Monopoly Pricing when Consumers are Antagonized by Unexpected Price Increases: A "Cover Version" of the Heidhues-Koszegi-Rabin Model*

Ran Spiegler†

March 15, 2010

Abstract

This paper reformulates and simplifies a recent model by Heidhues and Koszegi (2005), which in turn is based on a behavioral model due to Koszegi and Rabin (2006). The model analyzes optimal pricing when consumers are loss averse in the sense that an unexpected price hike lowers their willingness to pay. The main message of the Heidhues-Koszegi model, namely that this form of consumer loss aversion leads to rigid price responses to cost fluctuations, carries over. I demonstrate the usefulness of this "cover version" of the Heidhues-Koszegi-Rabin model by obtaining new results: (1) loss aversion lowers expected prices; (2) the firm's incentive to adopt a rigid pricing strategy is stronger when fluctuations are in demand rather than in costs.

1 Introduction

Like other creative artists, economic theorists value originality. When we construct a model of a certain economic phenomenon, we try to distance and differentiate our

*Financial support from ERC grant no. 230251 and the ESRC (UK) is gratefully acknowledged. I thank Ayala Arad, Eddie Dekel, Kfir Eliaz, Yves Guéron, Paul Heidhues and Ariel Rubinstein for extremely helpful comments.

†Tel Aviv University and University College London. E-mail: r.spiegler@ucl.ac.uk. URL: http://www.homepages.ucl.ac.uk/~uctprsp/

personal creation from previous work, unless we wish to present it as an extension, application or foundation of an existing model. However, there is also value in a different sort of exercise, in which the theorist takes an existing model and tries to rewrite it from scratch. When playing this game, the economic theorist steps into the shoes of another theorist, takes her economic motivation and basic modelling idea as given, but then tries to "do it his way". The result of such a re-modelling exercise can be viewed as a "*cover version*" of the original model. It is not an entirely new creation, because it is a deliberate and explicit variant on an existing model. However, it is also not entirely derivative, because the variant is often sufficiently different to merit separate attention. By offering a different way of formalizing the same economic phenomenon, the cover version contributes to our understanding of the phenomenon. Furthermore, the cover version may be technically easier to apply in certain domains, in which case it expands the original model's scope of applications.

This paper proposes a "cover version" of a recent model of optimal pricing when consumers are loss averse, due to Heidhues and Kőszegi (2005) - which in turn builds on a behavioral model proposed by Kőszegi and Rabin (2006). These two papers are referred to as HK and KR in the sequel, and the model in its totality is referred to as HKR. The motivation behind the model is an idea of long standing: consumers are antagonized by unexpected price hikes, and this may cause firms to be cautious when adapting prices to demand or cost shocks. (See Hall and Fitch (1939), Okun (1981) and Kahneman, Knetch and Thaler (1986).) In other words, consumers' distaste for unpleasant price surprises acts as a "menu cost" that deters firms from changing prices in response to exogenous shocks. Similar claims have been made in relation to wage rigidity in labor markets (see Fehr, Goette and Zehnder (2009)).

The HKR model formalizes this idea. The distaste for unexpected price hikes is viewed as a manifestation of loss aversion (as in Kahneman and Tversky (1979)): unpleasant price surprises are perceived as a loss relative to the price the consumer expected, which acts as his reference point. HK apply the model of loss aversion preferences with expectations-driven reference points proposed by KR, and analyze optimal monopoly pricing when consumers behave according to this model.[1]

I present a variation on the HKR model of optimal pricing when consumers are averse to unpleasant price surprises. I borrow the economic motivation as well as the main modelling idea from HKR, and do not claim any originality in this regard.

[1] Heidhues and Koszegi (2008) and Karle and Peitz (2008) extend this analysis to oligopolistic settings. For a field study which is skeptical of price-variation antagonism, see Courty and Pagliero (2009).

However, I depart from the HKR model in several dimensions. Although the concept of loss aversion has been backed up by decades of experimental research, to turn it into a workable economic model the theorist has to make crucial modelling choices that themselves lack experimental support. Specifically, any model of loss-averse consumers with expectations-driven reference points has to address the following questions:

1. Which aspects of the market outcome are evaluated as gains or losses relative to a reference point?

2. Does the consumer's expectation regarding his own consumption decision enter the specification of the reference point?[2]

3. When the reference point is stochastic, how should we "sum over" all the possible values it can get?

My variation on the HKR model departs from the original model in all three dimensions - mostly in the direction of simplifying it. First, while HKR assume that loss aversion affects the consumer's evaluation of both the price paid and the quantity consumed, I allow only the former. Second, while HKR assume that the reference point incorporates the consumer's expectation regarding his own action, I assume that only his expectation regarding the price that the firm charges determines his reference point. Finally, HKR assume that the consumer computes his expected utility from a given action for every possible reference point, and integrates over all reference points to obtain an evaluation of the action. In contrast, I assume that the consumer has a single reference point in mind, but this reference point is drawn from the price distribution that characterizes the market. In other words, while HKR "sum over" reference points at the utility level, I do so at the demand level.

I reproduce the main insights in HKR. Optimal prices are rigid in two respects. First, the price range is cramped relative to the benchmark without loss aversion: mark-ups are higher in low-cost states and lower in high-cost states. Second, prices can be sticky in the sense of not responding at all to small cost shocks. (I devote a lot less space to this effect than HKR.) I then proceed to demonstrate the usefulness of the reformulated model with a pair of new results. First, I show that the expected price that firms charge is weakly lower than in the benchmark without loss aversion. However, this effect is not monotone with respect to the magnitude of consumers' loss

[2] As we shall see, this question is relevant even if the consumer's actual consumption is not evaluated in terms of gains or losses relative to a reference point.

aversion. Second, I show that the price rigidity effect is stronger in some sense when shocks are in demand rather than in costs.

The rest of the paper is structured as follows. In Section 2 I present the model. Section 3 provides a partial characterization of optimal pricing strategies and demonstrates price rigidity effects. Sections 4 and 5 are devoted to the new results: the impact of loss aversion on expected prices, and the difference between cost and demand shocks. For expositional clarity, the detailed comparison with the original HKR model is deferred to Section 6.

2 The Model

A monopolistic firm sells a single unit of a product to a continuum of measure one of consumers. The firm's marginal cost c is distributed uniformly over some set of possible values C. Suppose for now that C is finite. Denote $|C| = m > 1$. Let c_h and c_l denote the highest and lowest cost values in C, $1 > c_h > c_l > 0$. Let $\bar{c}$ denote the average cost in C. The firm commits to a deterministic pricing strategy $P : C \to \mathbb{R}$, where $P(c)$ is the price the monopolist charges when the marginal cost is c. Note that because of the randomness of the firm's marginal cost, its pricing strategy induce a probability measure μ_P over stated prices, where

$$\mu_P(x) = \frac{|\{c \in C \mid P(c) = x\}|}{m}$$

The consumers' choice model is as follows. Each consumer first draws a cost state c^e from the uniform distribution over C, and sets his reference price to be $p^e = P(c^e)$. Thus, the reference point is essentially drawn from μ_P. The consumer buys the product if and only if the actual price p satisfies $p \leq u - L(p, p^e)$, where u is the consumer's "raw" willingness to pay for the product, and L is a loss function that satisfies several properties: (i) it weakly increases with p and weakly decreases with p^e; (ii) $L(p, p^e) = 0$ whenever $p \leq p^e$. This structure of L captures the loss-aversion aspect of consumer behavior: consumers react to unpleasant price surprises with a reduced willingness to pay, while pleasant price surprises do not change their willingness to pay. The raw willingness to pay u is distributed uniformly - and independently of the reference point - over $[0, 1]$.

The assumption that p^e is randomly drawn from μ_P requires justification. Clearly, there are many other ways to model the formation of the reference points. For instance, one could assume that the expected price according to μ_P is the consumer's reference point. The interpretation of "sampling-based" reference point formation is that the consumer creates his expectations on the basis of a random past market experience (his own or via word of mouth). Since this is one of the main departures from HKR, it will be discussed further in Section 6.

"Sampling-based" reference point formation implies that consumers can differ in two dimensions. First, they have different "raw" willingness to pay u for the product. Second, they have different market experiences that lead to different reference prices. Thus, the "sampling-based" formation of reference points enriches our conception of a consumer's "type"; moreover, it means that the firm can influence, through its pricing decision, the distribution of consumer types.

The monopolist's maximization problem can be written as a discrete optimal control problem. The firm chooses the function P to maximize

$$\Pi(P) = \frac{1}{m^2} \cdot \sum_{c} \sum_{c'} [P(c) - c] \cdot \max[0, 1 - P(c) - L(P(c), P(c'))] \tag{1}$$

To see why this is the objective function, note that for each realization of the marginal cost c, the firm is uncertain about the consumer's willingness to pay. Since his "raw" willingness to pay u is distributed uniformly over $[0, 1]$, in the absence of loss aversion (i.e., if L always takes the value zero) the probability that the consumer buys the product at a price $P(c)$ is $1 - P(c)$. However, if the consumer's expected price is $P(c') < P(c)$, the probability he will buy the product at $P(c)$ drops to $\max[0, 1 - P(c) - L(P(c), P(c'))]$. Since the reference price $P(c')$ is drawn from μ_P, we need to sum over all possible values of c'.

When L always gets the value zero, the optimal pricing strategy is

$$P^0(c) = \frac{1 + c}{2} \tag{2}$$

for every $c \in C$. In comparison, if the firm were restricted to charging a constant price for all cost values, the optimal price would be

$$\bar{p} = \frac{1 + \bar{c}}{2} \tag{3}$$

In the presence of loss aversion, consumer demand is a function of the price distribution induced by the firm's strategy. Specifically, given a pricing strategy P, the probability that the consumer buys the firm's product at a price p is

$$D_P(p) = \frac{1}{m}\sum_{c'} \max[0, 1 - p - L(p, P(c'))] \tag{4}$$

Thus, consumer demand depends on the price distribution as a whole. When the firm changes its price in one cost state, this affects (probabilistically) the consumer's reference point, hence consumer demand at other prices. In this sense, local prices changes have a global effect on consumer demand. When we hold P fixed, D_P decreases with p, as usual. Note that as long as $P(c) < 1$ for all c, $D_P[P(c)] > 0$ for all c. The reason is simple: when $P(c) < 1$, "raw" willingness to pay exceeds $P(c)$ with positive probability, and the reference price is equal to $P(c)$ (or higher) with positive probability.

3 Price Rigidity

In this section I begin characterizing optimal pricing strategies. The first main result is preceded by a pair of lemmas.

Lemma 1 *Let P be an optimal pricing strategy. Then, for every $c \in C$, $P(c) < 1$ and consumer demand (given by* (4)*) is strictly positive at $P(c)$.*

Proof. Let P be an optimal pricing strategy. Let us first show that $P(c) < 1$ for all $c \in C$. Assume, contrary to this claim, that $P(c^*) \geq 1$ for some $c^* \in C$. Then, consumer demand is zero at $P(c^*)$. Suppose that the firm deviates to a pricing strategy P' such that $P'(c^*) = 1 - \varepsilon$, where $\varepsilon > 0$ is arbitrarily small, and $P'(c) = P(c)$ for all $c \neq c^*$. Following the deviation, consumer demand at $P'(c^*)$ is strictly positive, because $1 - P'(c^*) - L(P'(c^*), P'(c^*)) = \varepsilon > 0$. Furthermore, for every $c \neq c^*$,

$$\max[0, 1 - P'(c) - L(P'(c), P'(c^*))] = \max[0, 1 - P(c) - L(P(c), P(c^*))]$$

hence the deviation does not change consumer demand at any $c \neq c^*$. It follows that the deviation is profitable. Then, for every $c \in C$, $P(c) < 1$. Since $L(P(c), P(c)) = 0$, consumer demand is strictly positive at $P(c)$. ■

The intuition for the result that aggregate consumer demand is always strictly positive is simple: for every price realization, there is positive probability that an individual consumer's reference price is weakly higher. Thus, any optimal pricing strategy P has the property that $D_P[P(c)] > 0$ for all c. The next lemma relies on this result.

Lemma 2 *Every optimal pricing strategy P is weakly increasing in c and satisfies $P(c) \geq c$ for every $c \in C$.*

Proof. Let P be an optimal pricing strategy. Suppose that $c_1 > c_2$ and yet $P(c_2) > P(c_1)$. Denote $P(c_1) = p_1$ and $P(c_2) = p_2$. Suppose that the firm switches to a pricing strategy P' that coincides with P for all $c \neq c_1, c_2$, and $P'(c_1) = p_2$, $P'(c_2) = p_1$. The change in the firm's objective function as a result of the deviation is

$$\frac{1}{m}\sum_c [(P'(c) - c) \cdot D_{P'}(P'(c)) - (P(c) - c) \cdot D_P(P(c))]$$

The deviation has the property that it does not alter the induced price distribution. That is, $\mu_P = \mu_{P'}$. Therefore, the deviation does not change consumer demand: $D_{P'}(x) = D_P(x)$ for every price x. Accordingly, let us omit the subscript of the demand function. We can thus rewrite the above expression as follows:

$$\frac{1}{m}\sum_c \{[P'(c) - c]D(P'(c)) - [P(c) - c]D(P(c))\}$$

Since P and P' coincide over $c \neq c_1, c_2$, this expression is strictly positive if and only if

$$[p_2 - c_1]D(p_2) - [p_1 - c_1]D(p_1) + [p_1 - c_2]D(p_1) - [p_2 - c_2]D(p_2) > 0$$

which simplifies into

$$(c_2 - c_1)(D(p_2) - D(p_1)) > 0$$

Recall that by assumption, $c_1 > c_2$ and $p_2 > p_1$. Since D is strictly decreasing in the range in which it is strictly positive, and since we have established in the previous lemma that $D(P(c)) > 0$ for all $c \in C$, the inequality holds. Therefore, the deviation to P' is strictly profitable.

Let $C^* = \{c \in C \mid P(c) < c\}$. Suppose that C^* is non-empty. Consider a deviation to a pricing strategy P' that satisfies $P'(c) = c$ for all $c \in C^*$ and $P'(c) = P(c)$ for all $c \notin C^*$. For every $c \in C^*$, the firm's profit conditional on being chosen ceases to be strictly negative as a result of the deviation. For every $c \notin C^*$, the deviation weakly lowers the loss aversion term because it raises the consumers' reference price with positive probability. Since we have established that consumer demand is strictly positive at all cost states, the deviation is strictly profitable. Therefore, C^* must be empty. ■

In the absence of loss aversion, this result followed trivially from (2). In the presence of loss aversion, it is a non-trivial result because in general, the firm's pricing strategy influences consumer demand. However, when the firm deviates to a pricing strategy that does not change the overall price distribution, consumer demand is unaffected. The proof is based on this type of deviation and thus demonstrates the tractability of "sampling-based" reference point formation.

We are now ready to prove our first main result. When the consumer is loss averse, the firm raises the price at the lowest level of marginal cost and lowers the price at the highest level of marginal cost, relative to the benchmark with no loss aversion. This "cramped" price range can be viewed as an instance of *price rigidity.* The firm does not want the consumer to experience large losses that will reduce his willingness to pay, and therefore shrinks the price range.

Proposition 1 *Let P be an optimal pricing strategy P. Then:*

$$P^0(c^l) \leq P(c^l) \leq P(c^h) \leq P^0(c^h)$$

Proof. Let P be an optimal pricing strategy. By Lemma 2, P is weakly increasing and satisfies $P(c) \geq c$ for every $c \in C$, hence $P(c^h) \geq P(c^l) \geq c^l$ and $P(c^h) \geq c^h$.

(i) $P^0(c^l) \leq P(c^l)$. Recall that $P^0(c^l) = \frac{1}{2}(1 + c^l)$. In the absence of loss aversion, the hump shape of the firm's objective function implies that for every $c \geq c^l$, the price $P^0(c^l)$ yields a higher profit than any $p' < P^0(c^l)$. Define c^0 to be the highest cost c for which $P(c) < \frac{1}{2}(1 + c^l)$. Since P is weakly increasing, $P(c) < \frac{1}{2}(1 + c^l)$ for all $c \leq c^0$. Suppose that the firm deviates to a pricing strategy P' that satisfies $P'(c) = \frac{1}{2}(1 + c^l)$ for $c \leq c^0$ and coincides with P for $c > c^0$. The "bare" profit excluding the loss aversion term goes up. As to the loss aversion component, because P' is flat over $c \leq c^0$, $L(P'(c_1), P'(c_2)) = 0$ whenever $c_1, c_2 \leq c^0$. Moreover, because $P'(c) > P(c)$ for $c \leq c^0$ and $P'(c) = P(c)$ for $c > c^0$, $L(P'(c_1), P'(c_2)) \leq L(P(c_1), P(c_2))$ whenever $c_2 \leq c^0 < c_1$. Since P and P' coincide at all $c > c^0$, $L(P'(c_1), P'(c_2)) = L(P(c_1), P(c_2))$ when $c_1, c_2 > c^0$. Finally, when $c_1 \leq c^0 < c_2$, $L(P'(c_1), P'(c_2)) = 0$ because by the definitions of c^0 and P', $P'(c_2) > P'(c_1)$. It follows that the deviation is profitable.

(ii) $P(c^h) \leq P^0(c^h)$. Consider two cases. First, suppose that P is flat - i.e., $P(c^l) = P(c^h) = \bar{p}$. We have seen that in this case, the optimal price $\bar{p}$ is given by (3),

which is strictly below $\frac{1}{2}(1 + c^h)$. Second, suppose that $P(c^l) < P(c^h) = p^h$. Let C^* be the set of cost values c for which $P(c) = P(c^h) = p^h$. Since P is weakly increasing in c, there exists $c^* \in C$ such that $C^* = \{c \in C \mid c > c^*\}$. By (4), consumer demand at p given P is strictly lower than $1 - p$ for every $p > p^l$. Assume that $p^h > \frac{1}{2}(1 + c^h)$. Suppose that the firm switches to a pricing strategy P' such that $P'(c) = P(c) - \varepsilon$ for all $c \in C^*$ and $P'(c) = P(c)$ for all $c \notin C^*$. If $\varepsilon > 0$ is sufficiently small, $P'(c) > P'(c')$ for every $c \in C^*$, $c' \notin C^*$. By the hump shape of the firm's profit function in the absence of loss aversion, this deviation strictly raises this "bare" profit in all states $c \in C^*$. In addition, the deviation weakly lowers the loss aversion term in those states, without changing the loss aversion term in the states outside C^*. Therefore, the deviation is strictly profitable. ■

Thus, when the consumer is loss averse, the firm raises the price at the lowest level of marginal cost and lowers the price at the highest level of marginal cost, relative to the benchmark with no loss aversion. The firm does not want the consumer to experience large losses that will reduce his willingness to pay, and therefore shrinks the price range.

3.1 Maximal Flexibility vs. Maximal Rigidity: An Example

To get a better grasp of the forces that determine price rigidity, let us examine the following example. Assume that $L(p, p^e) = \delta$ whenever $p > p^e$, where $\delta > 0$ is an exogenous parameter. This is an extreme case of loss aversion, in the sense that the loss function is discontinuous at $p = p^e$: the slightest price increase relative to the consumer's reference price generates a disutility. However, this disutility is insensitive to the magnitude of the price increase. This special case lends itself to a clean demonstration of the forces that influence price rigidity.

Let us put aside the quest for an optimal pricing strategy, and instead compare two extreme strategies: the optimal, fully flexible pricing strategy in the absence of loss aversion, denoted P^0 and given by (2), and the optimal fully rigid (i.e., constant) price, given by (3). The latter eliminates the loss aversion term because it has no price variation. When we calculate the loss in "bare" expected profits (i.e., ignoring the loss aversion term) as a result of switching from P^0 into the constant price $\bar{p}$, we see that it is

$$\frac{1}{4}[(\frac{1}{m}\sum_{c \in C} c^2) - (\bar{c})^2]$$

which is proportional to the *variance* of c. On the other hand, by switching from P^0

into the optimal constant price $\bar{p}$, the firm eliminates the expected loss due to loss aversion. Since the probability of trade may be zero for some realizations of actual and expected prices, this expected loss is bounded from above by

$$\begin{aligned} \frac{1}{m^2}\sum_{c}\sum_{c'<c}(P^0(c)-c)\cdot\delta &= \frac{1}{m^2}\sum_{c}\sum_{c'<c}(\frac{1-c}{2})\cdot\delta < \\ \frac{1}{m^2}\cdot\frac{1-\bar{c}}{2}\cdot\frac{m(m-1)}{2}\cdot\delta &= \frac{(m-1)(1-\bar{c})\cdot\delta}{4m} \end{aligned}$$

Holding the parameters m, $\bar{c}$ and δ fixed, we can see that whether the firm will prefer P^0 or the optimal constant price $\bar{p}$ depends on the variance of c.

The insight from this comparison is that the firm's incentive to introduce rigidity into its pricing strategy depends on two aspects of cost fluctuation: (i) the variance of costs, which measures the global spread of fluctuations; (ii) the number of cost realizations, which captures the frequency of fluctuations. When the firm experiences a large number of small fluctuations, it will tend towards a rigid pricing strategy because the attempt to avoid antagonizing consumers will be the dominant consideration. On the other hand, when fluctuations are large and infrequent, the firm will tend towards a flexible pricing strategy, because the "bare" incentive to allow prices to respond to cost shocks outweighs the incentive to minimize unpleasant surprises for consumers.

3.2 Price Stickiness

So far, I used the term "price rigidity" to describe a cramped range of prices compared with the benchmark pricing strategy P^0. However, when economists discuss price rigidity, they often have in mind a "stickiness" property - namely, lack of response to small shocks. In order to be able to address this issue, we need to switch from a finite set of cost values to a continuum. In this sub-section, I assume that c is drawn from the uniform distribution over $[0, 1]$. The firm's strategy is a function $P : [0, 1] \to \mathbb{R}$, and its objective function should be rewritten as follows:

$$\Pi(P) = \int_0^1 \int_0^1 [P(c) - c] \cdot \max[0, 1 - P(c) - L(P(c), P(c'))]dcdc'$$

All the results obtained in Section 3 for the finite case extend (subject to standard modifications) to the continuum case.

I now show that in the special case of extreme loss aversion, optimal pricing strategies indeed display price stickiness, in the sense that prices do not always respond to small cost shocks.

Proposition 2 *Let $L(p, p^e) = \delta$ whenever $p > p^e$. Then, any optimal pricing strategy is constant over some interval of cost values.*

Proof. Assume the contrary. By Lemma 2, this means that there is an optimal pricing strategy P that is strictly increasing over $[0, 1]$. Define

$$p^* = \frac{2 + \varepsilon}{4}$$

for any $\varepsilon > 0$. Suppose that the firm deviates to a strategy P' defined as follows:

$$P'(c) = \begin{cases} p^* & , \quad c < \varepsilon \\ \max[P(c), p^*] & , \quad c \geq \varepsilon \end{cases}$$

Note that P' is a weakly increasing function. Our objective is to show that if ε is sufficiently small, the deviation is profitable.

Let us first establish an upper bound on the loss in "bare" expected profits caused by the deviation. If $P(c) < p^*$ for some $c \geq \varepsilon$, then p^* is closer to $P^0(c) = \frac{1}{2}(1 + c)$ than $P(c)$. By the hump shape of the firm's "bare" profit function for any given c, it follows that $P'(c)$ attains a weakly higher "bare" profit at $c \geq \varepsilon$ than $P(c)$. Therefore, the loss in "bare" expected profits is bounded from above by

$$\begin{aligned} & \int_0^{\varepsilon} \{[P^0(c) - c] \cdot [1 - P^0(c)] - [p^* - c] \cdot [1 - p^*] dc \\ = & \ \frac{1}{48}\varepsilon^3 \end{aligned}$$

Let us now turn to the effect that the deviation has on the loss aversion term. By the construction of P' and the definition of L, $L(P'(c), P'(c^e)) \leq L(P(c), P(c^e))$ for all c, c^e. Since P is strictly increasing while P' is flat over the interval $[0, \varepsilon)$, the reduction in the loss aversion term as a result of the deviation is at least

$$\int_0^{\varepsilon} \int_{c'}^{\varepsilon} \delta dc dc' = \frac{1}{2}\delta\varepsilon^2$$

Thus, if ε is sufficiently small, the firm's gain from curbing the loss aversion term outweighs the loss in "bare" profits. Therefore, the deviation is profitable. ■

The case of extreme loss aversion is relatively straightforward to analyze because the loss aversion term only depends on whether the actual price exceeds the reference price. When the magnitude of the price surprise also affects the loss aversion term, calculations become more complicated and require optimal-control techniques. I conjecture that the price stickiness result extends to any loss function that is very steep for small price surprises and quickly becomes flat.

4 Impact of Loss Aversion on Expected Prices

What is the effect of loss aversion on the *expected* price that the monopolist charges in optimum? To address this problem, let us consider a special case, in which C is finite and L is only a function of $\Delta = p - p^e$. This is a common specification in the literature (including HKR) due to its tractability. If the firm adopts a constant function $P(c) = p^*$ for all c, then by definition L is always zero. We saw that in this case $p^* = \frac{1}{2}(1 + \bar{c})$ - i.e., the expected price is exactly as in the benchmark with no loss aversion.

For simplicity, assume that loss aversion is not too strong, in the following sense:

$$L(\frac{1+c^h}{2}, \frac{1+c^l}{2}) < \frac{1-c^h}{2} \tag{5}$$

Given Proposition 1, this restriction ensures that $1 - P(c) - L(P(c), P(c')) > 0$ for all c, c' under any optimal pricing strategy P. In other words, consumer demand is strictly positive conditional on any realization of actual and reference prices. The average price that the firm charges under P is simply

$$\frac{1}{|C|} \sum_{c \in C} P(c)$$

Proposition 3 *Under every optimal strategy, the average price the firm charges is weakly lower than* $\bar{p} = \frac{1}{2}(1 + \bar{c})$.[3]

Proof. Let P be an optimal pricing strategy. If P is a constant function - i.e., $P(c) = \bar{p}$ for all c - then we have seen that $\bar{p} = \frac{1}{2}(1+\bar{c})$. Assume that P is not a constant function. Since it maximizes the firm's expected profit, it satisfies the following inequality for

[3]This proof follows a suggestion by Ariel Rubinstein, which supplanted a previous proof based on differentiability assumption.

every alternative strategy Q:

$$\sum_c \sum_{c'} [P(c) - c] \cdot \max[0, 1 - P(c) - L(P(c), P(c'))]$$
$$\geq \sum_c \sum_{c'} [Q(c) - c] \cdot \max[0, 1 - Q(c) - L(Q(c), Q(c'))]$$

Define $Q(c) \equiv P(c) - \varepsilon$, where $\varepsilon > 0$ is arbitrarily small. Then, since L is only a function of the difference between the actual and expected price, the following inequality holds:

$$\sum_c \sum_{c'} [P(c) - c] \cdot \max[0, 1 - P(c) - L(P(c), P(c'))]$$
$$\geq \sum_c \sum_{c'} [P(c) - \varepsilon - c] \cdot \max[0, 1 - P(c) + \varepsilon - L(P(c), P(c'))]$$

Condition (5) ensures that $1 - P(c) + \varepsilon - L(P(c), P(c')) > 0$ for all c, c'. It follows that we can rewrite the above inequality as follows:

$$\sum_c \sum_{c'} \{[P(c) - c] \cdot [1 - P(c) - L(P(c), P(c'))] - [P(c) - \varepsilon - c] \cdot [1 - P(c) + \varepsilon - L(P(c), P(c'))]\} \geq 0$$

This inequality is simplified into

$$2 \sum_c \sum_{c'} P(c) \leq \sum_c \sum_{c'} [1 + c + \varepsilon - L(P(c), P(c'))]$$

Since P is not a constant function, $\sum_c \sum_{c'} L(P(c), P(c')) > 0$. Therefore, if ε is sufficiently close to zero, we can write

$$2 \sum_c \sum_{c'} P(c) < \sum_c \sum_{c'} [1 + c]$$

which immediately implies the result. ■

The rough intuition for this result is as follows. Since loss aversion diminishes willingness to pay, the firm effectively faces lower aggregate demand than in the absence of loss aversion, and this is a force that impels the firm to lower its price on average. Note that the average price does not decrease monotonically with the intensity of loss aversion. When loss aversion is sufficiently strong, the firm's optimal pricing strategy is the constant price $\bar{p}$ given by (3), which is equal to the average price under P^0.

When condition (5) fails to hold, consumer demand is not necessarily positive under every realization of reference and actual prices. In this case, I am able to prove a slightly modified version of Proposition 3: the average price conditional on consumer demand being strictly positive is below $\bar{p} = \frac{1}{2}(1 + \bar{c})$. The proof of this result is similar to the proof of Proposition 3 and therefore omitted, for the sake of brevity.

5 Cost vs. Demand Fluctuations

Throughout the paper, I have followed HKR by assuming that fluctuations occur in the cost dimension only. It is straightforward to extend the model by introducing aggregate demand fluctuations as well, and allowing the firm to condition its price on both cost and demand shocks. Let Ω be a finite set consisting of m *states*. The firm's prior over Ω is uniform. Each state $\omega \in \Omega$ is characterized by a cost-demand pair of parameters (c_ω, v_ω), where $v_\omega > c_\psi$ for all $\omega, \psi \in \Omega$. In state ω, the measure of the consumer population is v_ω, and the consumers' "raw" willingness to pay is uniformly drawn from the interval $[0, v_\omega]$. Thus, an increase in v corresponds to an upward shift in the demand function faced by the monopolist. Let $\bar{c}$ and $\bar{v}$ denote the average values of c and v across all states.

A pricing strategy is a function $P : \Omega \to \mathbb{R}$. The firm's objective function is:

$$\Pi(P) = \frac{1}{m^2} \cdot \sum_{\omega\in\Omega} \sum_{\psi\in\Omega} [P(\omega) - c_\omega] \cdot \max[0, v_\omega - P(\omega) - L(P(\omega), P(\psi))]$$

In the benchmark without loss aversion (i.e., $L \equiv 0$), the firm's optimal pricing strategy is

$$P^0(\omega) = \frac{v_\omega + c_\omega}{2}$$

The optimal constant price is

$$\bar{p} = \frac{\bar{v} + \bar{c}}{2}$$

Thus, in the absence of loss aversion, the optimal pricing strategy treats demand and cost shocks symmetrically. This raises the question of whether this equal treatment property carries over when consumers display loss aversion. In this section, I show that the answer is negative. Moreover, the firm's incentive to employ a rigid pricing strategy is stronger in some sense when fluctuations occur in the demand dimension.

5.1 A Two-State Example

Assume that there are two states of nature, l and h. Let $v_h \geq v_l$ and $c_h \geq c_l$, with at least one strict inequality. A pricing strategy is thus represented by a pair of prices

(p_l, p_h), where $p_\omega = P(c_\omega, v_\omega)$. It can be shown (in the manner of Lemma 2) that any optimal pricing strategy satisfies $p_h \geq p_l$. Since $L(p, p^e) = 0$ whenever $p \leq p^e$, the objective function can be simplified into:

$$(p_l - c_l)\max(0, v_l - p_l) + (p_h - c_h)[\frac{1}{2}\max(0, v_h - p_h) + \frac{1}{2}\cdot\max(0, v_h - p_h - L(p_h, p_l))]$$

When loss aversion is sufficiently weak, consumer demand is guaranteed to be strictly positive for all realizations of actual and reference prices, such that the objective function is further simplified into:

$$(p_l - c_l)(v_l - p_l) + (p_h - c_h)[v_h - p_h - \frac{1}{2}L(p_h, p_l)]$$

Let us adopt the following specification of the loss function:

$$L(p, p^e) = \delta + \lambda \cdot (p - p^e)$$

whenever $p > p^e$. When $\lambda = 0$, we are back with the case of extreme loss aversion introduced in Sections 3.1-3.2. When λ and δ are sufficiently small, the solution to the firm's maximization problem, denoted (p_l^*, p_h^*), satisfies $p_h^* > p_l^*$ and is given by first-order conditions:

$$\begin{aligned}
\frac{\partial\Pi(p_l^*,p_h^*)}{\partial p_l} &= v_l + c_l - 2p_l^* + \frac{\lambda}{2}(p_h^* - c_h) = 0 \\
\frac{\partial\Pi(p_l^*, p_h^*)}{\partial p_h} &= v_h + c_h - 2p_h^* - \frac{\lambda}{2}(p_h^* - c_h) - \frac{1}{2}[\delta + \lambda\cdot(p_h^* - p_l^*)] = 0
\end{aligned}$$

leading to the solution:

$$\begin{aligned}
p_l^* &= \frac{1}{8\lambda - \lambda^2 + 16}\{[(8 + 4\lambda)(v_l + c_l) + 2\lambda(v_h + c_h) - \lambda\delta] - c_h(4\lambda + \lambda^2)\} \\
p_h^* &= \frac{1}{8\lambda - \lambda^2 + 16}\{[8(v_h + c_h) + 2\lambda(v_l + c_l) - 4\delta] + c_h(4\lambda - \lambda^2)\}
\end{aligned}$$

These expressions allow us to compare two environments, A and B, that are identical in every respect except that in environment A fluctuations are in costs whereas in environment B fluctuations are in demand:

$$\begin{aligned} v_l^A &= \bar{v} & \qquad v_l^B &= \bar{v} - \varepsilon \\ v_h^A &= \bar{v} & \qquad v_h^B &= \bar{v} + \varepsilon \\ c_l^A &= \bar{c} - \varepsilon & \qquad c_l^B &= \bar{c} \\ c_h^A &= \bar{c} + \varepsilon & \qquad c_h^B &= \bar{c} \end{aligned}$$

Both environments share the same expected values of c and v. Moreover, $v_l^A + c_l^A = v_l^B + c_l^B = \bar{v} + \bar{c} - \varepsilon$ and $v_h^A + c_h^A = v_h^B + c_h^B = \bar{v} + \bar{c} + \varepsilon$. When $\lambda > 0$, in environment B, p_h^* is lower and p_l^* is higher than in environment A. In other words, the optimal pricing strategy displays greater rigidity in environment B. Note, however, that if $\lambda = 0$ - i.e., we are in the case of extreme loss aversion - this difference disappears.

The intuition for this result is as follows. The loss aversion term is contributed by the event in which the consumer's reference price is p_l and the actual price he faces is p_h. When the firm contemplates raising p_l or lowering p_h, it does so to curb the loss aversion term. The mark-up in state h, $p_h - c_h$, is higher when fluctuations are in demand rather than in costs. Therefore, the firm's gain from curbing the loss aversion term and thus raising expected demand in state h is higher in environment B. As a result, the firm's incentive to narrow the price range is stronger in the case of demand shocks.

5.2 Preference for a Fully Rigid Price

Let us now turn back to the case of an arbitrary number m of states, and compare two environments, A and B, which are related to each other as follows. First, the state space in both environments has the same cardinality m. Second, $v_\omega^A + c_\omega^A = v_\omega^B + c_\omega^B$ for every ω. Third, $\bar{c}^A = \bar{c}^B = \bar{c}$ and $\bar{v}^A = \bar{v}^B = \bar{v}$. Finally, $v_\omega^A = \bar{v}^A$ and $c_\omega^B = \bar{c}^B$ across all ω. Thus, the two environments are identical in every respect, except that in environment A the fluctuations are in costs whereas in environment B the fluctuations are in demand.

Our next result reveals a sense in which the incentive to employ rigid pricing strategies is stronger under demand fluctuations. We will say that a pricing strategy P is *regular* if it satisfies two properties: *(i)* it weakly increases with c and v; *(ii)* it induces strictly positive consumer demand for all realizations of actual and reference prices.

Proposition 4 *If the firm prefers the optimal constant price to a regular pricing strategy P in environment A, then it must have the same preference in environment B.*

Proof. The firm's expected profit from the optimal constant price $\bar{p}$ is the same under both environments. Thus, we only need to show that the firm's expected profit from an arbitrary regular pricing strategy P is higher in environment A than in environment B. Since P is regular, we can rewrite its expected profit in each environment (omitting the multiplicative term $1/m^2$) as follows:

$$\begin{aligned}\Pi(P) &= \sum_{\omega}[P(\omega)-c_\omega][v_\omega - P(\omega)] - \sum_{\omega}(P(\omega)-c_\omega)\sum_{\psi} L(P(\omega),P(\psi)) \\ &= \sum_{\omega} P(\omega)[v_\omega + c_\omega - P(\omega)] - \sum_{\omega} c_\omega v_\omega \\ &\quad -\sum_{\omega}\sum_{\psi} P(\omega)\cdot L(P(\omega),P(\psi)) + \sum_{\omega}\sum_{\psi} c_\omega \cdot L(P(\omega),P(\psi))\end{aligned}$$

The first term in the final expression for $\Pi(P)$ is identical for both environments, by the assumption that $v_\omega^A + c_\omega^A = v_\omega^B + c_\omega^B$ for every ω. The second term is identical for both environments, by the assumption that $\bar{c}^A = \bar{c}^B = \bar{c}$ and $\bar{v}^A = \bar{v}^B = \bar{v}$, coupled with the assumption that $v_\omega^A = \bar{v}^A$ and $c_\omega^B = \bar{c}^B$ across all ω. The third term is identical for both environments because it is only a function of the pricing function and not of underlying cost and demand parameters. It therefore remains to compare the fourth term under both environments. Let us rewrite this fourth term as follows:

$$\sum_{\omega} c_\omega L_\omega^* \tag{6}$$

where

$$L_\omega^* = \sum_{\psi} L(P(\omega),P(\psi))$$

Order the states in Ω according to their $v_\omega + c_\omega$ (which is the same in both environments A and B). By assumption, P is weakly increasing in $v_\omega + c_\omega$. Therefore, L_ω^* is also increasing in $v_\omega + c_\omega$. By assumption, both environments A and B share the same $\sum_\omega c_\omega$, yet $c_\omega^B = \bar{c}$ for all ω, whereas c_ω^A increases with $v_\omega + c_\omega$. It follows that expression (6) is higher in environment A. Therefore, the firm's expected profit from P is higher in environment A. ■

The intuition for this result - as in the two-state example - is that when the firm contemplates a small change in the price in some state, it is mindful of the implications of this price change for the loss aversion term in other states associated with higher prices. When fluctuations are in demand, the mark-up in those high-price states is higher than when fluctuations are in cost, and therefore the incentive to shrink the gap between the high and low prices is stronger.

6 Discussion

This section is divided into two parts. First, I conduct a detailed comparison between the model presented here and the original HKR model. Second, I briefly discuss an extension in which consumers have a taste for pleasant price surprises.

6.1 Comparison with HKR

As explained in the Introduction, the model presented in this paper shares the economic motivation and basic modelling idea with HK, yet it departs from their model in three dimensions.

Which aspects of the market outcome are relevant for loss aversion?
In the HKR model, consumers display loss aversion both in the price dimension and in the consumption quantity dimension. The former captures the distaste for unpleasant price surprises, which is the focus of this paper. The latter, however, captures a distinct phenomenon, which is close both formally and psychologically to the well-known "endowment effect": the consumer experiences a disutility if his consumption quantity is lower than expected.

The fact that both effects can be classified as instances of loss aversion attests to the power and generality of loss aversion as a theoretical construct. However, this does not change the fact that the two effects are distinct, and they may be relevant in different contexts. For instance, a sense of ownership seems more pertinent to durable goods than to perishable goods, and therefore I expect loss aversion in the consumption quantity dimension to be more relevant in the former case. Since we should not expect the two effects to be equally applicable to a given market situation, I see no obvious reason for incorporating both of them into the same model. As it happens, the two effects have contradictory pricing implications: distaste for price surprises leads to price rigidity, whereas the attachment effect may give the firm an incentive to randomize over prices. Thus, focusing on one of the effects while ignoring the other also makes it easier to obtain clear-cut results.

Does the consumer's expected action enter the specification of the reference point?
HKR assume that the consumer's reference point takes into account his own expectation of his own consumption decision. The consumer's decision is thus a "personal

equilibrium": the action he chooses maximizes the consumer's reference-dependent utility given the reference point induced by the very same action. In our context, this means that the consumer's reference price is not the price he expects the firm to charge, but the price he expects to end up paying. In particular, if he expects not to buy the product, then an unpleasant price surprise does not generate a loss.

How should we "sum over" multiple reference points?
The original specification of Prospect Theory assumed a single reference point. It is not clear how one should extend the model when there are multiple candidates for a reference point. A market environment with stochastic prices naturally generates a big set of possible outcomes that can act as reference points. HKR assume that the decision maker "sums over" them as follows. He computes his reference-dependent expected utility from a given action for any possible reference point r, and then integrates over all values of r to obtain his evaluation of the action.

The difficulty with this procedure is that it is hard to think of a concrete scenario that generates it. In addition, the procedure presumes that the consumer is aware of all possible reference points and nevertheless allows each of them to influence his evaluation of each action. I believe, however, that reference points are powerful when they are salient; the greater the number of possible reference points the decision maker is aware of, the lower the likelihood that his behavior will be sensitive to any of them. Of course, this is only my opinion and it is debatable.

The reason for my first two departures is clear: one can fruitfully analyze consumers' aversion to unpleasant price surprises and its implications for price rigidity, without entering the complications that arise when we (1) introduce loss aversion into the consumer's evaluation of other aspects of the market outcome, and (2) allow the consumer's expectation of his own consumption decision to enter the specification of the reference point. In this sense, the model in this paper is a simplification of HKR, which is therefore useful for pedagogical purposes and for certain applications.

The third departure from HKR is a modification of their model of the formation of reference points. One advantage of the "sampling-based" model of reference-point formation is that it is based on a concrete scenario that describes how each consumer gets to form his reference point: the consumer observes (directly or through word of mouth) one past market outcome and this becomes his reference point. Another advantage is that the resulting model of consumer choice treats the reference point as if it were a "consumer type". To obtain aggregate consumer demand, we simply integrate over all possible reference points, just as we integrate over all possible "standard" preference types. The difference is that the distribution of reference points, unlike the distribution of preference types, is influenced by the firm's strategy.

6.2 Pleasant Price Surprises[4]

Throughout this paper, we assumed that consumers' willingness to pay reacts only to unpleasant surprises - i.e. cases in which the actual price exceeds the expected price. One could argue that in many real-life situations, consumers' willingness to pay also reacts to pleasant surprises. For instance, they may have a "taste for bargains". when a consumer encounters a price that is lower than expected, he may view it as a "bargain" and this may spuriously enhance his willingness.

Consider the following simple version of our model. The firm's marginal cost is $c = 0$ with certainty. The consumer's "raw" willingness to pay is $u = 1$ with certainty. The firm can employ a *random* pricing strategy - namely, a lottery μ over prices. In this case, the consumer's reference point is p^e with probability $\mu(p^e)$. Let $\lambda_g, \lambda_l \geq 0$. The consumer is willing to buy the product at a price p if $p \leq v + \lambda_g \cdot (p^e - p)$ when $p < p^e$, and if $p \leq v - \lambda_l \cdot (p - p^e)$ when $p > p^e$. Thus, the consumer's willingness to purchase is affected by both pleasant and unpleasant surprises.

If $\lambda_g = \lambda_l = 0$, the firm's optimal pricing strategy is $p^* = 1$. If $\lambda_g = 0$ and $\lambda_l > 0$, we are back with our original model (except that c and u are deterministic). The firm has a strict disincentive to randomize over prices in this case. To see why, observe that the consumer's willingness to pay cannot exceed one. If the firm assigns positive probability to prices above one, the consumer never buys at those prices, while if the firm assigns positive probability to prices below one, it earns by definition less than one (and in addition, the consumer may fail to buy the product, because even if $p < 1$, it is not necessarily the lowest price in the support of the price distribution, and can therefore constitute an unpleasant surprise that reduces the consumer's willingness to pay).

Now suppose that $\lambda_g > 0$, and consider a random pricing strategy that assigns probability α to p_1 and probability $1 - \alpha$ to p_2, where p_1 and p_2 satisfy

$$\begin{aligned} p_2 &> p_1 > 4 \\ 1 + \lambda_g \cdot (p_2 - p_1) &> p_1 \end{aligned}$$

[4] This sub-section is based on a suggestion by Kfir Eliaz.

Then, the firm's expected profit is

$$\alpha \cdot p_1 \cdot [\alpha \cdot 0 + (1-\alpha) \cdot 1] + (1-\alpha) \cdot p_2 \cdot 0 = \alpha(1-\alpha)p_1$$

Thus, since $p_1 > 4$, the firm's expected profit exceeds the level it can reach without randomization. Note that since there are only two price levels in the support of the price distribution, and since both price levels exceed the consumer's raw willingness to pay, the consumer's loss aversion parameter λ_l is irrelevant. The consumer experiences a loss when the expected price is p_1 and the actual price is p_2. But since p_2 exceeds the consumer's "raw" willingness to pay, he would not buy the product at this price even if we set the loss aversion parameter to $\lambda_l = 0$.

The feature in this example that gives the firm a strict incentive to randomize is that the high price p_2 can be arbitrarily high, such that the contribution of the pleasant surprise to the consumer's willingness to pay is unbounded. However, the consumer does not buy the product at p_2. One could argue that prices that lead to no trade should not serve as reference points. This criticism makes sense if the sampling process that generates the reference points is based on actual transaction prices rather than stated prices. This would be in the spirit of the HKR assumption that the consumer's reference point incorporates his expectation regarding his own consumption decision. If we accepted this critique, then the rationale for randomization in this example would disappear.

7 Conclusion

My aim in this paper was to conduct a re-modelling exercise which I referred to as a "cover version" of the HKR model. Hopefully, the variation analyzed here contributes a different way of looking at the notion of reference-dependent consumer preferences. Its usefulness was demonstrated with a pair of novel results, concerning the impact of loss aversion on expected price and the difference between cost and demand fluctuations. An important ingredient in the HKR model, namely the notion of personal equilibrium, has been suppressed, because it is not needed in order to capture the implications of loss aversion for price rigidity. Combining such a notion with the "sampling-based" method of aggregating reference points is left for future research. See Spiegler (2010) for a proposal for such an extended model.

References

[1] Courty, P. and M. Pagliero (2008), "Price Variation Antagonism and Firm Pricing Policies", mimeo.

[2] Fehr, E., L. Goette and C. Zehnder (2009), "A Behavioral Approach to the Labor Market: The Role of Fairness Concerns", Annual Review of Economics 1, 355-84.

[3] Hall, R. and C. Hitch (1939), "Price Theory and Business Behavior", Oxford Economic Papers 2, 12–45.

[4] Heidhues, P. and B. Kőszegi (2005), "The Impact of Consumer Loss Aversion on Pricing", mimeo.

[5] Heidhues, P. and B. Kőszegi (2008), "Competition and Price Variation when Consumers are Loss Averse", American Economic Review 98, 1245-1268.

[6] Karle, H. and M. Peitz (2008), "Pricing and Information Disclosure in Markets with Loss-Averse Consumers", mimeo.

[7] Kőszegi, B. and M. Rabin (2006), "A Model of Reference-Dependent Preferences", Quarterly Journal of Economics 121, 1133-1166.

[8] Kahneman, D. and A. Tversky (1979), "Prospect Theory: An Analysis of Decision under Risk", Econometrica 47, 263–291.

[9] Kahneman, D., J. Knetsch and R. Thaler (1986), "Fairness as a Constraint on Profit Seeking: Entitlements in the Market", American Economic Review 76, 728–741.

[10] Okun, A. (1981), Prices and Quantities: A Macroeconomic Analysis (Washington, DC: The Brookings Institution).

[11] Osborne, M. and A. Rubinstein (1998), "Games with Procedurally Rational Players", American Economic Review 88, 834-849.

[12] Spiegler, R. (2006), "The Market for Quacks", Review of Economic Studies 73, 1113-1131.

[13] Spiegler, R. (2010), *Bounded Rationality and Industrial Organization*, Oxford University Press, forthcoming.