



Consumer bounded rationality and rigidity/flexibility retail price patterns[☆]

Ran Spiegler^{*}

Tel Aviv University, Israel
University College London, United Kingdom

ARTICLE INFO

Article history:

Received 9 March 2012

Accepted 2 April 2012

Available online 12 April 2012

JEL classification:

C79

D43

Keywords:

Bounded rationality

Industrial organization

Sales

Regular prices

Price rigidity

Sampling

ABSTRACT

I revisit the model of market competition with boundedly rational consumers due to Spiegler (2006), in which firms compete in price distributions and consumers use a naive sampling procedure to evaluate them. I assume that firms can assign weight to arbitrarily low prices, and consumers have a non-trivial ex ante outside option. In symmetric Nash equilibrium, firms charge a high “regular price” with positive probability, and in addition randomize continuously over an interval of “sale” prices that are bounded away from the regular price. Sales become less frequent but more drastic as the number of competitors increases and as the consumer’s outside option becomes more attractive.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

Retail prices seem to display patterns that combine rigidity and flexibility. Evidence generated from product-specific supermarket prices suggests that products are often characterized by a “regular price” which remains fixed for long stretches of time, yet subjected with some frequency to sales of varying magnitude. Thus, the regular price is rigid while the sales prices are flexible.¹ A number of explanations have recently been proposed. Kőszegi and Heidhues (2011) derive the rigidity/flexibility pattern as part of an optimal pricing strategy for a monopolist who faces loss averse consumers. Stevens (2012) argues that the pattern is an outcome of sellers’ optimal adaptation to a changing environment when information acquisition is costly.

In this brief paper, I propose an alternative interpretation of the rigidity/flexibility pattern, which is based on a model

due to Spiegler (2006).² In this model, n firms compete in probability distributions over prices, facing a homogeneous population of boundedly rational consumers. A price distribution can represent cross-section dispersion or variation over time. Both interpretations fit retail price settings. The cross-section interpretation is appropriate because most retailers (especially supermarkets) sell multiple products and can introduce sales in an arbitrary subset of those. The temporal-variation interpretation is appropriate because sales are typically temporary.

Evaluating the price distribution that characterizes each seller is a hard task, and in response boundedly rational consumers resort to simplifying heuristics. Following Osborne and Rubinstein (1998), I assume that each consumer evaluates any given seller by means of a single sample point drawn from the seller’s price distribution. The consumer naively extrapolates from the sample, as if she regards it as being “representative” of the distribution from which it was drawn. The consumer proceeds to choose the best alternative in her sample. However, the actual price that she ends up paying is a new, independent draw (or any sequence of such draws) from the chosen distribution.

This model fits retail environments in which the shopping experience is complex and involves many products. For example, in the case of supermarket shopping, the sampling procedure can be interpreted as follows. The consumer selects an arbitrary product

[☆] I thank Kfir Eliaz, Ariel Rubinstein and Jidong Zhou for helpful comments. Financial support from the European Research Council, Grant no. 230251, is gratefully acknowledged.

^{*} Correspondence to: Tel Aviv University, Israel.

E-mail addresses: rani@post.tau.ac.il, r.spiegler@ucl.ac.uk.

URL: <http://www.tau.ac.il/~rani>.

¹ Kőszegi and Heidhues (2011) survey the relevant literature. See Klenow and Malin (2010) for a large review which is oriented towards macroeconomic, business-cycle implications.

² See Spiegler (2011, Ch. 7) for a textbook exposition of this model.

(which the firm cannot predict) as a “representative” product, and decides to do her entire shopping at the supermarket that sells this product at the lowest price. The sampling procedure can also be interpreted as a trial period in which the consumer tries the various retailers, and decides to do her regular shopping at the retailer who came out cheapest in the trial phase. The trial phase is very short relative to the regular shopping phase.

The main finding in Spiegel (2006) was that in the unique symmetric Nash equilibrium, increasing n subjects the equilibrium price distribution to a mean-preserving spread. However, that paper made expositional assumptions that eliminated certain realistic features, and ended up obscuring flexibility/rigidity price patterns. This brief paper brings these features back into the model. First, sellers are allowed to assign positive probability to prices that are arbitrarily lower than marginal cost. Second, consumers have an ex ante outside option – that is, they may choose not to shop at any of the competing retailers. Spiegel (2006) allowed for these aspects in isolation, but not in combination. For expositional ease, I describe this outside option as a perfect substitute that consumers can obtain outside the market at a fixed price p_0 .³

The mean-preserving-spread effect continues to hold in symmetric equilibrium. However, now the equilibrium strategy exhibits a clear rigidity/flexibility pattern: it places an atom on a high “regular” price, which is the consumers’ underlying willingness to pay. In addition, it induces a continuous density over an interval of prices, the upper bound of which is p_0 . This interval represents variable “sale” prices, and if p_0 is strictly below the consumers’ willingness to pay, they are bounded away from the regular price. The equilibrium is robust to an extension in which sellers are uncertain of the outside option price, provided that its distribution is weakly first-order stochastically dominated by the uniform distribution over $[p_0, 1]$.

The intuition for this result is simple. Sale prices function as baits, and this is where competition among sellers takes place. If sellers placed an atom on a price in this range, it would be vulnerable to a standard undercutting deviation. In contrast, price realizations above p_0 cannot function as baits – consumers always prefer the outside option when they sample such prices. Therefore, there are no competitive pressures in the high-price range, and sellers prefer to shift weight from this range to the regular price.

I believe that this intuition captures some of the economics of rigidity/flexibility patterns in retail pricing. I should emphasize that unlike that of Stevens (2012), the present model is unable to say anything about the forces that determine the frequency with which the regular price itself changes. And unlike that of Kőszegi and Heidhues (2011), the present model requires competition to generate the rigidity/flexibility pattern. Indeed, the model predicts that as competition gets tougher – that is, when n increases or p_0 decreases – sales become less frequent but larger in magnitude.

2. The model

A market consists of a set $\{1, \dots, n\}$ of expected-profit maximizing firms and one consumer. The firms produce a homogeneous product at a constant marginal cost that is normalized to zero. They play a simultaneous-move, complete information game. A strategy for a firm is a cumulative distribution function (cdf) G_i over the set of feasible prices $(-\infty, 1]$. The consumer also has access to an ex ante outside option, denoted as 0, that gives her an equivalent product at a fixed price $p_0 \in [0, 1]$. The outside option becomes unavailable once the consumer chooses a firm. The assumption that prices have an upper bound is

³ It would be formally equivalent to assume that the outside option is an imperfect substitute that gives the consumer a net utility of $v_0 = 1 - p_0$.

interpreted as a limited liability or an ex post individual rationality constraint. Even if the consumer is somehow tricked into choosing a sub-optimal retailer, she cannot be forced to buy at a price that exceeds her willingness to pay or resources.

After the firms make their decisions, the consumer chooses an alternative from the set $\{0, 1, \dots, n\}$, according to the following sampling procedure, borrowed from Osborne and Rubinstein (1998). She draws one sample point p_i from each G_i , $i = 1, \dots, n$. She then chooses an alternative $i^* \in \arg \min_{i=0,1,\dots,n} p_i$ (resolving ties in favor of the outside option, and symmetrically among firms). If $i^* \neq 0$, the outcome of the consumer’s choice is a *new, independent draw* from G_{i^*} .

Let us construct firm i ’s payoff function, fixing the profile $(G_j)_{j \neq i}$. Define $H_i(p)$ as the probability that the consumer chooses firm i , conditional on $p_i = p$ in her sample. One may view H_i as the firm’s “residual demand function” given the opponents’ strategies. Let Ep_i denote the expected price according to G_i , and let EH_i denote the expectation of H_i with respect to G_i , namely firm i ’s market share. Then, firm i ’s payoff function is

$$u_i(G_1, \dots, G_n) = Ep_i \cdot EH_i.$$

Note that the payoff function is quadratic in dG_i . Thus, although strategies in this game are probability distributions, they are not “mixed strategies”. Rather, this is a game in which the players’ pure strategies are probability distributions. As Spiegel (2006) demonstrates, firms may strictly prefer to randomize. This strict preference for randomization will disappear in equilibrium. However, the familiar property of mixed-strategy Nash equilibrium, namely that every element in the support of an equilibrium strategy is a best reply, fails to hold in equilibrium.

3. Equilibrium

The following result characterizes symmetric Nash equilibrium in this game.

Proposition 1. *There is a unique symmetric Nash equilibrium in the game. Each firm plays the cdf*

$$G^*(p) = 1 - \left[\frac{1 - \min(p, p_0)}{1 - p^*} \right]^{1/(n-1)} \quad (1)$$

defined over $p \in [p^*, 1]$, and $G^*(1) = 1$, where $p^* < p_0$ is uniquely determined by the equation

$$(1 - p^*) \left[1 - \left(\frac{1 - p_0}{1 - p^*} \right)^{n/(n-1)} \right] = \frac{n}{2}.$$

Proof. The proof is an extension of the proof of Proposition 1 in Spiegel (2006). For every $i = 0, 1, \dots, n$ and every p_i , define $x_i = 1 - p_i$. We may interpret x_i as the consumer’s net utility from alternative i , if we equate the upper bound on prices with the consumers’ willingness to pay. It will be more convenient to translate the firms’ strategies into cdfs over x . That is, for every $p \in (-\infty, 1]$, define

$$F(1 - p) = 1 - G(p).$$

Note that the support of F is restricted to being contained in $[0, \infty)$. Let F^* be a symmetric Nash equilibrium strategy. Let S denote the support of F^* , and define $x^* = \sup(S)$. Slightly abusing notation, let H^* be the residual demand function induced by F^* . The proof proceeds stepwise.

Step 1: $S = \{0\} \cup [x_0, x^*]$. Furthermore, $x = 0$ is the only point at which F^* may place an atom.

Proof. By definition, $H^*(x) = 0$ for every $x \leq x_0$. If $x^* \leq x_0$, then firms earn zero profits, and it is profitable to deviate by shifting some weight to some $x > x_0$. Therefore, $x^* > x_0$. Similarly, if $(0, x_0) \cap S$ is non-empty, then it is profitable to deviate by shifting weight from $(0, x_0)$ to $x = 0$; this deviation does not change the deviating firm’s market share and it increases its expected share in the surplus conditional on being chosen.

If F^* places an atom on some $x \geq x_0$, then it is profitable to deviate by shifting weight from x slightly upward. Define $\tilde{x} = \inf(S \setminus \{0\})$. We have seen that $\tilde{x} \geq x_0$. Suppose that this inequality is strict. Consider a deviation by firm 1, say, that shifts all the weight that F^* assigns to $(\tilde{x}, \tilde{x} + \varepsilon)$ to x_0 . This deviation reduces Ex_1 by at least $(\tilde{x} - x_0) \cdot (F^*(\tilde{x} + \varepsilon) - F^*(\tilde{x}))$. At the same time, it reduces EH_1 by at most $(F^*(\tilde{x} + \varepsilon) - F^*(\tilde{x}))^n$. If $\varepsilon > 0$ is sufficiently small, the reduction in EH_1 is negligible in comparison to the reduction in Ex_1 , such that the deviation is profitable. A similar argument establishes that S contains no holes in $[x_0, x^*]$. \square

Step 2: The residual demand function H^* is defined over S by

$$H^*(x) = \frac{x}{x^*}. \tag{2}$$

Proof. Consider three elements in S , $x_1 < x_2 < x_3$, and suppose that the point $(x_2, H^*(x_2)) \in \mathbb{R}^2$ lies above the line connecting the points $(x_1, H^*(x_1))$ and $(x_3, H^*(x_3))$. Then, it is profitable to deviate by shifting weight from the neighborhoods of x_1 and x_3 to x_2 , in a way that preserves Ex . This deviation increases EH and therefore the deviating firm’s payoff. Conversely, if $(x_2, H^*(x_2)) \in \mathbb{R}^2$ lies below the line connecting $(x_1, H^*(x_1))$ and $(x_3, H^*(x_3))$, it is profitable to deviate by shifting weight from the neighborhood of x_2 to x_1 and x_3 in a way that preserves the expectation of x . It follows that H^* is linear over S . The formula for H^* then follows immediately from the facts that $H^*(0) = 0$ and $H^*(x^*) = 1$. \square

Step 3: $Ex = \frac{1}{2}$.

Proof. Since $H^*(x) < x/x^*$ for $x \notin S$, the support of any best-replying strategy for any firm i must be contained in S . Thus, any firm i chooses its cdf F_i over S to maximize $(1 - Ex_i) \cdot EH_i$, where the expectations are taken w.r.t. F_i . By Step 2, the objective function can be rewritten as $(1 - Ex_i) \cdot Ex_i$, because x^* is a constant as far as firm i is concerned. It follows that any cdf F_i over S that satisfies $Ex_i = \frac{1}{2}$ is a best reply. \square

Define $A = F^*(0)$. To derive the formula for F^* , observe that by Step 3,

$$0 \cdot A + \int_{x_0}^{x^*} x dF^*(x) = \frac{1}{2}. \tag{3}$$

By Step 1, F^* is continuous over $(x_0, x^*]$; hence for every x in this interval,

$$H^*(x) = [F^*(x)]^{n-1}. \tag{4}$$

In particular,

$$\lim_{x \rightarrow x_0^+} H^*(x) = A^{n-1}. \tag{5}$$

The expression for F^* can be derived by plugging (2) and (4) into (3), and using (5) to solve for A and x^* . Translating back into the language of cdf s over prices, we obtain (1). \square

Thus, symmetric equilibrium exhibits a rigidity/flexibility pattern whenever $p_0 < 1$ – that is, when consumers have a non-trivial outside option. The equilibrium strategy places an atom of size

$$A = \left(\frac{1 - p_0}{1 - p^*} \right)^{1/(n-1)} \tag{6}$$

on the price $p = 1$. Note that A can be written as the unique solution in $[0, 1]$ of the equation

$$nA^{n-1} = 2(1 - p_0)(1 - A^n).$$

The equilibrium strategy assigns zero probability to the interval $[p_0, 1)$, and it is continuous and strictly increasing over the interval $[p^*, p_0)$, such that firms face a linear residual demand function over this interval. When $p_0 = 1$, this rigidity/flexibility pattern disappears and the equilibrium price distribution is entirely smooth. Indeed, it collapses to the basic characterization in Spiegel (2006).

The atom on $p = 1$ can be interpreted as a “regular” price, whereas the realizations in $[p^*, p_0)$ can be interpreted as variable “sale” prices. In particular, note that when $n > 2$ or $p_0 < 1$, we have $p^* < 0$ – that is, firms assign positive probability to price realizations below marginal cost, in a way that is reminiscent of loss-leader pricing.

The intuition behind this pattern is simple. Every realization of a cdf potentially has two functions: it attracts clients to the firm and generates revenues from the firm’s clientele. The consumers’ sampling procedure implies that these two functions are somewhat disconnected. When a firm contemplates assigning weight to a realization $p \in (p^*, p_0)$, it knows that this realization will not generate a clientele – that is, if a consumer samples this realization, she will necessarily prefer the outside option. Therefore, such a realization only generates revenue for the firm. But this means that it is better for the firm to shift this weight to $p = 1$. In other words, there are no competitive forces in the region $(p_0, 1)$.

In contrast, the realizations $p < p_0$ generate a clientele in equilibrium. Moreover, a realization $p < 0$ generates losses and its only function can be to generate a clientele. There is fierce competition among firms in this region. An atom on some $p < p_0$ is inconsistent with equilibrium, because any firm can profitably deviate by shifting this weight to a slightly lower realization, thereby increasing its clientele (this is a conventional “undercutting” argument).

The following corollary presents a few comparative statics results.

Corollary 1. *When n increases or p_0 decreases:*

- (i) G^* undergoes a mean-preserving spread, such that Ep remains constant at $\frac{1}{2}$.
- (ii) A increases.
- (iii) $E(p \mid p < p_0)$ and p^* decrease.

The first result merely extends the main finding of Spiegel (2006) to the current setting. The other results are novel. As competition gets tougher (either because the number of competitors rises, or because consumers have a better ex ante outside option), The mean-preserving-spread property implies that A goes up – that is, the regular price is more ubiquitous and sales are less frequent. At the same time, when sales do occur, they are more drastic: the expected realization p conditional on being in the “sales” region $[p^*, p_0)$ is lower, and p^* itself is lower.

Comment: a stochastic outside option

The assumption that firms know p_0 is of course strong. However, the equilibrium that we derived would continue to hold if we assumed instead that firms have a common belief that the outside option price is drawn from the interval $[p_0, 1]$ according to some cdf G_0 , as long as G_0 is first-order stochastically (weakly) dominated by $U[p_0, 1]$ – that is,

$$G_0(p) \geq \frac{p - p_0}{1 - p_0}$$

for every $p \in [p_0, 1]$. This property guarantees that the residual demand function coincides with the linear formula obtained in the proof of the proposition for $p \in \{1\} \cup [p^*, p_0]$, but lies below the extension of that formula to the range $(p_0, 1)$. The implication is that firms would never want to shift weight from the support of G^* to the interval $(p_0, 1)$.

References

- Klenow, P., Malin, B., 2010. Microeconomic evidence on price-setting. In: Friedman, B., Woodford, M. (Eds.), *Handbook of Monetary Economics*. Elsevier Press, Amsterdam.
- Kőszegi, B., Heidhues, P., 2011. Regular prices and sales, mimeo.
- Osborne, M., Rubinstein, A., 1998. Games with procedurally rational players. *American Economic Review* 88, 834–849.
- Spiegler, R., 2006. Competition over agents with boundedly rational expectations. *Theoretical Economics* 1, 207–231.
- Spiegler, R., 2011. *Bounded Rationality and Industrial Organization*. Oxford University Press, New York.
- Stevens, L., 2012. Price adjustment in a model with multiple-price policies, mimeo.