde{w})}{h(\tilde{w})^2} > -\frac{qp''(q)}{p'(q)}$.

Lemma 1. *Under Assumptions 1 and 2, a pricing equilibrium exists. Moreover, there are $a \leq \underline{w} \leq \bar{w} < b$ such that for any p on the support of an optimal pricing distribution, the critical shock at p lies on the interval $[\underline{w}, \bar{w}]$.*

We turn to presenting our first price stickiness result. Since this form of the result is easiest to state, in the proposition we give a *global* statement, conditions under which discrete pricing does or does not obtain over the entire support of prices. But we also provide and use a *local* statement (conditions for price stickiness for a range of price levels) in the text and interpretation.

Proposition 2. (Price Stickiness) *Suppose Assumptions 1 and 2 hold.*

Part I. If for all $c \in [\underline{c}, \bar{c}]$ and all $\tilde{w} \in [\underline{w}, \bar{w}]$,

$$\theta(c) \left(1 + \frac{1 - \Theta(c)}{1 + \lambda} \right) > \frac{(\lambda + 1)^2}{\lambda - 1} \cdot \frac{h(\tilde{w})}{(1 - H(\tilde{w}))^2}, \quad (11)$$

then in a pricing equilibrium, the set of prices the firm sets with positive probability is finite.

Part II. If for all $c \in [\underline{c}, \bar{c}]$ and all $\tilde{w} \in [\underline{w}, \bar{w}]$,

$$\theta(c) \left(1 + \frac{1 - \Theta(c)}{2} \right) < \frac{4}{\lambda - 1} \cdot \frac{h(\tilde{w})}{(1 - H(\tilde{w}))^2},$$

then in a pricing equilibrium, the firm's pricing distribution is nonatomic.

Although our formal result is about pricing in a static model, we can interpret it in terms of the monopolist's price series over time. Under that interpretation, Part I means that the monopolist's prices form a step function. In other words, even though its cost is continuously distributed, the firm does not change its price for most small cost changes. And when it does adjust the price, it does so in discrete steps, even if the cost change that triggers the price adjustment is very small.

The key force behind our price stickiness result is the consumer's loss aversion in money: Her probability of buying the good at price p is negatively affected by the comparison of p to lower prices in the distribution. In particular, if the distribution of prices were continuous, comparing p to slightly lower possible prices would decrease the consumer's demand in the region. In order to avoid this unfavorable comparison, the firm has an incentive to "lump together" nearby prices into a single one. If this incentive outweighs the firm's benefit from differentiating production according to cost, the price distribution is discrete.

Proposition 2 also illustrates, however, that the firm does not always want to absorb small cost shocks. The proposition makes it clear that four factors combine to determine whether price stickiness occurs.¹⁶ First, the firm’s cost distribution needs to be “sufficiently dense” for the price distribution to be discrete. If the cost is highly variable, in the optimal pricing strategy the probability that the price is within a narrow range is small. Thus, comparisons between prices in these states play a relatively small role in determining the consumer’s overall sensation of loss in money. As a result, the firm does not benefit as much from lumping together nearby prices, and so chooses to differentiate production according to cost.

Second, the firm’s tendency to set sticky prices depends on the consumer’s loss aversion, measured by λ . As the buyer’s loss aversion disappears ($\lambda \rightarrow 1$), the firm sets a non-atomic price distribution in equilibrium. Intuitively, if the consumer does not find losses very painful, the unfavorable comparison of a price p to lower prices in the distribution has little impact on her demand, so the firm does not benefit much from eliminating such comparisons. For non-trivial loss aversion, Condition 11 is more likely to be satisfied, reflecting the firm’s increased incentive to absorb small cost shocks. Our conditions, however, leave substantial ambiguity as to whether price stickiness occurs for high λ ’s. The reason is that an increase in λ not only increases the firm’s revenue-side incentive to absorb small cost shocks, it also increases its production-side costs of doing so,¹⁷ and neither effect seems to systematically dominate for large λ .

To interpret the other determinants of price stickiness, we provide a local analogue of Condition 11 in terms of the properties of the agent’s equilibrium demand curve (which depends on F). Let

$$D_F(p) \equiv 1 - H(w_{F,q^*(F)}(p)).$$

Consider an interval (p_1, p_2) of an optimal pricing distribution F , and let $[c_1, c_2]$ be the closure of the costs corresponding to prices in (p_1, p_2) . By the same argument as in the proof of Proposition

¹⁶ The arguments below on the properties of the model as a function of Θ and λ are not completely precise, because changes in the parameters in general change \underline{w} and \bar{w} . But it is easy to show that \underline{w} and \bar{w} can be chosen so that the statements hold.

¹⁷ As λ increases, the elasticity of demand typically also increases, as the consumer becomes more sensitive to losses in the money dimension. Thus, as explained below, the firm benefits more from differentiating production according to marginal cost.

2, if for all $p \in (p_1, p_2)$, $c \in [c_1, c_2]$,

$$\frac{(\lambda - 1)D_F(p)}{\lambda + 1} \cdot \left(1 + \frac{1 - F(p)}{\lambda + 1}\right) > \frac{1}{\theta(c)} \cdot \frac{D'_F(p)}{D_F(p)} \quad (12)$$

whenever $D'_F(p)$ exists, then the pricing distribution on (p_1, p_2) is discrete.¹⁸

Third, then, whether the firm chooses to absorb small cost shocks depends on the consumer's demand $D_F(p)$ (the probability with which she buys the good). If she is very unlikely to buy, her gain-loss utility from buying it is dominated by comparing this to not buying, and not by comparing the price to other possible purchase prices. Thus, the monopolist gains little from offering the same price for nearby cost levels.

Crucially, this comparative static depends on an *individual* consumer's frequency of purchase, not the overall frequency of purchase in the population. We can distinguish two kinds of firms. One firm has a small customer base of people who buy relatively frequently. The other firm sells the same amount overall, but to a larger population who buy less frequently. According to our model, the former firm would have a greater incentive to absorb small cost shocks.

Fourth, the firm is more likely to absorb small cost changes if the consumer's proportional responsiveness to prices—as measured by $\frac{D'_F(p)}{D_F(p)}$ —is small. If consumers are very responsive to price changes, the firm can differentiate production levels in different states of the world without charging very different prices in them. Price comparisons being small, the cost saving comes with little impact on the revenue side.

An interesting consequence of Condition 12 is that price stickiness is more likely to occur in regions where the price is more likely to fall. That the price is often in a given region is an indication that the cost distribution is quite dense there ($\theta(c)$ is high), making it more likely that the condition is satisfied. Loosely speaking, this means that if many price realizations are observed in a region of possible prices, the price is unlikely to fall there. Paradoxically, the price more often falls in a region where only a single isolated price has been observed.

¹⁸ The condition for the pricing distribution to be nonatomic on (p_1, p_2) is

$$\frac{(\lambda - 1)D_F(p)}{2} \cdot \left(1 + \frac{1 - F(p)}{2}\right) < \frac{1}{\theta(c)} \cdot \frac{D'_F(p)}{D_F(p)}.$$

Condition 12 also includes the term $1 + \frac{1-F(p)}{\lambda+1}$, indicating that price stickiness is more likely to occur at prices that fall in the lower part of the price distribution. Absorbing small cost shocks allows the firm to raise its average price while holding demand constant. This exerts a positive externality on profits in higher-price states, because it decreases the consumer’s sense of loss from paying the price in those states. And the lower the current price is in the distribution, the more important is this “externality effect”.¹⁹

The same intuitions for price stickiness apply to the alternative specification of our model in which the underlying uncertainty is about the amount of inframarginal demand. The firm increases revenues from the marginal consumer by lumping prices together into atoms. Lumping these prices together, however, comes at the cost of decreasing revenues from inframarginal consumers, because there is a benefit from charging higher prices for higher demand realizations.

Even in circumstances where the firm does not absorb small cost shocks, its price distribution is different in an interesting way from that of a firm facing standard consumers. Of course, since a standard firm faces consumers who have different preferences, it is not immediately clear how to make such a comparison. But there is a natural way to do it using market observables. The equilibrium demand function $D_F(\cdot)$ is (at least in principle) observable, and if F is continuously distributed, it is differentiable. A classical firm facing this demand function would solve

$$\max_p (p - c)D_F(p)$$

for any c . Let $L(p, c) \equiv \frac{p-c}{p}$ be the Lerner Index and $\eta_F(p)$ the absolute value of the demand elasticity. Then we can write the first-order condition for the optimal price at cost c as $L(p, c) - \frac{1}{\eta_F(p)} = 0$. This condition is the result of a classic tradeoff a firm faces: By raising its price, it makes money on its inframarginal consumers, but loses its profits on the marginal ones. At the optimum, these two effects cancel. The same is not true when consumers are loss averse:

Proposition 3. (Countercyclical Markups) *Suppose Assumptions 1 and 2 hold. In a pricing equilibrium with a nonatomic price distribution, $p \cdot (L(p, c) - [1/\eta_F(p)])$ is strictly decreasing in c .*

The intuition for this result derives from a variant of the externality effect we have discussed

¹⁹ We are not aware of any empirical work that tests some of these more specific predictions of the model.

above. The firm’s temptation to decrease prices and reap higher profits at lower cost realizations is constrained by the concern that a low price forms the basis for an unfavorable comparison when the consumer sees a higher price. In other words, lowering the price for a low cost may increase profits in this state of the world, but it decreases demand in the states when the firm needs to charge higher prices. And the lower is a cost realization, the more likely that the firm would want to charge higher prices at other times, so the less aggressively it wants to cut prices.

This intuition, however, ignores another kind of externality that affects the firm’s price distribution. By increasing demand at any price, the monopolist increases the consumer’s anticipated probability of buying, increasing her “attachment” to the good and therefore demand at all other prices. The above logic is correct only once this consideration is held constant, giving rise to the specific expression in Proposition 3. The expression is equivalent to $\frac{(p-c)D'_F(p)+D_F(p)}{-D'_F(p)}$, the firm’s marginal profit from increasing the price, discounted by the effect of the change in price on demand.

Using the empirical observation that marginal costs move procyclically (Bils 1987), a natural interpretation of the above proposition is that the firm’s prices follow a more countercyclical pattern than those of a monopolist facing standard consumers. But the exact same intuition suggests that this result also obtains in the variant of our model in which the underlying uncertainty is about the amount of inframarginal demand. Thus, our model predicts a tendency towards countercyclical markups whether or not business cycles originate in cost or demand shocks.²⁰

Although Proposition 3 is stated for continuous price distributions, an appropriate analogous statement is true in cases when the firm sets a finite set of prices. In that case, however, the equilibrium demand curve is kinked, so the derivatives used to define the price elasticity do not exist everywhere. The proposition remains true if the derivative $D'_F(p)$ in the definition of the elasticity is replaced by the *left* derivative, and the cost at a kink is replaced by the *average* cost.

²⁰ This discussion ignores some serious conceptual issues in connecting our model to macroeconomic settings. First, an open issue is whether loss aversion should be defined over real or nominal prices. Second, if booms and recessions affect gain-loss utility — e.g if wealth changes associated with a boom mitigate potential losses — business cycles cannot fully be modeled as changes in marginal cost or inframarginal demand. Addressing these and other issues is beyond the scope of the current paper.




Figure 3: Timing without Commitment

4 The Role of Commitment and the Optimal Price

Our formal analysis so far has assumed that the firm can commit to a pricing strategy before observing its cost realization. This framework applies mostly to firms that are in the market for the long run and therefore “invest” in influencing consumer expectations. If either consumers do not use past prices to update beliefs, or the firm heavily discounts future profits, the monopolist has no incentive to invest in changing expectations. Therefore, we now look at market outcomes when the firm is unable to commit to a price distribution, and instead chooses its price after observing its cost, and after the consumer has formed expectations about market prices and her behavior.

The timing of the game is now slightly different from Section 2, and is illustrated in Figure 3. The consumer first forms expectations about market prices and her own behavior in response to those prices. Then, the firm observes its cost and chooses its price. Next, the consumer observes \tilde{w} and decides whether or not to buy the good. A pricing equilibrium is defined as a situation where (1) for each c , the firm maximizes profits, given the consumer’s expectations; and (2) the consumer plays a personal equilibrium, and correctly anticipates the distribution of prices.

We first establish that equilibrium exists in this version of our model.

Lemma 2. *Suppose Assumptions 1 and 2 hold. In the no-commitment model, a pricing equilibrium exists. For any pricing equilibrium, there are $a \leq \underline{w} \leq \bar{w} < b$ such that for any p on the support of the pricing distribution, the critical shock at p lies on the interval $[\underline{w}, \bar{w}]$.*

It is easy to see that even without commitment, it can be in the firm’s interest to absorb small cost shocks. Suppose that the consumer expects the firm to charge price p with probability one,

and let the equilibrium probability of purchase with this expectation be $q^*(p)$. Then, the firm's profits from charging price p' at cost c are $(p' - c) \cdot [1 - H(w_{p,q^*(p)}(p'))]$, where (by Expression 5) the critical shock is given by

$$w_{p,q^*(p)}(p') = \begin{cases} 2p' - 2v + (1 - q^*(p))(\lambda - 1)p' - q^*(p)(\lambda - 1)v & \text{if } p' \leq p \\ 2p' - 2v + (1 - q^*(p))(\lambda - 1)p' - q^*(p)(\lambda - 1)v + q^*(p)(\lambda - 1)(p' - p) & \text{if } p' > p. \end{cases}$$

Given the firm's maximization problem, if

$$\frac{1 - H(w_{p,q^*(p)}(p))}{h(w_{p,q^*(p)}(p)) \cdot (2 + (\lambda - 1))} \leq p - c \leq \frac{1 - H(w_{p,q^*(p)}(p))}{h(w_{p,q^*(p)}(p)) \cdot (2 + (1 - q^*(p))(\lambda - 1))} \quad (13)$$

for all $c \in [\underline{c}, \bar{c}]$, the firm charges price p for all cost realizations. Intuitively, if the consumer expects to see the price p when she gets to the store, there is a kink in her demand function at p due to the loss she experiences if the price (to her surprise) turns out to be greater than p . Because demand is more sensitive to increases in the price than to decreases of it, for a range of cost levels p becomes the optimal price for the monopolist to charge.

This logic also indicates that if there is a deterministic-price pricing equilibrium, there are typically multiple ones. In fact, the set of possible equilibrium price levels is a closed interval.²¹ The lack of commitment reintroduces multiple equilibria into our model, because a kinked demand curve creates an incentive for the firm to “comply” with the consumer's expectations—expectations she forms without observing a pricing policy.

Although the above argument demonstrates that the firm may have an incentive to set sticky prices even without commitment, it may seem that there should also be equilibria with random pricing. After all, if the consumer expects the firm to charge different prices for different cost realizations, there is no kink in her demand function. Then, a profit-maximizing firm does in fact charge different prices for different cost realizations. Nevertheless, the following proposition shows that if the cost distribution is dense, there are only deterministic-price pricing equilibria.

Proposition 4. (Price Stickiness without Commitment) *Suppose Assumption 2 holds.*

²¹ To see this, first note that $q^*(p)$ is decreasing in p and $w_{p,q^*(p)}(p)$ is increasing in it, and both are continuous. Together with the Regularity Condition, this implies that both the left-hand side and right-hand side in Condition 13 are decreasing and continuous in p . Since $p - c$ is increasing and continuous in p , the set of prices for which Condition 13 is satisfied must be a closed interval.

Part I. If for all $c \in [\underline{c}, \bar{c}]$ and $\tilde{w} \in [\underline{w}, \bar{w}]$,

$$\theta(c) > \frac{(\lambda + 1)^2}{\lambda - 1} \cdot \frac{h(\tilde{w})}{(1 - H(\tilde{w}))^2},$$

then in any pricing equilibrium, the firm's price is deterministic.

Part II. If for all $c \in [\underline{c}, \bar{c}]$ and $\tilde{w} \in [\underline{w}, \bar{w}]$,

$$\theta(c) < \frac{4}{\lambda - 1} \cdot \frac{h(\tilde{w})}{(1 - H(\tilde{w}))^2},$$

then in any pricing equilibrium, the firm's prices are non-atomic.

Strikingly, even without commitment, and even though there are generally a continuum of equilibria in this case, with a sufficiently dense cost distribution the firm sets a deterministic price. The intuition for this result is easiest to see by first assuming that the consumer expects the firm's prices to be continuously distributed. Recall that the consumer's willingness to pay is negatively affected by comparisons to lower prices. The higher the price, the more she compares this price to lower prices and, thus, the greater is the reduction in her willingness to pay in response to a price increase. Hence, the firm's demand curve is relatively inelastic for prices lower than the range of prices the consumer expects, but then becomes more elastic as we move through the price distribution. Because of this, a firm with a sufficiently dense cost distribution offers a narrower price distribution than the consumer expects (although this price distribution is still continuous). But if it is this denser price distribution that the consumer expects, her demand function becomes elastic more quickly. The firm responds to this by condensing its price distribution further, and so on ... the equilibrium collapses to a single point.

Our model thus endogenously generates a market with a kinked demand curve. This feature of markets was emphasized (for completely different reasons) in some early models of oligopolistic competition, but was later dismissed as poorly founded game theoretically (Tirole 1988, page 243). Loss aversion may be a foundation for kinked demand curves.

It bears emphasizing that our theory *does not assume* a kinked demand curve from the start, but instead derives it endogenously as a property of market equilibrium. As we have mentioned, if consumers expected the firm to charge different prices for different cost realizations, then the

firm’s demand function would be a smooth one. Neither is the kink in demand a straight-forward consequence of reference-dependent models in general. For example, if there is a continuum of consumers with different reference points, the aggregate demand curve is again smooth, although individual ones are not. The assumption that expectations are the reference point, together with the equilibrium conditions, yield a deterministic price and create a kink in demand.

The conditions for price stickiness in Proposition 4 are very similar to those in Proposition 2, but do not include a term with $1 - \Theta(c)$. The reason is that without commitment, there is no “externality effect”: For any given cost realization, the firm ignores the effect of its pricing decision on demand in other states of the world. All our other comparative statics results, however, continue to hold. Namely, price stickiness is more likely to occur when the cost distribution has high density, for goods that the consumer buys with higher probability, and for more elastic goods. Furthermore, these comparative statics statements hold globally as well as locally.

An interesting difference between Propositions 2 and 4, however, is that the latter does not require uniqueness of personal equilibrium (Assumption 1). In Proposition 2, uniqueness is necessary so that changes in the firm’s pricing distribution do not induce jumps in the personal equilibrium played by the consumer. In the current model, the firm chooses its price after the consumer has formed expectations, so such jumps cannot happen.

To summarize, the firm has similar incentives to stabilize the price with and without commitment. Since the firm in the no-commitment model is assumed to maximize its profit given the (equilibrium) market demand curve, however, there is no analogue to Proposition 3. In addition, the two models are in general different in the price *levels* they generate. When the price is deterministic with and without commitment, there is an unambiguous relationship between the two price levels.

Proposition 5. (Price Comparison) *Suppose Assumptions 1 and 2 hold. If the deterministic price p_c is charged in a pricing equilibrium with commitment, and the deterministic price p_{nc} is charged in a pricing equilibrium without commitment, then $p_c < p_{nc}$.*

The proof of Proposition 5 proceeds by showing that if the consumer expects to be offered the price p_c with probability one, a non-committed firm prefers to charge a higher price. This is equiva-

lent to showing that the firm’s incentive to raise the price above p_c is lower under commitment than under no commitment. By raising the price under commitment, the firm deflates the probability with which the consumer expects to buy, which in turn decreases the consumer’s loss in the good dimension if she does not buy, and increases her loss in the money dimension if she does. This feedback of behavior into expectations further decreases the buyer’s “attachment” to the good and thus her willingness to pay. Since without commitment, the firm does not take such feedback into account, it sets a higher price in that case.

While this intuition for Proposition 5 may seem straightforward, there is actually an effect running in the opposite direction. If the firm raises the price above p_c under no commitment, the consumer experiences a loss due to the surprisingly high price. In contrast, committing to a higher price leads her to expect to pay more, and thus not to be surprised to have to pay more. The latter effect depends on the consumer’s inframarginal demand, while the former depends on her marginal demand. But the consumer’s attachment is relatively more responsive to marginal changes in demand than the firm’s profits, because her attachment depends on her expectation to pay the price as well as to get the good. Thus, at the optimal deterministic price, the former effect outweighs the latter. By a similar logic, the same is true in general for the highest price in the firm’s price distribution. But since lower prices in non-deterministic price distributions affect comparisons the consumer draws at higher prices, this does not imply an analogue to Proposition 5 for such distributions.²²

Proposition 5 has two potentially important implications. First, a firm that “invests” in influencing consumer expectations sets a lower price than one that cannot influence expectations. Second, since a firm under commitment could increase profits ex post from raising the price, it *underprices at the top* relative to the equilibrium demand curve in the market. That is, given the consumer’s expectations, the highest price it chooses would appear (to an observer assuming that consumers are classical) to be below the profit maximizing price. Both of these predictions are testable.

²² Suppose the firm sets non-deterministic prices, and consider its pricing decision for cost \underline{c} . With commitment, the firm takes into account that the price it sets for \underline{c} will form a basis for comparison at higher prices. This leads to an incentive to increase prices at low costs that acts against the effect above. If this effect is strong enough, it increases demand at higher prices, allowing it to charge higher prices in those states as well.

5 The “Attachment Effect” and Promotional Sales

Recall the example introduced in Section 2.3, which illustrates that absent a shock to the consumer’s willingness to pay, a basic consequence of our model is the possibility of multiple personal equilibria for a given pricing distribution. We now consider the implications of multiple personal equilibria for pricing equilibria, because we believe it suggests an important novel effect for optimal pricing. We focus on a model in which the consumer’s willingness to pay is not subject to a shock, $v = 1$, c is deterministic, and the firm can commit to a pricing distribution. For a moment, suppose the firm is restricted to charge deterministic prices. Then, for any $p \in [p_{min}, p_{max}]$ for which $p \geq c$, it would be a pricing equilibrium for the firm to charge price p and the consumer to buy at prices $p' \leq p$. This suggests that the consumer’s loss aversion significantly hurts the firm if she is predisposed to not buying, but it significantly benefits the firm if she is predisposed to buying. This simple intuition, however, is partially misleading if the monopolist can offer random prices.

To understand the central idea, suppose the consumer would not buy the good for any deterministic price above p_{min} . Then, it is not optimal for the firm to offer a deterministic price. In particular, consider setting $p = p_{min} - \epsilon$ and $p = p_{min} + 2\epsilon$ with probability one-half each. By our calculation in Section 2.3, it is not a personal equilibrium for the consumer to buy with probability zero; if she expected to do so, she would still prefer to buy at the lower price. But for a sufficiently small ϵ , neither is it an equilibrium for her to buy only at the low price: If she expected to do so, her utility from buying at the high price would be $1 - (p_{min} + 2\epsilon) - \frac{3}{2}\lambda\epsilon + \frac{1}{2}(1 - \lambda(p_{min} + 2\epsilon))$, which is greater than $\frac{1}{2}(p_{min} - \epsilon - \lambda)$, her utility from not buying. Thus, the unique equilibrium is for her to buy at both prices, for an expected revenue over p_{min} .

Intuitively, by offering a sufficiently low price with some probability, the firm makes sure that the consumer buys the good at least in that state of the world, and—through the resulting increase in her anticipated probability of buying—increases her willingness to pay. The firm takes advantage of this “attachment effect” by sometimes setting a higher price. However, the above pricing policy is far from the best the firm can do.

Proposition 6. (Sales and Minimum Profits) *Suppose H is degenerate and c is deterministic, and that the firm can commit to a pricing strategy. The lowest possible expected profit the firm can*




Figure 4: The Firm’s Profits in the Worst-Case Scenario for $c = 0$

earn in a pricing equilibrium is $\max \{0, \pi(\lambda) - c\}$, where

$$\pi(\lambda) = \left(\exp \left(\frac{\lambda - 1}{\lambda + 1} \right) - 1 \right) \cdot \frac{3 + \lambda}{\lambda - 1} - 1. \quad (14)$$

When the consumer systematically plays the worst possible continuation equilibrium—which is when the monopolist earns lowest profits—the firm wants to commit to a pricing distribution for which the *unique* personal equilibrium is to buy the good with probability one. Furthermore, it prefers to “just unravel” any other personal equilibrium: For any price p on the support of the distribution, if the consumer expected to buy the good for prices below p , she should (barely) be willing to buy it at price p . This yields a differential equation at the heart of the proof of Proposition 6.

The firm’s minimum expected profits for $c = 0$ are illustrated in Figure 4 as a function of λ . As is clear from the figure, the guaranteed expected profit is decreasing in λ . If the consumer is more loss averse, it becomes more difficult to counteract her tendency to play the “unfavorable” personal equilibrium, decreasing the lower bound on profits. However, the bound decreases quite

slowly in λ . For $\lambda = 3$ —a parameter value suggested by the experimental evidence—expected profits are still about 0.95, even though for those preferences $p_{min} = \frac{1}{2}$. Thus, the firm goes a long way in eliminating the consumer’s predisposition not to buy, and makes profits close to her intrinsic valuation of the good.

Surprisingly, even as the consumer becomes infinitely loss averse ($\lambda \rightarrow \infty$), the firm cannot make expected profits less than $e - 2 \approx 0.72$. Since the consumer’s sense of loss from paying any positive amount when she was not expecting to do so becomes infinitely large, it would seem that with an unfavorable equilibrium selection, the firm should not be able to make any profits. But as the consumer’s loss aversion increases, so does the strength of the attachment effect on which the firm’s pricing strategy relies: Any sale-induced increase in her anticipated probability of buying has a large effect on her willingness to pay for the good. Thus, the firm’s minimum expected profits remain bounded far away from zero, even as the best possible profits increase without limit.

Proposition 6 therefore implies that the expected profit the firm can guarantee itself with the possibility to offer random sales is much higher than the profits it can be certain of without sales. In other words, sales play a crucial role in determining the range of market outcomes possible with loss averse consumers. But neither Proposition 6, nor our theory, can make precise predictions on how often sales will actually be observed in equilibrium.

We propose, however, that psychologically it is natural for consumers to go about their lives playing a no-purchase personal equilibrium for many products, whenever such equilibrium exists. In these circumstances, sales will be observed in pricing equilibrium, with some interesting economic interpretations. First, occasional sales in this equilibrium not only boost immediate demand, they also increase consumers’ willingness to buy at higher prices. Thus, such sales can be interpreted as an investment into future market size. This view makes it clear that only a patient firm should care to offer random sales; indeed, in the no-commitment version of our model, random prices would not be observed with deterministic costs. Second, note that if $\pi(\lambda) > c > p_{min}$, the firm earns positive profits in equilibrium. More interestingly, it sets the price below marginal cost with positive probability. Therefore, this equilibrium provides a non-competitive market reason for what would in a standard setting be interpreted as predatory pricing.

Nevertheless, a drawback of the model without shocks, and the reason we have not employed it in earlier sections, is that there is little we can definitively say about the firm’s strategy in a pricing equilibrium. Using Proposition 6, a pricing policy F is sustainable in an equilibrium if and only if there is a continuation personal equilibrium that yields profits $\pi \geq \{\pi(\lambda) - c, 0\}$.²³ As a result, a wide range of equilibrium profit levels are possible, and most profit levels can result from both deterministic and random prices.²⁴ But many of the possible pricing equilibria can only be supported by making rather ad hoc assumptions about how the pricing distribution affects the consumer’s selection between personal equilibria, raising the concern that any results we reach would be driven by these selection rules.

6 Related Industrial Organization Literature

The understanding of strategic pricing behavior is a central theme in industrial organization. Within this theme, our paper is related to the extensive work on collusion, countercyclical markups, kinked demand curves, and sales. We briefly discuss these connections and empirical evidence relating to our model. One major advantage of our theory over others is that it explains multiple pricing phenomena under a single umbrella. Our results regarding price rigidity and countercyclical markups have important implications for other fields of economics, such as macroeconomics.²⁵

Chevalier, Kashyap and Rossi (2000) investigate the pricing behavior of a large supermarket chain in Chicago, and find that most price series are characterized by only a few prices over a five-year horizon. A typical series is that of 9.5 ounce Triscuit crackers, which features only eight “regular prices” over 7.5 years. They also find overwhelming evidence of temporary sales: The price drops for a short period of time, after which it returns to its original level.

It seems hard to explain the price stickiness in their data based on menu or fixed adjustment costs. When the price moves back to its original level after a temporary sale, the firm incurs the

²³ A simple way to support such an F is the following. If the firm commits to the pricing distribution F , the consumer plays the personal equilibrium that yields profits π . If the firm chooses any other price distribution, the consumer plays the lowest-profit personal equilibrium.

²⁴ For example, for $\lambda = 3$ and $c = 0$, anything in the (approximate) interval $[0.95, 2]$ is an equilibrium profit level.

²⁵ For example, Stiglitz (1984) writes that “[c]onventional wisdom has it that a large part of the explanation of Keynesian unemployment is the observed rigidity of wages and prices.”

menu cost anyhow and should move the price to the new optimal level. It seems implausible that this level is identical before and after each of the many sales. From the perspective of our model, the fact that prices move back to the original level can be interpreted as the firm absorbing small cost shocks in order to increase sales.²⁶

Slade (1999) investigates the prices set for saltine crackers by grocery stores in a small US town. In her data, prices change infrequently and often by large amounts. Furthermore, she finds evidence of what she calls the effect of “goodwill”, that low past prices increase current sales. Our model provides an *endogenous* rationale for the price rigidity and the goodwill-effect present in her data: Sales increase the consumer’s expectation of buying the good and thereby increase demand in higher-price states of the world.

To explain these and other pricing phenomena, many theories have been developed based on strategic interactions (competition or collusion) between oligopolistic firms. In contrast to all of these models, our results are driven by the demand side. As such, we obtain our results even in a monopoly setting.

Empirical work by Mills (1927), Means (1935), and Carlton (1986, 1989) indicate that prices are less volatile in more concentrated industries. As Carlton (1989) points out, there is a widespread informal view that the lack of price movement in concentrated industries is the result of collusive behavior. The idea is that firms avoid changing prices, lest this be “misinterpreted” by other firms and lead to a price war. Athey, Bagwell and Sanchirico (2004) show that this intuition holds in the optimal symmetric collusive equilibria of a repeated game in which firms have private cost shocks.²⁷ If firms’ current behavior, however, can depend on past market shares, there exist more efficient collusive equilibria without price stickiness.²⁸

²⁶ Similarly, Kashyap (1995) finds that prices in retail catalogs often remain constant for years, even though new catalogs are designed and printed every six months. Once again, since the retailer prints a new catalog anyway, it incurs no menu costs for changing the price in response to (presumably) at least slightly different circumstances.

²⁷ If firms responded to their respective cost positions, high cost firms would have to be prevented from setting the low cost firms’ equilibrium price through a loss in continuation value when lower prices are set. With symmetric strategies in each period, this must entail an increase in the likelihood of a price war. Thus, in the optimal collusive equilibrium, rigid prices may prevail.

²⁸ In the optimal collusive equilibrium, low prices and high market shares today lead to lower equilibrium market shares in the future. Athey and Bagwell (2001) and Aoyagi (2002) show this in collusive models with explicit communication. Skrzypacz and Hopenhayn (2004), and Blume and Heidhues (2002, 2003) show that the same insight holds in repeated auction environments with limited observability and no explicit communication.

Our model is also related to an older literature on kinked demand curves (Hall and Hitch 1939, Sweezy 1939). In these models, each firm believes that if it lowers its price, rivals will do the same, while if it increases its price, rivals will not follow—leading to a kinked demand curve. Maskin and Tirole (1988) provide a game-theoretic foundation for these beliefs in a repeated alternating-move pricing game, but do not investigate the impact of cost shocks on pricing behavior.

Proposition 3 relates our paper to Rotemberg and Saloner (1986), who study the effect of demand shocks on markups of colluding firms. During a “boom” potential short run profits are high, increasing the temptation to defect from the long run optimal collusive strategy. As a result, markups must be countercyclical. In Rotemberg and Saloner’s model, an increase in the demand intercept has the same effect as a decrease in marginal cost. Empirical evidence suggests that marginal costs move procyclically,²⁹ so for industries in which demand is relatively stable over the business cycle, they in fact predict procyclical markups. In contrast, our model predicts countercyclical markups independent of whether the business cycle originates in cost or demand shocks.

Chevalier and Scharfstein (1996) develop an explanation for countercyclical markups based on capital market imperfections and the assumption that low current prices serve as market share investments because customers have to pay a cost when switching between sellers.³⁰ In a recession the default risk is higher, leading to lower investments and higher prices. Our model can be viewed as providing a psychological foundation for switching costs, since expecting to buy a good makes it painful not to buy it. In contrast to their model, however, we predict countercyclical markups for firms whose financial constraints do not change over the business cycle.³¹

There is also a considerable industrial-organization literature investigating why firms engage in sales or promotions—that is, vary prices across locations or time even when demand and cost conditions are similar. The literature developed two common rationales for sales. One set of models develops an equilibrium rationale for playing a mixed strategy in oligopoly situations.³²

²⁹ See, for example, Bilal (1987).

³⁰ Klemperer (1995) discusses under what conditions switching costs lead to countercyclical markups.

³¹ Nevertheless, more myopic firms (possibly due to financial distress) have a tendency to set higher prices in our model than patient ones (Proposition 5), as Chevalier and Scharfstein’s (1996) empirical evidence indicates. There is also evidence for an implication of Proposition 5, that firms underprice (Hoch, Dréze, and Purk 1994).

³² See, for example, Shilony (1977), Varian (1980) and Gal-Or (1982).

Other models focus on firms' incentives to price discriminate between groups of consumers.³³ In stark contrast to this literature, in our model of Section 5 the existence of low prices (weakly) increases demand in high price states.

7 Conclusion

Loss aversion is one of the most well-documented traits of human preferences. It has been observed in a variety of experimental situations, including monetary gambles and the risky or riskless buying, selling, and exchange of goods. It also affects three of the most important economic markets, residential housing markets (Genesove and Mayer 2001), labor markets (Bewley 1998), and financial markets (Odean 1998). Finally, it has been argued theoretically that much of the observed avoidance of risk should be attributed to loss aversion (Rabin 2000, Barberis, Huang, and Thaler 2003).

Despite this large and growing literature on loss aversion, the implications of the phenomenon for major market outcomes (outside finance) have not been investigated. Our paper attempts to start filling this gap. Based on Kőszegi and Rabin (2004), we develop a model of consumer behavior with loss aversion, a model that we believe transfers to many possible settings. We then embed these consumers in a marketplace with a monopolistic seller operating in a partially unpredictable economic environment, and derive the properties of the monopolist's pricing strategy.

One natural extension of this paper is to analyze the impact of loss-averse consumers on the pricing of a multi-product monopolist such as a supermarket. If consumers isolate each of their individual purchases from the others, the single-good model we have developed still applies to these sellers. However, if (as may sometimes be plausible to assume) consumers budget some joint purchases together, the analysis changes because the firm might choose to correlate prices across goods. Another natural extension is to analyze the impact of loss averse consumers on pricing in oligopoly environments. At first blush, it may seem that competition mitigates the phenomena we have identified in this paper, because oligopolists will be more worried about outcompeting each other than about responding optimally to consumers' loss aversion. While this may be true with

³³ See, for example, Salop (1977), Salop and Stiglitz (1982), Conlisk, Gerstner and Sobel (1984), and Sobel (1984, 1991).

perfect competition, in an imperfectly competitive environment loss aversion also affects how the firms compete. For example, a firm does not want to ask for higher prices than its competitors if that creates a sense of loss in its potential consumers. Thus, firms may use consumer loss aversion as a competitive weapon.

A Proofs

Proof of Proposition 1.

We prove that if h satisfies the condition 10, the function $q \mapsto \int_0^\infty (1 - H(w_{F,q}(p)))dF(p)$ has slope less than 1. Take any $q, q' \in [0, 1]$ with $q > q'$. Then, using equation 5, $w_{F,q'}(0) - w_{F,q}(0) = v(\lambda - 1)(q - q')$. Using equation 8, (in the following we use the right hand side limit whenever the derivative is not defined at p) one has

$$\frac{\partial w_{F,q'}(p)}{\partial p} - \frac{\partial w_{F,q}(p)}{\partial p} = (\lambda - 1) \left\{ (q - q') - \int_0^p (1 - H(w_{F,q}(p')))dF(p') + \int_0^p (1 - H(w_{F,q'}(p')))dF(p') \right\}.$$

Now we prove by contradiction that $0 < w_{F,q'}(p) - w_{F,q}(p) \leq (\lambda - 1)(q - q')(v + p)$ for all $p \geq 0$. Suppose by contradiction that this is not true, and let \bar{p} be the least upper bound of prices p for which it is true. Using that $w_{F,q}(p), w_{F,q'}(p) \geq 2p - (1 + \lambda)v$, and the condition on h , this implies that $(1 - H(w_{F,q}(p))) - (1 - H(w_{F,q'}(p))) < q - q'$ for all prices $p \leq \bar{p}$. But then clearly

$$0 < \frac{\partial w_{F,q'}(p)}{\partial p} - \frac{\partial w_{F,q}(p)}{\partial p} \leq (\lambda - 1)(q - q')$$

for all $p \leq \bar{p}$. Furthermore, since $w_{F,q'}(\bar{p}) - w_{F,q}(\bar{p}) \geq 0$ and $\frac{\partial w_{F,q'}(\bar{p})}{\partial p} - \frac{\partial w_{F,q}(\bar{p})}{\partial p} > 0$, the difference in the derivatives is less than or equal $(\lambda - 1)(q - q')$ for p slightly above \bar{p} . This implies that $0 < w_{F,q'}(p) - w_{F,q}(p) \leq (\lambda - 1)(q - q')(v + p)$ holds for p slightly above \bar{p} as well, a contradiction.

To complete the proof, note that similarly to the above, $0 < w_{F,q'}(p) - w_{F,q}(p) \leq (\lambda - 1)(q - q')(v + p)$ implies that

$$\int_0^\infty (1 - H(w_{F,q}(p)))dF(p) - \int_0^\infty (1 - H(w_{F,q'}(p)))dF(p) < (q - q'),$$

which proves that the map $q \mapsto \int_0^\infty (1 - H(w_{F,q}(p)))dF(p)$ has slope less than 1. \square

Proof of Lemma 1.

Let $A = \frac{b+(1+\lambda)v}{2}$. Then, for any expectations by the consumer, she buys with probability zero for any $p \geq A$. Thus, we can restrict attention to pricing distributions on the interval $[0, A]$.

Let $P(c)$ be a function that maps marginal cost into a price $p \in [0, A]$, and let F_P be the cumulative distribution function over prices induced by P and the distribution of costs. It is easy to see that we can further restrict attention to pricing functions that are non-decreasing in cost: Since

for any pricing distribution, demand is decreasing in p , in an optimal pricing strategy production decreases in costs (that is, the firm produces more in low-cost states, and less in high-cost states).

By Tychonoff's theorem, the product of compact topological spaces is compact. The space of functions from a compact set to a compact set, endowed with the topology of pointwise convergence, is isomorphic to a product of compact sets. Therefore, the space of such functions is compact in the pointwise topology. Furthermore, the set of increasing functions is a closed subset of this space (the pointwise limit of an increasing function is increasing), so is itself compact.

Let \rightarrow stand for pointwise convergence. Below, we show that if $P_i \rightarrow P$, then demand at any price $[1 - H(w_{q^*(F_{P_i}), F_{P_i}}(p))] \rightarrow [1 - H(w_{q^*(F_P), F_P}(p))]$. Hence, by Lebesgue's dominated convergence theorem, and using that the composite of two pointwise continuous functions is pointwise continuous,

$$\int_{\underline{c}}^{\bar{c}} [P_i(c) - c] \cdot [1 - H(w_{q^*(F_{P_i}), F_{P_i}}(P_i(c)))] d\Theta(c) \rightarrow \int_{\underline{c}}^{\bar{c}} [P(c) - c] \cdot [1 - H(w_{q^*(F_P), F_P}(P(c)))] d\Theta(c).$$

Therefore profits are pointwise continuous in P_i , and hence by Weierstrass' theorem (that continuous functions over a compact set have a maximum) there exists an optimal pricing function.

We are left to show that if $P_i \rightarrow P$, then

$$[1 - H(w_{q^*(F_{P_i}), F_{P_i}}(p))] \rightarrow [1 - H(w_{q^*(F_P), F_P}(p))].$$

Observe that $F_{P_i}(p)$ converges at all but countably many points to $F_P(p)$. Namely, it converges at all points p where $P^{-1}(p)$ is unique; since P is non-decreasing, this is true unless P has a flat part at p , which can only happen at countably many points. Furthermore, recall that $w_{q, F(\cdot)}(0)$ is independent of $F(\cdot)$, and that

$$\frac{\partial w_{F, q}}{\partial p} = 2 + (\lambda - 1)(1 - q) + (\lambda - 1) \int_0^p (1 - H(w_{F, q}(p'))) dF(p').$$

Since we can "build up" $w_{q, F(\cdot)}$ using the Fundamental Theorem of Calculus, the fact that $H(\cdot)$ is continuous allows us to apply Lebesgue's dominated convergence theorem repeatedly by transfinite induction,³⁴ to conclude that $\frac{\partial w_{F_{P_i}, q}}{\partial p}$ is almost everywhere equal to $\frac{\partial w_{F_P, q}}{\partial p}$. Therefore, $w_{q, F_{P_i}} \rightarrow w_{q, F_P}$.

We complete the proof by establishing that $q^*(F_{P_i}) \rightarrow q^*(F_P)$. This is sufficient because $w_{q, F}$ is continuous in q .

The personal equilibrium probability of buying is determined by the property that

$$q = \int_{-\infty}^{\infty} [1 - H(w_{q, F}(p))] dF(p).$$

We prove that the right-hand side of this equation is continuous in F . Combined with the fact that the equilibrium is unique for all pricing distributions, this implies that $q^*(F)$ is continuous in F .

³⁴ Here, transfinite induction is applied to the property that $w_{F_{P_i}, q}$ converges pointwise to $w_{F_P, q}$ on $[0, p]$.

Using integration by parts,

$$\int_{-\infty}^{\infty} [1 - H(w_{q,F}(p'))] dF(p') = - \int_{-\infty}^{\infty} \frac{\partial[1 - H(w_{q,F}(p'))]}{\partial p} F(p') dp'.$$

We have shown above that the derivative of demand converges pointwise almost everywhere. Therefore, the result follows from again applying Lebesgue's dominated convergence theorem.

This completes the proof that a pricing equilibrium exists.

We now find a \bar{w} that satisfies the requirements of the Lemma. Suppose that the optimal pricing distribution is F . Denote the maximum price in the support of F by $\bar{p} \leq A$. If $\bar{p} \leq \bar{c}$, the result is obvious by the assumption on the support of H . Suppose therefore that $\bar{p} > \bar{c}$.

We prove by contradiction that

$$(\bar{p} - \bar{c}) \cdot 2h(w_{F,q^*(F)}(\bar{p})) < 1 - H(w_{F,q^*(F)}(\bar{p})).$$

Suppose the opposite is the case. Then, since $w'_{F,q^*(F)}(\bar{p}) > 2$, holding the consumer's expectations constant (so that $w_{F,q^*(F)}$ determines demand), the firm could increase its revenues at price \bar{p} by lowering this price. Of course, this change also affects the consumer's continuation personal equilibrium by increasing the probability of buying. If the firm makes positive average profits at all prices, this further increases profits. If there is a positive measure of prices where the firm makes negative profits, we construct a new price distribution with which the probability of buying is the same as with F in the following way. First, we raise prices in the region where profits are negative to keep the overall probability of buying the same. Second, we raise prices above this region to keep the probability of buying at each of those prices constant. Each of these changes increases profits, contradicting that F was an optimal pricing strategy.

Finally, we show that $\underline{w} = a$ satisfies the requirements of the proposition. Suppose the lowest price on the support of F is \underline{p} . Suppose by contradiction that $w'_{F,q^*(F)}(\underline{p}) < a$. Then, for prices at and near \underline{p} , the consumer buys with probability one, and she would still do so if the firm increased prices in this range. Furthermore, by a proof identical to that in Proposition 2, this increase in prices can be used to increase profits at higher prices as well. This contradicts that F is profit maximizing. \square

Proof of Proposition 2.

Part I. Let q^* be the probability of buying with the price distribution F , and G^* the distribution of buying prices conditional on buying. In the proof, we will use the convention that a subscript on a probability (“*Prob*”) denotes the distribution with respect to which the probability should be taken.

We start by restricting attention to discrete distributions with finitely many atoms. For any positive integer J , notice that there is a profit-maximizing discrete price distribution with at most J atoms. As in Lemma 1, we can restrict attention to distributions on $[0, A]$. Now, the set of discrete price distributions on $[0, A]$ with at most J atoms is compact. Furthermore, the firm's profits are continuous in the pricing distribution. Thus, there is a profit-maximizing distribution

among them. We prove that there exists a $d_{min} > 0$ such that for any J , the profit-maximizing price distribution with at most J atoms satisfies the property that its atoms are at least d_{min} apart.

We prove this by contradiction. Suppose that the price distribution F has two neighboring atoms p_1 and p_2 of F , with $p_1 < p_2$ and $p_2 - p_1 = 2d$. We show that if d is sufficiently small, F cannot be profit-maximizing.

We start by constructing the distribution F_1 from F in the following way. We take a small weight $\epsilon > 0$ (to be specified below) from each of the atoms p_1 and p_2 , and replace it with an atom \bar{p} that satisfies

$$\frac{1 - H(w_{F,q^*}(p_1)) + 1 - H(w_{F,q^*}(p_2))}{2} = 1 - H(w_{F_1,q^*}(\bar{p})),$$

which implies that the overall probability of buying the good (conditional on price being on this interval) stays fixed.

Define the price p implicitly through

$$\frac{1 - H(w_{F,q^*}(p_1)) + 1 - H(w_{F,q^*}(p_2))}{2} = 1 - H(w_{F,q^*}(p)).$$

Since H is differentiable, for any $\kappa < 1$ there is a $d_\kappa > 0$ such that if $d < d_\kappa$, then $\frac{1}{\kappa}d > p - p_1 > \kappa d$. Furthermore, since H'' is bounded, we can choose d_κ to be independent of p_1 and p_2 .

It follows from equation 8 that on the interval $(0, p_1)$, $w_{F,q^*} = w_{F_1,q^*}$ and on the interval (p_1, p) , $w_{F,q^*} > w_{F_1,q^*}$. Hence, since by definition $w_{F,q^*}(p) = w_{F_1,q^*}(\bar{p})$, one has $\bar{p} > p$. Now, note that on the interval (p_1, p) both curves w_{F,q^*} and w_{F_1,q^*} are linear, and using equation 8 we have that the difference in slopes is $(\lambda - 1)\epsilon[1 - H(w_{F,q^*}(p_1))]$. If using $\epsilon = \min\{Prob_F(p_1), Prob_F(p_2)\}$ yields $\bar{p} \in (p_1, p_2)$, then we use that weight. Otherwise, we use the weight ϵ that yields $\bar{p} = p_2$.

Since difference in slopes on the interval (p_1, \bar{p}) is $(\lambda - 1)\epsilon[1 - H(w_{F,q^*}(p_1))]$ and since $p - p_1 > \kappa d$, the difference $w_{F,q^*}(p) - w_{F_1,q^*}(p) \geq (\lambda - 1)\kappa\epsilon d[1 - H(w_{F,q^*}(p_1))]$. Using this and the fact that $w_{F,q^*}(p) = w_{F_1,q^*}(\bar{p})$ and that the slope of the curve w_{F_1,q^*} is at most $\lambda + 1$, we have

$$\bar{p} - p \geq \frac{\lambda - 1}{\lambda + 1}\kappa\epsilon d[1 - H(w_{F,q^*}(p_1))]. \quad (15)$$

Furthermore, notice that by the Regularity Condition on H ,

$$p(1 - H(w_{F,q^*}(p))) > \frac{p_1(1 - H(w_{F,q^*}(p_1))) + p_2(1 - H(w_{F,q^*}(p_2)))}{2}.$$

Since by definition $(1 - H(w_{F,q^*}(\bar{p}))) = (1 - H(w_{F_1,q^*}(p)))$, this implies that the change increases revenues on this interval by at least

$$\frac{\lambda - 1}{\lambda + 1}\kappa\epsilon d(1 - H(w_{F,q^*}(p_1)))(1 - H(w_{F,q^*}(p))). \quad (16)$$

The above calculation also implies that the change increases purchase probability-weighted price by at least 2ϵ times Expression 16 (the probability of the price being on the changed interval, times the minimum revenue change on that interval). Hence, $w_{F,q^*}(p_2) > w_{F_1,q^*}(p_2)$, and $w_{F,q^*}(p_2) - w_{F_1,q^*}(p_2)$ is at least 2ϵ times Expression 16. Furthermore, equation 8 implies that above p_2 the

slopes of w_{F,q^*} and w_{F_1,q^*} are identical. Hence, w_{F_1,q^*} leads to a purchasing probability $q > q^*$. We thus construct a new pricing distribution F_2 , which is identical to F_1 for prices $p \leq \bar{p}$, but for which $q = q^*$.

Suppose p_3 is the smallest atom of F_1 above \bar{p} . (Note that p_3 may equal p_2 .) We construct the distribution of prices F_2 from F_1 by shifting the *entire* distribution starting at p_3 by the amount Δ satisfying $w_{F_2,q^*}(p_3 + \Delta) = w_{F,q^*}(p_3)$. Given our lower bound for $w_{F,q^*}(p_2) - w_{F_1,q^*}(p_2)$, we must have

$$\Delta \geq 2\epsilon \cdot \frac{1}{1+\lambda} \cdot \frac{\lambda-1}{\lambda+1} \kappa \epsilon d (1 - H(w_{F,q^*}(p_1)))(1 - H(w_{F,q^*}(p))).$$

Since by replacing (part of) the atoms p_1 and p_2 with an atom at \bar{p} , we have kept the probability of buying the same, this ensures that for any $p \geq p_3$, we have $w_{F_2,q^*}(p + \Delta) = w_{F,q^*}(p)$. Thus, since we have kept the probability of buying at all atoms the same, q^* is the equilibrium probability of buying with the price distribution F_2 . All these changes have increased the firm's revenues by at least 2ϵ times

$$\frac{\lambda-1}{\lambda+1} \kappa \epsilon d (1 - H(w_{F,q^*}(p_1)))(1 - H(w_{F,q^*}(p))) \left(1 + \frac{1}{1+\lambda} Prob_{F_2}(p \geq \bar{p}) \right).$$

We next consider production costs. To minimize production costs the firm associates with each price an interval of marginal costs, with higher prices being associated with higher marginal costs. Thus, when using the original pricing distribution F , there exists an interval of marginal costs $[\underline{c}_1, \bar{c}_1]$ associated with the price p_1 and an interval $[\bar{c}_1, \bar{c}_2]$ associated with the price p_2 . When the firm moves mass ϵ from prices p_1 and p_2 to price \bar{p} , it moves the highest marginal costs originally associated with p_1 and the lowest marginal cost originally associated with p_2 to a new interval $[c', c'']$ that now is associated with \bar{p} such that $\int_{c'}^{\bar{c}_1} \theta(c) dc = \epsilon$ and $\int_{\bar{c}_1}^{c''} \theta(c) dc = \epsilon$. Since θ is continuous, we can choose ϵ small enough so that

$$\int_{\bar{c}_1}^{c''} c\theta(c) dc - \int_{c'}^{\bar{c}_1} c\theta(c) dc \leq \frac{1}{\kappa} \cdot \frac{\epsilon}{\theta(\bar{c}_1)}.$$

When using the pricing distribution F , the difference in production between the two prices p_1 and p_2 is $H(w_{F,q^*}(p_2)) - H(w_{F,q^*}(p_1))$. Thus, when switching to the new pricing distribution, the firm's production cost increase by less than

$$\frac{1}{\kappa} \cdot \frac{\epsilon}{\theta(\bar{c}_1)} \cdot \frac{[H(w_{F,q^*}(p_2)) - H(w_{F,q^*}(p_1))]}{2}.$$

For any $\kappa < 1$, we can choose d'_κ so that

$$H(w_{F,q^*}(p_2)) - H(w_{F,q^*}(p_1)) < \frac{1}{\kappa} h(w_{F,q^*}(p_1)) 2d(\lambda + 1). \quad (17)$$

Thus, the firm's saving is at most

$$\frac{\epsilon h(w_{F,q^*}(p_1)) d(\lambda + 1)}{\theta(\bar{c}_1) \kappa^2}.$$

In sum, for $d < d_\kappa, d'_\kappa$, profits on the interval increase if

$$\frac{\lambda - 1}{(\lambda + 1)^2} \kappa \epsilon d (1 - H(w_{F,q^*}(p_1))) (1 - H(w_{F,q^*}(p))) \left(1 + \frac{1}{1 + \lambda} Prob_{F_2}(p \geq \bar{p})\right) > \frac{\epsilon h(w_{F,q^*}(p_1)) d}{\theta(\bar{c}_1) \kappa^2}$$

or

$$\theta(\bar{c}_1) \left(1 + \frac{1}{1 + \lambda} Prob_{F_2}(p \geq \bar{p})\right) > \frac{h(w_{F,q^*}(p_1))}{(1 - H(w_{F,q^*}(p_1))) (1 - H(w_{F,q^*}(p)))} \cdot \frac{1}{\kappa^3} \cdot \frac{(\lambda + 1)^2}{\lambda - 1}.$$

Since higher prices are associated with higher costs, we have $Prob_{F_2}(p \geq \bar{p}) = 1 - \Theta(c'')$. Thus, the above becomes

$$\theta(\bar{c}_1) \left(1 + \frac{1 - \Theta(c'')}{1 + \lambda}\right) > \frac{h(w_{F,q^*}(p_1))}{(1 - H(w_{F,q^*}(p_1))) (1 - H(w_{F,q^*}(p)))} \cdot \frac{1}{\kappa^3} \cdot \frac{(\lambda + 1)^2}{\lambda - 1}.$$

By assumption, this is true if κ is sufficiently close to 1 and ϵ is sufficiently small.

To complete the proof, we use the fact that all distributions F^* can be approximated arbitrarily closely (for the purposes of profits) by discrete distributions. Namely, for any j , we can break up the support of F^* into intervals of length $\frac{1}{2j}$, and on each interval, replace the prices with an atom at the average. Call this discrete distribution F_j . As $j \rightarrow \infty$, the profits with F_j approach the profits with F^* .

Now, using the above procedure, we can replace each F_j by a discrete distribution F'_j that yields higher profits than F_j and has atoms at least d_{min} apart. Hence, there is a maximum number of mass points that a distribution F'_j can have. Thus, the distributions F'_j have a convergent subsequence of pricing distribution which have the same number of mass points. Consider this subsequence and order the mass points from the lowest to the highest. As the first mass points lies in the interval $[0, A]$, this sequence has a convergent subsequence in which the first mass point converges. Consider this subsequence, it must have a subsequence in which the second mass point converges, etc... . Hence, the distributions F'_j have a convergent subsequence, whose atoms are at least d_{min} apart. Obviously, this limit yields at least as high profits as F^* . In fact, if F^* does not have the property that each of its atoms are at least d_{min} apart, then the limit yields strictly higher profits. (Use the above argument to construct a sequence whose prices are at least $d_{min} - \eta$ apart for some sufficiently small η but whose prices are not at least d_{min} apart, which weakly increases profits. Then use the above change to construct a distribution whose prices are at least d_{min} , which strictly increases profits.) \square

Part II. We prove by contradiction. Suppose that the price distribution F has an atom at p , with probability r . We prove that it is in the firm's interest to break this atom into two atoms with weight $\frac{r}{2}$ each, and move the two atoms slightly apart.

Assume for a moment that p is an "isolated" atom: there is a neighborhood of p with measure r . In that case, we can do the reverse of the estimation in Part I with the following key changes. Namely, we need a an upper bound on the revenue decrease from splitting up a pricing atom, and a lower bound on the cost savings. Therefore, the "revenue inequality" corresponding to Inequality 15, which bounds the revenue decrease from above, will read

$$\bar{p} - p \leq \frac{\lambda - 1}{2} \cdot \frac{1}{\kappa} \epsilon d [1 - H(w_{F,q^*}(p_1))].$$

The right-hand side of this inequality differs from that of Inequality 15 in two ways. First, we use a lower bound on the slope of w_{F_1, q^*} , and therefore replace $1 + \lambda$ by 2. Second, to get an upper bound, we have $\frac{1}{\kappa}$ instead of κ . On the other hand, the “cost inequality” corresponding to Inequality 17 becomes

$$H(w_{F, q^*}(p_2)) - H(w_{F, q^*}(p_1)) > \kappa h(w_{F, q^*}(p_1)) 4d.$$

Here, we have used a lower bound on the slope of w_{F, q^*} , replacing $1 + \lambda$ with 2, and exchanged $\frac{1}{\kappa}$ with κ .

To complete the proof, note that as $\delta \rightarrow 0$, the measure of the interval of $(p - \delta, p + \delta)$ under F approaches r . Thus, we can ignore the weight around p . \square

Proof of Proposition 3. Let q^* be the consumer’s probability of purchase under the optimal pricing distribution.

Note that

$$p \cdot \left(L(p, c) - \frac{1}{\eta_D(p)} \right) = (p - c) + \frac{D_F(p)}{D'_F(p)},$$

so we will prove that this expression is strictly decreasing.

Let $p(c)$ be the firm’s pricing function, and that the optimal price distribution is F . Suppose by contradiction that there are some p_1 and p_2 with $p_2 = p(c_2) > p(c_1) = p_1$ and

$$(p_2 - c_2) + \frac{D_F(p_2)}{D'_F(p_2)} \geq (p_1 - c_1) + \frac{D_F(p_1)}{D'_F(p_1)}, \quad (18)$$

Furthermore, since F is differentiable almost everywhere, we can choose p_1 and p_2 so that F is differentiable at both points. The idea of the proof is to show that profits can be increased by pushing the price distribution closer together. More specifically, we take some weight from the left of p_1 and put it on p_1 , and take some weight from the right of p_2 , and put it on p_2 in a way that increases profits.

Formally, we start by forming the price distribution F' in the following way. We take the weight F places on the interval $(p_1 - \epsilon_1, p_1)$ and place it on p_1 , and take the weight F places on $(p_2, p_2 + \epsilon_2)$ and place it on p_2 . Furthermore, we choose ϵ_1 and ϵ_2 so that

$$\begin{aligned} & (F(p_1) - F(p_1 - \epsilon_1))(1 - H(w_{F, q^*}(p_1))) + (F(p_2 + \epsilon_2) - F(p_2))(1 - H(w_{F, q^*}(p_2))) \\ &= \int_{p_1 - \epsilon_1}^{p_1} (1 - H(w_{F, q^*}(p))) dF(p) + \int_{p_2}^{p_2 + \epsilon_2} (1 - H(w_{F, q^*}(p))) dF(p) \end{aligned}$$

That is, under the demand function that would prevail with F , this change does not change total demand. This means that we are moving demand from the low-price state to the high-price state.

Suppose the above changes increase average probability-weighted prices on the interval $(p_1 - \epsilon_1, p_1)$ by dp_1 , and decrease probability-weighted prices on the interval $p_2, p_2 + \epsilon$ by dp_2 . We next establish that by Inequality 18, under the demand curve D_F , this either creates a gain in profits for sufficiently small dp_1 and dp_2 , or the decrease in profits is second-order in dp_1 .

In the limit as dp_1 approaches zero, we must have $dp_2 = \frac{D'_F(p_1)}{D'_F(p_2)} dp_1$ in order to keep overall demand under D_F constant. The change in profits must be equal to

$$dp_1(D_F(p_1) + (p_1 - c_1)D'_F(p_1)) - dp_2(D_F(p_2) + (p_2 - c_2)D'_F(p_2))$$

up to a term that is at most second-order in dp_1 . Substituting for dp_2 and using 18 yields that the above expression is greater than or equal to zero.

Now, it is easy to see that $w_{F,q^*}(p_1) > w_{F',q^*}(p_1)$, and by Expression 5, the difference is first-order in dp_1 . To hold overall sales fixed, we first create the price distribution F'' from F' by shifting all prices on the interval $[p_1, p_2]$ by Δ satisfying

$$w_{F,q^*}(p_1) = w_{F'',q^*}(p_1 + \Delta).$$

This ensures that $w_{F,q^*}(p) = w_{F'',q^*}(p + \Delta)$ for all $p \in [p_1, p_2]$. In addition, Δ is first-order in dp_1 .

Now, we argue the average purchase price below $(p_2 + \epsilon_2)$ is greater under w_{F'',q^*} than under w_{F,q^*} . First, we constructed F'' so that demand given by w_{F'',q^*} and w_{F,q^*} is the same in this region. Furthermore, we achieved this by moving demand from low-price states to high-price states, and shifting the price distribution to the right. By Expression 5, both of these changes increase the average purchase price.

Therefore,

$$w_{F,q^*}(p_2 + \epsilon_2) > w_{F'',q^*}(p_2 + \epsilon_2).$$

To complete the argument, we create the distribution F''' from F'' by shifting all prices above $p_2 + \epsilon_2$ by Δ' so that

$$w_{F,q^*}(p_2 + \epsilon_2) = w_{F''',q^*}(p_2 + \epsilon_2 + \Delta').$$

Then, the probability of purchase under price distribution F''' is also q^* . Since Δ is first-order in dp_1 , F''' clearly yields greater profits than F . \square

Proof of Lemma 2. As in Lemma 1, let $A = \frac{b+(1+\lambda)v}{2}$. Then, for any expectations by the consumer, the firm does not want to choose a price higher than A , since for any $p \geq A$, the consumer's purchase probability is zero.

Consider the space \mathcal{P} of increasing pricing functions from $[\underline{c}, \bar{c}]$ to $[0, A]$. We define the function $\phi : \mathcal{P} \rightarrow \mathcal{P}$ in the following manner. $\phi(P)$ is the firm's optimal pricing strategy, if consumers expect the pricing distribution induced by the pricing strategy P . (We will establish in a moment that ϕ is indeed a function.) A fixed point of ϕ defines a pricing equilibrium, so we prove that ϕ has a fixed point.

For any price distribution expectations F by the consumer and any realized c , the firm's profit maximization problem is strictly concave by the strict convexity of $w_{F,q^*(F)}$ and the regularity condition on H . Thus, ϕ is indeed a function.

We have established in the proof of Lemma 1 that the demand curve $1 - H(w_{F_P,q^*(F_P)}(p))$ is continuous in the pricing strategy P . Since the firm's profit maximization problem is strictly concave, the firm's optimal price for any realized c is also continuous in P . Thus, ϕ is a pointwise

continuous function. By the Kakutani-Fan-Glicksberg theorem, it has a fixed point, completing the proof. \square

Proof of Proposition 4.

Part I. We prove by contradiction. Suppose that the firm's price is not deterministic. First, it is easy to show that there is a non-trivial interval I such that for any $p \in I$, there is a cost $c \in [\underline{c}, \bar{c}]$ such that the firm sets price p for that cost. To see this, note simply that the firm's optimal price $p(c)$ is continuous in c .

Suppose the consumer expects to buy with probability q , and expects the distribution of buying prices conditional on purchase to be G . Since the demand function is decreasing in p , it is differentiable almost everywhere. Let p be a price for which the demand function is differentiable. Then, the cost c for which the firm sets price p satisfies

$$h(w_{F,q,G}(p))(p-c)[2+(1-q)(\lambda-1)+qG(p)(\lambda-1)] = 1 - H(w_{F,q,G}(p)). \quad (19)$$

Now take a sequence $p_i \searrow p$ such that the demand function is differentiable at each p_i . Let the corresponding costs be c_i . Then, for each i we must have

$$h(w_{F,q,G}(p_i))(p_i - c_i)[2 + (1 - q)(\lambda - 1) + qG(p_i)(\lambda - 1)] = 1 - H(w_{F,q,G}(p_i)).$$

For p_i close to p , the difference in the two equations above can be approximated by

$$\begin{aligned} & h(w_{F,q,G}(p))(p_i - p - c_i + c)[2 + (1 - q)(\lambda - 1) + qG(p)(\lambda - 1)] \\ & + h(w_{F,q,G}(p))(p - c)(\lambda - 1)[qG(p_i) - qG(p)] \\ & + h'(w_{F,q,G}(p))(p_i - p)(p - c)[2 + (1 - q)(\lambda - 1) + qG(p)(\lambda - 1)]^2 \\ & = -h(w_{F,q,G}(p))[2 + (1 - q)(\lambda - 1) + qG(p)(\lambda - 1)](p_i - p). \end{aligned}$$

Since $p(c)$ is increasing in c , $qG(p_i) - qG(p)$ is the probability that the agent buys at prices between p and p_i . This can be approximated by $(1 - H(w_{F,q,G}(p)))(\Theta(c_i) - \Theta(c))$, which itself can be approximated by $(1 - H(w_{F,q,G}(p)))\theta(c)(c_i - c)$. Using this, the above becomes

$$\begin{aligned} & \frac{p_i - p}{c_i - c} \cdot [2 + (1 - q)(\lambda - 1) + qG(p)(\lambda - 1)] \cdot \\ & \cdot [2h(w_{F,q,G}(p)) + h'(w_{F,q,G}(p))(p - c)[2 + (1 - q)(\lambda - 1) + qG(p)(\lambda - 1)]] \\ & = h(w_{F,q,G}(p))[2 + (1 - q)(\lambda - 1) + qG(p)(\lambda - 1)] - (p - c)(1 - H(w_{F,q,G}(p)))(\lambda - 1)\theta(c). \end{aligned}$$

The left-hand side of the above equality is positive. To see this, note that all multiplicative terms other than the last one are clearly positive. To prove that the last one is positive, differentiate the hazard rate and substitute the first-order condition 19.

To arrive at a contradiction, we show that the right-hand side of this inequality is less than zero. Using the first-order condition 19 to substitute for $p - c$, we want to show that

$$(\lambda - 1)\theta(c) \frac{(1 - H(w_{F,q,G}(p)))^2}{h(w_{F,q,G}(p))[2 + (1 - q)(\lambda - 1) + qG(p)(\lambda - 1)]} \geq 2 + (1 - q)(\lambda - 1) + qG(p)(\lambda - 1).$$

Now clearly $2 + (1 - q)(\lambda - 1) + qG(p)(\lambda - 1) \leq 1 + \lambda$, so this inequality follows from the condition of the proposition. This completes the proof.

Part II. To show that any pricing equilibrium features a non-atomic price distribution, we do the converse estimation to Part I. \square

Proof of Proposition 5.

Suppose the optimal commitment price is p_c . We prove that if the consumer expects the firm to charge p_c with probability one, the firm would like to choose a higher price ex post. This implies that the equilibrium Condition 13 is violated in the following way:

$$p_c - \bar{c} < \frac{1 - H(w_{p_c, q^*(p_c)}(p_c))}{h(w_{p_c, q^*(p_c)}(p_c)) \cdot (2 + (\lambda - 1))}.$$

Recall that the right-hand side of the above inequality is decreasing in p , while the left-hand side is increasing, so the interval of no-commitment equilibrium prices lies above p_c .

Given that the firm charges a deterministic price, the equilibrium probability of buying, $q^*(p)$, is determined by

$$q^*(p) = 1 - H[w_{p, q^*(p)}(p)],$$

where $w_{p, q^*(p)}(p) = 2p - 2v + (1 - q^*(p))(\lambda - 1)p - q^*(p)(\lambda - 1)v$. Let κ be the average marginal cost of the firm.

With commitment, the firm solves

$$\max_p (p - \kappa) \cdot [1 - H(w_{p, q^*(p)}(p))],$$

leading to the first-order condition for the optimal commitment price p_c of

$$1 - H(w_{p_c, q^*(p_c)}(p_c)) - (p_c - \kappa)h(w_{p_c, q^*(p_c)}(p_c)) \left(2 + (1 - q^*(p_c))(\lambda - 1) - \frac{\partial q^*(p_c)}{\partial p_c}(\lambda - 1)(p_c + v) \right) = 0. \quad (20)$$

We will show that the firm can gain ex post from charging a higher price than p_c for *all* costs. This implies that there is a profitable deviation. For $p' \geq p_c$, the firm's average profit from charging p' after the consumer has been expecting p_c is

$$(p' - \kappa) \cdot [1 - H(2p' - 2v + (1 - q^*(p_c))(\lambda - 1)p' - q^*(p_c)(\lambda - 1)v + (\lambda - 1)q^*(p_c)(p' - p_c))].$$

The derivative of this with respect to p' at p_c is

$$1 - H(w_{p_c, q^*(p_c)}(p_c)) - (p_c - \kappa)h(w_{p_c, q^*(p_c)}(p_c))(2 + (1 - q^*(p_c))(\lambda - 1) + (\lambda - 1)q^*(p_c)). \quad (21)$$

We will argue that Expression 21 is strictly greater than zero. By Equation 20, this is equivalent to proving that

$$q^*(p_c) < -\frac{\partial q^*(p_c)}{\partial p_c}(p_c + v). \quad (22)$$

By the firm's optimization problem under commitment,

$$(p_c - \kappa) \cdot \frac{\partial q^*(p_c)}{\partial p_c} + q^*(p_c) = 0,$$

which implies Inequality 22. □

Proof of Proposition 6.

For any given reference lottery, it is clear that there is a price p_r (not necessarily in the support of the expected price distribution) below which the consumer buys the good, and above which she does not. Consequently, equilibrium behavior also has this property. Now, analogously to Expression 5, if the consumer expects to buy for prices strictly below p_r , then the difference in utility between buying and not buying at p_r is

$$B_f(p_r) \equiv 1 - p_r - (\lambda - 1) \int_0^{p_r} (p_r - p') dF(p') + F(p_r)(\lambda - p_r) + (1 - F(p_r))(1 - \lambda p_r). \quad (23)$$

Clearly, p_r is the reservation price in a continuation equilibrium if $B_f(p_r) = 0$, and if $B_f(p) > 0$ for all p in or below the support of F , in the unique personal equilibrium the consumer buys with probability one.

We will solve for a continuous price distribution F_{min} that has a connected support with a lower bound of p_{min} , which satisfies

$$B_f(p) = 1 - p - (\lambda - 1) \int_0^p (p - p') dF(p') + F(p)(\lambda - p) + (1 - F(p))(1 - \lambda p) = 0, \quad (24)$$

for any p on its support. Let π_{min} be the profits if the firm chooses F_{min} and the consumer always buys the good. Our strategy for the proof is the following: First, we argue that π_{min} is a lower bound for equilibrium profits. Second, if $\pi_{min} > 0$, we construct an equilibrium in which the firm earns π_{min} . To show the second step, observe that the consumer is willing to buy with probability one if the firm selects F_{min} . We will show that the firm cannot do better if for any other chosen pricing distribution F the consumer plays the worst possible personal equilibrium. This step will also imply that if $\pi_{min} \leq 0$, then the firm's minimum profit is zero.

Step 1. For any $\epsilon > 0$, if the firm shifts the distribution F_{min} to the left by ϵ then Expression 23 is strictly greater than zero for any price on the support. Therefore in the unique continuation equilibrium the consumer always buys the good and profits are $\pi_{min} - \epsilon$. Thus in equilibrium the firm's profits cannot be lower π_{min} .

Step 2. We show that for any distribution F , if the consumer plays the worst possible continuation equilibrium and the firm's profit is positive, it is less than π_{min} . This implies that if $\pi_{min} \leq 0$, the lowest possible profit level is zero, and (in combination with Step 1) it also means that if $\pi_{min} > 0$, it is the lowest possible profit level.

We carry out the proof by performing a set of changes to F , each of which increases the firm's profits, and which approach F_{min} . First, notice that for any distribution of prices F which leads to selling the good with probability less than 1, the firm can do better. To see this, suppose that

the reservation price in the lowest-profit personal equilibrium is p_r , and $F(p_r) < 1$. Since this equilibrium leads to positive profits for the firm, we must have $p_r > c$. Then, consider putting the weight $1 - F(p_r)$ on the price $p_r - \epsilon$. By Expression 23, for a sufficiently small ϵ , the consumer buys with probability one, increasing expected profits.

Thus, we look for price distributions F for which the unique equilibrium is for the consumer to buy the good at any price in the support of F . Among these pricing choices, profit maximization is equivalent to revenue maximization. Let $[\underline{p}, \bar{p}]$ be the smallest closed interval that contains the support of F . The constraint that the consumer buys with probability one implies that for any $p \in [0, \bar{p})$, the equilibrium reservation price has to be greater than p . To check this, it must be the case that if the consumer expects to buy the good at prices below p , she would be willing to buy it at price p . Using Expression 23, this is equivalent to saying that

$$1 - p - (\lambda - 1) \int_0^p (p - p') dF(p') + F(p)(\lambda - p) + (1 - F(p))(1 - \lambda p) > 0. \quad (25)$$

Now take any $\epsilon > 0$. We prove that for any distribution for which the above expression is greater than ϵ on a positive measure subset of $[\underline{p}, \bar{p}]$, the firm can increase profits by bringing it below ϵ everywhere. First, if for any $\delta > 0$, $F(p_{min} - \delta) > 0$, the firm can increase profits by replacing all weight below $p_{min} - \delta$ by a point mass at $p_{min} - \delta$. (By Expression 23, this does not decrease the reservation price $p_r \geq p_{min}$, and hence must increase profits.) By the definition of p_{min} the left-hand-side limit of Expression 23 is zero at p_{min} for any distribution that places no weight below p_{min} . By continuity, for a sufficiently small δ , this will result in a distribution for which Expression 23 is less than ϵ for any $p < p_{min} - \delta$ sufficiently close to $p_{min} - \delta$.

Now suppose that Expression is less than or equal to ϵ for any $p < p_1$ with $p \in [\underline{p}, \bar{p}]$, but that it goes above ϵ on some interval $[p_1, p_2]$. For this to be the case, either p_1 has to be a point mass of F , or F must put positive weight on the interval (p_1, p_2) . Consider changing the distribution of prices in the following way. If p_1 is an atom of F , take off a weight from p_1 so that with the remaining weight, Expression 23 equals ϵ for $p = p_1$. Take this weight, and whatever other weight F places on the on $(p_1, p_2]$, and put all of that weight on p_2 . This change increases the left-hand side of Expression 24 for any $p > p_2$, and does not affect it for $p < p_1$. Furthermore, the left-hand-side of Expression 24 is continuously decreasing over any interval in which there is no probability weight and hence, if $p_2 - p_1$ is sufficiently small, the expression will be (weakly) less than ϵ and greater than zero on $[p_1, p_2]$. Thus, the new price distribution still satisfies the condition that the consumer buys with probability one, and it has higher expected profits. Using this step, by transfinite induction we can increase profits by constructing a price distribution for which Expression 23 is less than or equal to ϵ everywhere.³⁵ As ϵ converges to zero, F converges pointwise to F_{min} , and therefore the associated profits converge to π_{min} , completing Step 2.

Finally, we characterize F_{min} . Taking the derivative of both sides of Equation 24 yields

$$-1 - (\lambda - 1)F(p) - F(p) - \lambda(1 - F(p)) + f(p)(\lambda - 1)(p + 1) = 0.$$

³⁵ To apply transfinite induction on $(p_{min} - \delta, \infty)$, let $P(a)$ be the property that there is an F' constructed from F using our step such that either a is above the support of F' , or for F' , Expression 23 is between zero and ϵ for $p_{min} - \delta < p < a$.

This gives

$$f(p) = \frac{\lambda + 1}{(\lambda - 1)(p + 1)}.$$

Integrating gives

$$F(p) = \frac{\lambda + 1}{\lambda - 1} \ln(p + 1) + \text{constant}.$$

Using the condition that $F(p_{min}) = 0$, we get

$$F(p) = \frac{\lambda + 1}{\lambda - 1} \ln\left(\frac{p + 1}{p_{min} + 1}\right).$$

Setting $F(\bar{p}) = 1$ gives

$$\bar{p} = \exp\left(\frac{\lambda - 1}{\lambda + 1}\right) \cdot (p_{min} + 1) - 1.$$

Finally, the pieces are together to calculate the firm's expected revenue:

$$\int_{p_{min}}^{\bar{p}} p f(p) dp = \int_{p_{min}}^{\bar{p}} \frac{\lambda + 1}{\lambda - 1} \cdot \left(1 - \frac{1}{p + 1}\right) dp = \frac{\lambda + 1}{\lambda - 1} \cdot (\bar{p} - p_{min}) - 1.$$

Substituting for \bar{p} and rearranging completes the proof. □

References

- AOYAGI (2002): “Efficient Collusion in Repeated Auctions with Communication,” Working Paper.
- ATHEY, S., AND K. BAGWELL (2001): “Optimal Collusion with Private Information,” *RAND Journal of Economics*, 32(3), 428–465.
- ATHEY, S., K. BAGWELL, AND C. SANCHIRICO (2004): “Collusion and Price Rigidity,” *Review of Economic Studies*, 71(2), 317–349.
- BARBERIS, N., M. HUANG, AND T. SANTOS (2001): “Prospect Theory and Asset Prices,” *Quarterly Journal of Economics*, 116(1), 1–53.
- BARBERIS, N., M. HUANG, AND R. THALER (2003): “Individual Preferences, Monetary Gambles and the Equity Premium,” URL: <http://www.stanford.edu/~mhuang/papers/BHT.pdf>.
- BENARTZI, S., AND R. H. THALER (1995): “Myopic Loss Aversion and the Equity Premium Puzzle,” *Quarterly Journal of Economics*, 111(1), 73–92.
- BEWLEY, T. F. (1998): “Why Not Cut Pay?,” *European Economic Review*, 42(3-5), 459–490.
- BILS, M. (1987): “The Cyclical Behavior of Marginal Cost and Price,” *American Economic Review*, 77(5), 838–855.
- BLUME, A., AND P. HEIDHUES (2002): “Modeling Tacit Collusion in Auctions,” Working Paper.
- (2003): “Private Monitoring in Auctions,” Working Paper.
- CARLTON, D. W. (1986): “The Rigidity of Prices,” *American Economic Review*, 76(4), 637–658.
- (1989): “The Theory and the Facts of How Markets Clear: Is Industrial Organization Valuable for Understanding Macroeconomics?,” in *Handbook of Industrial Organization, Volume 1*, ed. by R. Schmalensee, and R. Willig, vol. 1, pp. 909–946. Elsevier Science Publishers B.V.
- CHEVALIER, J., AND D. SCHARFSTEIN (1996): “Capital-Market Imperfections and Countercyclical Markups: Theory and Evidence,” *American Economic Review*, 86(4), 703–725.
- CHEVALIER, J. A., A. K. KASHYAP, AND P. E. ROSSI (2000): “Why Don’t Prices Rise During Periods of Peak Demand? Evidence from Scanner Data,” NBER Working Paper Number 7981.
- CONLISK, J., E. GERSTNER, AND J. SOBEL (1984): “Cyclic Pricing by a Durable Goods Monopolist,” *Quarterly Journal of Economics*, 99(3), 489–505.
- FUDENBERG, D., AND D. K. LEVINE (1989): “Reputation and Equilibrium Selection in Games with a Patient Player,” *Econometrica*, 57(4), 759–778.
- GAL-OR, E. (1982): “Hotelling’s Spatial Competition as a Model of Sales,” *Economic Letters*, 9, 1–6.

- GENESOVE, D., AND C. MAYER (2001): “Loss Aversion and Seller Behavior: Evidence from the Housing Market,” *Quarterly Journal of Economics*, 116(4), 1233–1260.
- HALL, R. L., AND C. . J. HITCH (1939): “Price Theory and Business Behavior,” *Oxford Economic Papers*, 2, 12–45.
- KAHNEMAN, D., J. L. KNETSCH, AND R. H. THALER (1990): “Experimental Tests of the Endowment Effect and the Coase Theorem,” *Journal of Political Economy*, 98(6), 1325–1348.
- (1991): “Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias,” *Journal of Economic Perspectives*, 5(1), 193–206.
- KAHNEMAN, D., AND A. TVERSKY (1979): “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, 47(2), 263–291.
- KASHYAP, A. K. (1995): “Sticky Prices: New Evidence from Retail Catalogs,” *Quarterly Journal of Economics*, 110(1), 245–274.
- KŐSZEGI, B. (2004): “Anticipation in Observable Behavior,” Working Paper, UC Berkeley.
- KŐSZEGI, B., AND M. RABIN (2004): “A Model of Reference-Dependent Preferences,” Economics Department, University of California, Berkeley, Working Paper E04-337.
- KLEMPERER, P. (1995): “Competition when Consumers have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade,” *Review of Economic Studies*, 62(4), 515–539.
- MASKIN, E., AND J. TIROLE (1988): “A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles,” *Econometrica*, 56(3), 571–599.
- MEANS, G. (1935): “Industrial Prices and their Relative Inflexibility,” Senate Document 13, 74th Congress, 1st session. Washington D.C.: US Government Printing Office.
- MILLS, F. (1927): *The behavior of prices*. New York: National Bureau of Economic Research.
- ODEAN, T. (1998): “Are Investors Reluctant to Realize Their Losses?,” Mimeo, UC Davis.
- RABIN, M. (2000): “Diminishing Marginal Utility of Wealth Cannot Explain Risk Aversion,” in *Choices, Values, and Frames*, ed. by D. Kahneman, and A. Tversky, chap. 11, pp. 202–208. New York: Russell Sage Foundation.
- ROTEMBERG, J., AND G. SALONER (1986): “A Supergame Theoretic Model of Business Cycles and Price Wars During Booms,” *American Economic Review*, 76(3), 390–407.
- SALOP, S. (1977): “The Noisy Monopolist: Imperfect Information, Price Dispersion, and Price Diskrimination,” *Review of Economic Studies*, 44, 393–406.
- SALOP, S., AND J. STIGLITZ (1982): “The Theory of Sales: A Simple Model of Equilibrium Price Dispersion with Identical Agents,” *American Economic Review*, 72(5), 1121–1130.

- SHILONY, Y. (1977): “Mixed Pricing in Oligopoly,” *Journal of Economic Theory*, 14, 373–388.
- SKRZYPACZ, A., AND H. HOPENHAYN (2004): “Tacit Collusion in Repeated Auctions,” *Journal of Economic Theory*, 114, 153–169.
- SLADE, M. (1999): “Sticky Prices in a Dynamic Oligopoly: An Investigation of (s, S) Thresholds,” *International Journal of Industrial Organization*, 17, 477–511.
- SOBEL, J. (1984): “The Timing of Sales,” *Review of Economic Studies*, 51(3), 353–368.
- SOBEL, J. (1991): “Durable Goods Monopolist with Entry of New Consumers,” *Econometrica*, 59(5), 1455–1485.
- STIGLITZ, J. (1984): “Price Rigidities and Market Structure,” *American Economic Review*, 74(2), 350–355.
- SWEEZY, P. (1939): “Demand Conditions under Oligopoly,” *Journal of Political Economy*, 47, 568–573.
- TIROLE, J. (1988): *The Theory of Industrial Organization*. MIT Press.
- TVERSKY, A., AND D. KAHNEMAN (1991): “Loss Aversion in Riskless Choice: A Reference-Dependent Model,” *Quarterly Journal of Economics*, 106(4), 1039–1061.
- VARIAN, H. R. (1980): “A Model of Sales,” *American Economic Review*, 70(4), 651–659.