Competition over Agents with Boundedly Rational Expectations*

Ran Spiegler†

First version: April 2004
This version: November 2005

Abstract

I study a market model in which profit-maximizing firms compete in multi-dimensional pricing strategies over a consumer, who is limited in his ability to grasp such complicated objects and therefore uses a sampling procedure to evaluate them. Firms respond to increased competition with an increased effort to obfuscate, rather than with more competitive pricing. As a result, consumer welfare is not enhanced and may even deteriorate. Specifically, when firms control both the price and the quality of each dimension, and there are diminishing returns to quality, increased competition implies an efficiency loss which is entirely borne by consumers.

*I would like to thank the LSE for its stimulating hospitality during the fall term of 2003/4. I have greatly benefitted from numerous discussions with Eddie Dekel and Kfir Eliaz. Finally, I thank Gil Kalai, Michele Piccione, Ronny Razin, Ariel Rubinstein, Nicolas Sahuguet, Asher Wolinsky, numerous seminar audiences, two referees and especially Barton Lipman, for helpful comments.

†School of Economics, Tel Aviv University, Tel Aviv 69978, Israel. E-mail: rani@post.tau.ac.il. URL: Http://www.tau.ac.il/∼rani.

1 Introduction

Economists have grown familiar with market models with informational asymmetries between firms and consumers: markets for experience goods in which firms know their quality better than consumers, insurance markets in which consumers know their risks better than firms, etc. But while these models allow for informational asymmetries, they impose a *perceptual* symmetry between firms and consumers, because they retain the assumption that "the model itself is common knowledge".

In reality, firms and consumers often differ in their ability to understand the market model. Firms interact more frequently with the market, and pay closer attention to it, than most consumers. As a result, they have more opportunities to learn the market model and the market equilibrium. Moreover, because prices are typically set by firms, they are in a position to complicate the consumer's task of understanding the actual value of their products, by employing complex pricing schedules.[1]

For instance, consider the problem of choosing where to open a bank account. Banks offer a large number of financial services. At the time we open the account, we do not know yet which subset of services will be relevant for us. The bank can complicate our decision problem by adopting different fees for different transactions, or different interest rates for different types of saving accounts. Likewise, when we purchase life or health insurance, we need to calculate trade-offs across a large number of scenarios. Insurance companies can contribute to the difficulty of this task by applying different reimbursement policies to different contingencies.

In these examples, firms employ strategies with a potentially complex, multi-dimensional structure. Consumers find it difficult to grasp this structure in its entirety. Therefore, it is natural for them to resort to simplifying heuristics. An example of such a heuristic is to sample a small number of dimensions and choose the best-performing firm along the sampled dimensions. This heuristic can be applied in many market settings, hence it saves considerable cognitive resources.

One could argue that repeated exposure to the market would enable consumers to learn the true value of each alternative, thus saving them the need

[1] As Ellison and Ellison (2004) demonstrate, even in a potentially competitive environment such as e-commerce, internet retailers react to price search engines - whose objective is to reduce the complexity of consumers' decision problem - with a variety of obfuscation devices that partly regain this complexity.

to rely on simplifying short-cuts. However, these learning opportunities are scarce in situations such as those described above. To quote Camerer and Lowenstein (2003, p. 8-9), "many important aspects of economic life are like the first few periods of an experiment rather than the last".

My objective in this paper is to examine the implications of market competition, when consumers reason in this way about multi-dimensional goods and services. Here are some of the questions that I raise: What kind of "obfuscation devices" do firms use to manipulate consumers? Will competition among firms weaken, or rather strengthen the effort to obfuscate? What are the social costs of obfuscation? To the extent that firms take advantage of consumers, will competition among firms mitigate this exploitative effect?

To address these questions, I study a simple market model with one consumer and n firms. A strategy for a firm is a *cdf* F over a set of feasible prices. I interpret F as a reduced-form representation of a multi-dimensional pricing strategy. The firm provides a product having a large number of equally important, or equally likely dimensions. It prices each dimension independently, such that $F(p)$ is the probability that the price of a randomly sampled dimension is at most p.

Firms are standard expected-profit maximizers. The consumer, in contrast, employs a "boundedly rational" choice procedure, called $S(1)$ and borrowed from Osborne and Rubinstein (1998). Faced with a profile of *cdf*s, he draws one sample point from each *cdf*, and chooses the cheapest firm in his sample. The actual outcome of his choice is a *new*, independent draw from the chosen *cdf*. The interpretation is that the consumer evaluates a multi-dimensional pricing strategy by examining one dimension at random and choosing the best-performing firm along that dimension. Firms take into account the consumers' choice procedure when choosing their strategy.

If firms were restricted to degenerate distributions, the consumer would always make the optimal choice and the market would be truly competitive. This suggests that firms have an incentive to introduce variance into their *cdf*, in order to make it harder for the consumer to perceive their true expected price. Thus, the firms' strategic considerations involve two effects: competing over the consumer, and trying to take advantage of his inference errors. The question is how the "competitive" and "obfuscatory" effects interact.

The characterization of symmetric Nash equilibrium in Section 3 provides a sharp answer. There is a unique symmetric equilibrium, given by a simple formula. Expected price is independent of the number of competitors: the consumer receives half the surplus, regardless of n. Moreover, when we add a firm to the market, the equilibrium cdf is a mean-preserving spread of the original one. Thus, *firms respond to greater competition with greater obfuscation, rather than with more competitive pricing.*

In Sections 4 and 5, I analyze two extensions of the model. First, I assume that firms control both price and quality, and they choose a probability distribution over price-quality pairs. I assume diminishing returns to quality. This extended model has a unique symmetric Nash equilibrium. If n is sufficiently large, the outcome is *inefficient* in terms of expected surplus. Moreover, the efficiency loss increases with the number of firms, and is borne *entirely* by the consumer. Second, I endow the consumer with an outside option, and show that in equilibrium, he may be worse off than if he stayed out of the market. This effect, too, is exacerbated as n gets larger.

The lesson from these results is simple. Interventions that foster competition in a market with rational consumers (increasing the number of competitors, introducing an attractive outside option) may have adverse welfare effects when consumers are limited in their ability to evaluate complex objects. The reason is that firms respond to increased competition by obfuscating, rather than by acting more competitively.

Related literature

Osborne and Rubinstein (1998) formulated the $S(1)$ procedure in the context of strategic-form games, in which *all* players behave according to this procedure, which calls for a novel equilibrium concept. In contrast, in the present paper, the strategic agents (i.e., the firms) are rational, hence the need does not arise. In Spiegler (2003), I study a market in which providers of a worthless treatment (in the sense of having the same success rate as a default option) compete in prices over consumers who evaluate treatments according to $S(1)$. As a result, this "market for quacks" becomes active and displays non-standard welfare and comparative-statics properties. The main difference between Spiegler (2003) and the present study is that a pricing strategy in the former is a scalar, and the source of randomness for consumers is exogenous. As a result, the equilibrium characterization techniques are conventional and the welfare properties are milder: while the firms' ability to obfuscate in the present paper results in a non-competitive outcome for any n, in Spiegler (2003) the market outcome is competitive when $n \to \infty$.

This work belongs to a small group of papers which study market interaction between rational firms and consumers with cognitive imperfections: bounded ability to grasp intertemporal patterns in Piccione and Rubinstein

(2003); limited memory in Chen, Iyer and Pazgal (2003); and biased beliefs concerning future tastes in DellaVigna and Malmendier (2004) and Eliaz and Spiegler (2004).

Within this literature, some works have examined one of the themes of the present paper, namely the incentive to obfuscate when consumers have a limited perception of complex objects. Rubinstein (1993) demonstrates that a monopolist may use a probabilistic pricing strategy, in order to screen the consumer's ability to categorize the realization of a random variable. Erev and Haruvy (2001) argue that when consumers evaluate alternatives with double exponential noise, lower-quality firms have a stronger motive to increase the variance parameter of the noise. In Gabaix and Laibson (2005), firms offer two-attribute products, and one attribute is hidden from some consumers. Competitive forces bring equilibrium profits down to zero. However, the pricing structure reflects the consumers' bounded rationality: the attribute of which all consumers are aware is priced below marginal cost, while the price of the hidden attribute is maximal.

It should be noted that the randomization motive per se can be accounted for by models with rational consumers. For instance, Salop (1977) demonstrates that a monopolist may wish to randomize in order to discriminate between consumers with diverse search costs. Wilson (1988) derives a probabilistic pricing strategy from the assumption that consumers arrive in a random order and are served on a first-come-first-served basis.

Finally, this paper is related to the large literature on equilibrium price dispersion (e.g., Varian (1980), Burdett and Judd (1983), Rob (1985), McAfee (1995), Burdett and Coles (1997)), which analyzes price competition in the face of rational consumers with search costs. It should be emphasized that price dispersion in these models is an artifact of mixed-strategy equilibrium. In contrast, in the present paper firms have a strict incentive to randomize.

2 A basic model

A market consists of a set $\{1, ..., n\}$ of expected-profit maximizing firms and one consumer. The firms play a simultaneous-move, complete information game. A strategy for a firm is a cumulative distribution function (*cdf*) F_i over a set of feasible prices $(-\infty, 1]$. Let T_i and Ep_i denote the support of F_i and the expected price according to F_i, respectively.

After the firms make their decisions, the consumer chooses an alternative from the set $\{1, ..., n\}$, according to a procedure called $S(1)$. He draws one sample point from each F_i. Given a sample of prices $(p_1, ..., p_n)$, he chooses $i^* \in \arg\min_i p_i$ (with the usual symmetric tie-breaking rule). The outcome of the consumer's choice is a *new, independent draw* from F_{i^*}. Firms take the consumer's decision rule into account when calculating their expected profit from a given strategy profile.

This is a stylized model that is open to a number of interpretations. The primary interpretation that I adopt in this paper is that F_i is a reduced-form representation of a multi-dimensional pricing strategy. Firms offer a service that covers a continuum of equally likely contingencies. A strategy for a firm specifies a price for every contingency, such that $F_i(p)$ is the fraction of contingencies for which the price is at most p. Evaluating the firm's pricing strategy is a difficult task for the consumer, because it requires him to calculate trade-offs across an enormous number of equally likely contingencies. The consumer simplifies his decision problem by examining one contingency *at random* and choosing the firm that offers the best terms in this particular contingency. It is crucial for this interpretation that the firms cannot predict the contingency that the consumer will sample during his deliberation. If firms knew it in advance, they would compete fiercely along that dimension, while charging the maximal price in all other contingencies.[2]

Let us construct firm i's payoff function, fixing the profile $(F_j)_{j \neq i}$. Define $H_i(p)$ as the probability that the consumer chooses firm i, conditional on $p_i = p$ in his sample. One may view $H_i(p)$ as the demand for firm i, induced by the other firms' behavior. Let EH_i denote the expected value of H_i, where the expectation is taken with respect to F_i. When each firm plays a *cdf* with a well-defined density f, the definition of H_i and EH_i has a simple form: $H_i(p) = \Pi_{j \neq i}[1 - F_j(p)]$ and $EH_i = \int_{-\infty}^{1} H_i(p) f_i(p) dp$. For simplicity, I assume zero costs. Then, firm i's payoff function is $u_i(F_1, ..., F_n) = Ep_i \cdot EH_i$. In other words, the firm's payoff is equal to its expected price multiplied by the probability that it is chosen by the consumer.

The firms' strategies in this model are probability distributions. However, they are *not* "mixed strategies" (with $(-\infty, 1)$ being the purported set of "pure" strategies). Technically, this is because u_i is quadratic, rather than linear, in F_i. In particular, we should not expect firm i to be indifferent

[2] Alternatively, one could assume that firms know that the consumer samples a particular set of contingencies C^*, as well as another random contingency outside this set. In this case, the model describes how firms price the contingencies outside C^*.

between F_i and the individual elements in T_i. Indeed, the following example illustrates that firms may have a strict incentive to randomize.

Let $n = 3$. Suppose that firms 1 and 2 both play the *cdf* $F \equiv U[0, 1]$. If firm 3 assigns probability one to some price $\beta \in (0, 1)$, then its expected revenue is $\beta \cdot (1 - \beta)^2$. Now suppose that the firm switches to what might be called a "*quasi-bait-and-switch*" strategy: assigning probability β to $p = 0$ and probability $1 - \beta$ to $p = 1$. In this case, the firm's expected revenue is $[\beta \cdot 0 + (1 - \beta) \cdot 1] \cdot \beta$, a strict improvement.

I interpret this type of randomization as an obfuscation device. The firm prices some contingencies "competitively" and other contingencies "monopolistically". The former contingencies generate a clientele, whereas the latter contingencies generate revenue. The obfuscation inherent in this quasi-bait-and-switch strategy is extreme, because the consumer ends up choosing firm 3 if and only if $p_3 = 0$ in the contingency that he sampled, yet with probability $1 - \beta$, the actual price he ends up paying is the highest in the market. The quasi-bait-and-switch pricing strategy turns out to play an important role in the analysis.

This example also illustrates the fundamental difference between the role of sampling in this paper and in more conventional I.O. models with consumer search (e.g., Burdett and Judd (1983)). In both cases, the consumer chooses the cheapest firm i^* in his sample. However, in a search model the consumer ends up paying p_{i^*}, whereas in the present model, the actual price is a new, independent draw from F_{i^*}.

Explicit multi-dimensional pricing

The following, more elaborate model may be viewed as a "foundation" for the interpretation of F_i. Suppose that a pure strategy for firm i is explicitly modeled as a function from a continuum of contingencies to a set of feasible prices. Such a pure strategy induces a *cdf* F_i over prices, where $F_i(p)$ is the measure of contingencies for which the price assigned by the pure strategy does not exceed p. Our interpretation of the consumer's choice procedure becomes explicit.

If each firm i plays a mixed strategy which randomizes uniformly over all pure strategies that induce a given *cdf* F_i , then the elaborate formalism is reduced to our model. The set of pure-strategy equilibria in the original model is thus isomorphic to a natural subclass of mixed-strategy equilibria in the elaborate model.[3] What enables this linkage is that in both types

[3] It can be shown that there exists no pure-strategy equilibrium in the elaborate model.

of equilibria, firms know the fraction of contingencies in which opponents charge any given price, without knowing the exact contingencies for which this is the case.

An alternative interpretation

We may also think of F_i as a genuinely random pricing strategy, rather than a short-hand for a multi-dimensional pricing strategy. In this context, $S(1)$ captures the behavior of an inexperienced consumer who makes a once-and-for-all decision following a brief learning phase. In many situations, consumers enter the market to seek an ad-hoc solution to a specific problem, and therefore lack the opportunity to learn the expected value of the firms' products through repeated purchases.

For example, people typically seek services such as criminal defense litigation, real-estate appraisal, or building contracting very infrequently. Professionals in these fields effectively control the price (and effort level) that they apply to a given case. Moreover, as far as consumers are concerned, these are random variables because the firm can discriminate among cases on the basis of criteria that are hidden from the eyes of an inexperienced consumer. Therefore, consumers often have no choice but to evaluate each firm on the basis of a small collection of "anecdotes": random stories generated from fellow consumers' experience. It is natural for the consumer to pick the expert with the best anecdote in his sample. But the outcome of the consumer's decision will be an independent draw from the distribution associated with this expert.

Although I find both the "multi-dimensionality" and "learning" interpretations of the model plausible, I adhere to the former throughout the rest of the paper, for the sake of expositional clarity.

Negative prices

The assumption that the support of F_i can contain arbitrarily negative prices raises at least two concerns. First, if a consumer samples a negative price from F_i, he may realize that this sample point cannot represent the firm's pricing in the other dimensions, for the firm would go out of business if it did. Thus, a negative price may lead the consumer to rethink his choice procedure.[4] Second, when F_i represents how the firm prices a bundle of

[4]In reply, one may argue that since we have normalized the firms' marginal cost to zero, "negative" prices that are sufficiently close to zero represent positive prices below marginal cost. If the consumer is ignorant of the cost structure, he need not react to a

goods, negative prices expose the firm to counter-exploitation by a rational "parasite" who will purchase a large quantity of negatively priced goods.

At any rate, as we shall see in Section 3, the assumption plays a minor role in the analysis. It simplifies the formula of the symmetric equilibrium strategy as well as the proof, without affecting the main findings. Appendix II analyzes the model under the assumption that negative prices are ruled out.

3 Equilibrium

I begin with a characterization of symmetric Nash equilibrium in this model. Let $F(p,n)$ be the equilibrium *cdf*, and let $Ep(n)$ denote the expected price according to $F(p,n)$.

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

$$F(p,n) = 1 - \left[\frac{2(1-p)}{n}\right]^{\frac{1}{n-1}} \tag{1}$$

over the support $[1-\frac{n}{2}, 1]$.

Corollary 1 *For every* $n \geq 2$*:*

(i) $Ep(n) = \frac{1}{2}$

(ii) $F(p,n+1)$ *is a mean preserving spread of* $F(p,n)$.

Corollary 1 demonstrates that *an increase in* n *results in an increase in the variance of the equilibrium cdf, without affecting the expected price.* Thus, the number of competitors, normally an indicator of the market's competitiveness, has an orthogonal effect when the consumer evaluates firms according to the $S(1)$ procedure.

low sampled price by questioning his choice procedure.

Firms in the model have two strategic considerations. First, there is the usual competitive motive, which induces firms to offer attractive distributions. Second, there is an incentive to confuse the consumer by introducing greater variance. In this way, an expensive firm may increase the probability that the consumer will choose it over a cheaper firm. The surprising feature of Corollary 1 is that firms respond to greater competition by cultivating the "obfuscatory effect" only. Indeed, firms strictly prefer to obfuscate in equilibrium. For $n = 2$, it is easy to verify that although the support of $F(p, 2)$ is $[0, 1]$, the only degenerate *cdf* that yields the same expected profit for an individual firm is the one that assigns all weight to $p = \frac{1}{2}$. For $n > 3$, the firm strictly prefers $F(p, n)$ to *all* degenerate *cdf*s.

The reasoning behind part *(i)* of Corollary 1 consists of two parts. First, I argue that in order to calculate the expected price of $F(p, n)$, we can restrict attention to what was referred to earlier as "*quasi-bait-and-switch*" strategies. Moreover, $F(p, n)$ and the optimal quasi-bait-and-switch strategy share the same expected price. Second, I show that the optimal quasi-bait-and-switch strategy has an expected price of $\frac{1}{2}$, independently of the number of competitors.

The first part rests on a result that if F_i is a best-reply, then H_i must be *linear* on its support T_i. (This result has a precedent in Myerson (1993) - see Section 6.) Suppose that H_i were strictly convex (concave). Then, the firm could shift weight in a mean-preserving fashion from intermediate (extreme) price levels to extreme (intermediate) levels, and thus raise its expected revenue. The linearity of H_i has a useful corollary. Recall that firm i typically prefers its best-replying strategy to any distribution whose support consists of a single element. Nevertheless, an alternative *indifference principle* does hold in this model. If F_i is a best-reply, then firm i is indifferent between F_i and a quasi-bait-and-switch strategy - i.e., a lottery whose support consists of *two* elements, $p_* = \inf(T_i)$ and $p^* = \sup(T_i)$ - having the same expected price as F_i.

Let α denote the probability that the quasi-bait-and-switch strategy assigns to p_*. The firm's payoff is $[\alpha p_* + (1-\alpha)p^*] \cdot [\alpha H_i(p_*) + (1-\alpha)H_i(p^*)]$. A few straightforward steps in the proof of Proposition 1 establish that $p^* = 1$ and $H_i(1) = 0$. By an elementary calculation, the expected price of the optimal quasi-bait-and-switch strategy is $\frac{1}{2}$, independently of p_* and $H_i(p_*)$, hence of n. The "indifference principle" then implies that this must be a property of $F(p, n)$, too.

For a rough intuition behind part *(ii)* of Corollary 1, let us revisit the retail banking example of the Introduction. The pricing of a particular financial service has two independent functions: attracting clients and generating revenues. The service generates revenues from clients who chose the bank because it offers good terms for another service they happened to sample. As the number of competing banks increases, it becomes harder to generate a clientele from intermediately priced services. Therefore, the bank increasingly resorts to a strategy that relies on low-price services to attract clients and on high-price services to generate revenues.

Indeed, when $n > 2$, $F(p, n)$ assigns positive probability (increasing with n) to negative prices. The interpretation is that firms price some dimensions below marginal cost. These dimensions thus function as "loss leaders". Proposition 1 thus provides an explanation for this marketing tool as an obfuscation device. As competition becomes more intense, loss-leader pricing is more ubiquitous, reflecting a greater effort to obfuscate.[5]

Asymmetric equilibria

The next result shows that *if* asymmetric equilibria exist, they are *less competitive* (in terms of expected prices) than the symmetric equilibrium. Thus, the symmetric equilibrium has a special status in the model, because it is the *most competitive equilibrium.*

Proposition 2 *In Nash equilibrium,* $Ep_i \geq \frac{1}{2}$ *for every firm* i*, and* $Ep_i = \frac{1}{2}$ *for at least* $n - 1$ *firms.*

The reasoning behind this result is as follows. The "indifference principle" holds in asymmetric equilibrium. However, it turns out that we cannot rule out the possibility that $H_i(1) > 0$ for exactly firm i (because it is possible that i's competitors place an atom on $p = 1$). The implication is that among the

[5]Lal and Matutes (1994) provide an alternative account of loss-leader pricing, which focuses on the role of advertising. They assume that consumers can discover prices of unadvertised items only at the store, where they face a "hold-up". Consumers have rational expectations: they anticipate the hold-up problem and therefore reduce their willingness to shop. Firms use advertising as a commitment device that partially resolves the hold-up problem. As a result, firms compete fiercely over advertised items, while selling unadvertised items at the monopoly price. Lal and Matutes do not study how the number of competitors affects loss-leader pricing.

simple distributions with support $\{\inf(T_i), \sup(T_i)\}$, the optimal distribution has an expected price above $\frac{1}{2}$. By the "indifference principle", this is also a property of F_i.

The question whether asymmetric equilibria exist remains open for $n > 2$. However, for $n = 2$, the answer is negative.

Proposition 3 *When $n = 2$, there exist no asymmetric equilibria.*

The proof involves imitation arguments: each firm i can guarantee a payoff of $\frac{1}{2}Ep_j$, by imitating j's strategy. When $n > 2$, such imitation arguments are unavailable, and therefore I am unable to rule out asymmetric equilibria.

The role of the bounds on prices

Our model assumes that $p \in (-\infty, 1]$. The absence of a lower bound on prices guarantees the continuity of the equilibrium strategy. As Appendix II illustrates, if a lower bound is introduced, the equilibrium *cdf* places an atom on the lower bound for every sufficiently high n. However, as long as the lower bound does not exceed $\frac{1}{2}$, Corollary 1 continues to hold. Thus, the important aspects of Proposition 1 are insensitive to the lower bound.

If we assume that the upper bound on prices is any $s > 0$, then Proposition 1 continues to hold, except that p should be substituted everywhere with p/s. An upper bound is necessary for equilibrium existence. Otherwise, firms could charge a low price with positive probability in order to attract a clientele, while at the same time charging an arbitrarily high price with positive probability, thereby generating unbounded profits.

A comment on welfare

Any welfare analysis in a model with boundedly rational consumers must be conducted with caution, because we cannot infer their preferences from their behavior. Our $S(1)$ consumer chooses probabilistically, hence he does not behave as if he maximizes utility. Therefore, if we wish to analyze the welfare implications of market competition, we need to *assume* his "true" preferences over *cdf*s, because these are not revealed by his choices.

The most natural welfare criterion is that our consumer is risk-neutral, hence he only cares about the expected price that he ends up paying. The reason he does not choose the firm that offers the lowest expected price is that due to his bounded rationality, he is unable to identify that firm. If a "benevolent regulator" intervened and calculated the expected price of each alternative for him, the consumer would agree that the best alternative is the one that minimized expected price. Under this criterion, the conclusion from Proposition 1 is that an increased number of competitors does not affect consumer welfare.

This criterion does not take into account the consumer's risk attitudes. If his "true" preferences over *cdf*s display risk aversion, then the conclusion from Proposition 1 would appear to be that consumer welfare diminishes when n becomes larger, because $F(p, n+1)$ is a mean-preserving spread of $F(p, n)$. However, our $S(1)$ consumer is unaware of the actual structure of the *cdf*s he is facing. At the time he makes a decision, there is no reason to assume that the risk he perceives has anything to do with the actual risk. Therefore, welfare judgments that involve risk attitudes are problematic in this model.

4 Competition and inefficiency

As we saw in Section 3, when $n > 2$, firms assign positive probability to negative prices in equilibrium. However, suppose that firms were unable to use negative prices as a competitive device. An alternative way of attracting the consumer would be to invest in the *quality* of their product. Raising the product's quality level is a competitive strategy that is less vulnerable to the criticisms voiced in Section 2 against negative prices. The upshot is that investing in higher quality may introduce inefficiencies. The question is how the firms' pricing and quality decisions will respond to competition. In this section I analyze a model in which firms control both the price and the quality of their product.

Formally, assume that firms simultaneously choose probability distributions over quality-price pairs (q, p) satisfying $q \in [0, \infty)$ and $p \in [0, v(q)]$, where $v(q)$ is the consumer's willingness to pay for quality level q. If the consumer chooses a firm that offers the realization (q, p), his payoff is $v(q) - p$ and the firm's payoff is $p - q$. Assume that v is increasing and strictly concave, and that $v(q) - q$ attains a unique maximum at some $q^* > 0$. Denote the maximal surplus $v(q^*) - q^*$ by s.

It is easy to show that the Pareto frontier consists of all quality-price pairs (q, p) for which either *(i)* $q = q^*$ and $p \in [0, v(q^*)]$, or *(ii)* $q > q^*$ and $p = 0$. Thus, as long as we wish to sustain a consumer payoff below $v(q^*)$, it is efficient to produce the surplus-maximizing quality level q^*, and use prices to induce the desired payoff; whereas if we wish to sustain a consumer payoff $\bar{v} > v(q^*)$, it is efficient to produce q such that $v(q) = \bar{v}$ and set the price to zero. This characterization of the Pareto frontier is a consequence of *non-negativity* of prices. In this section, the firm's sole instrument for raising consumer payoff beyond $v(q^*)$ is to produce $q > q^*$. This feature will play an important role in the results of this section. If negative prices were allowed, we would have $q = q^*$ throughout the Pareto frontier.

Consumers choose according to the $S(1)$ procedure: they draw one sample point from each firm, and select the firm that maximizes $v(q) - p$ in their sample (with a symmetric tie-breaking rule). The outcome of the consumer's decision is a new, independent draw from his chosen distribution (or a sequence of such draws). The interpretations advanced for the basic model are appropriate here, too.

At first glance, this model looks like a considerable complication of the basic model, because now a strategy is a probability distribution over *pairs*. A simplification is immediately made possible thanks to the following observation. Because the firm is the residual claimant of any surplus that it produces, it has no reason to offer a quality-price pair outside the Pareto frontier. (The proof is elementary and therefore omitted.)

Remark 1 *In Nash equilibrium, firms assign probability zero to quality-price pairs that do not belong to the Pareto frontier.*

Thus, a strategy for a firm can be represented as a probability distribution over the unidimensional variable $\pi = \frac{1}{s}(p - q)$, which represents the firm's normalized profit conditional on being chosen by the consumer. For every π, the Pareto frontier assigns a unique consumer payoff $w(\pi)$, such that $w(\pi) = v(-s\pi)$ when $\pi \leq -\frac{q^*}{s}$, and $w(\pi) = s - s\pi$ when $-\frac{q^*}{s} < \pi \leq 1$. Observe that $w(\cdot)$ is strictly decreasing.

We have thus reduced the extended model into our basic formalism. Define firm i's strategy as a *cdf* F_i over π. Similarly, define $H_i(\pi)$ as the probability that in a random sample $(\pi_j)_{j \neq i}$ drawn from $(F_1, ..., F_n)$, $\pi_j > \pi_i$ for every $j \neq i$ (plus a term contributed by the breaking of ties). Firm i's payoff is $u_i(F_1, ..., F_n) = E\pi_i \cdot EH_i$, where $E\pi_i$ is the firm's expected profit conditional on being chosen, and EH_i is the probability that it is chosen. Both expectations are taken w.r.t F_i.

This simplification implies that the characterization of symmetric equilibrium is exactly the same as in Section 3. However, it has different welfare implications.

Proposition 4 *There is a unique symmetric Nash equilibrium in the game. Each firm plays the cdf given by expression (1).*

Corollary 2 *The symmetric equilibrium satisfies the following properties:*

(i) $E\pi = \frac{1}{2}$ *for every* $n \geq 2$.

(ii) For every $n > \frac{2v(q^*)}{s}$, $q > q^*$ *with positive probability.*

(iii) Expected surplus is strictly decreasing with n*, for* $n > \frac{2v(q^*)}{s}$.

The number of competitors does not affect the industry's expected equilibrium profits, which are equal to half the maximal surplus, regardless of n. However, when the number of firms is sufficiently high, the result is *inefficient* in terms of expected social surplus. Moreover, the loss in expected surplus is increasing with n, and borne entirely by the consumer. In this regard, greater competition has a negative effect on both consumer and social welfare.

The adverse welfare effects of competition are a consequence of the mean-preserving spread property of $F(\pi, n)$. Because there are diminishing returns to quality, the surplus function $v(q) - q$ is strictly concave. Recall that every $\pi \in [-\frac{q^*}{s}, 1]$ corresponds to the surplus-maximizing quality level q^* and every $\pi < -\frac{q^*}{s}$ corresponds to an "excessively high" quality level $q > q^*$. Thus, if we rewrite the surplus as a function of π, it is strictly increasing and strictly concave in the range $(-\infty, -\frac{q^*}{s}]$ and constant in the range $(-\frac{q^*}{s}, 1]$. Since $F(\pi, n+1)$ is a mean preserving spread of $F(\pi, n)$, expected surplus weakly decreases with n. Furthermore, because $F(\cdot)$ assigns positive probability to $q > q^*$ for $n > \frac{2v(q^*)}{s}$, expected surplus strictly decreases with n in this range.

Thus, the forces that drive the inefficiency result are: *(i)* the mean-preserving-spread property of $F(\pi, n)$, which reflects the firms' increased effort to obfuscate in response to greater competition, and *(ii)* diminishing returns for quality. Non-negativity of prices is crucial for this result. If firms could offer arbitrarily negative prices, the Pareto frontier would be linear, and expected surplus in equilibrium would be s, regardless of n.

The mean-preserving spread property has another implication. As n becomes larger, more weight is concentrated in two ranges of quality-price pairs: *(1)* the neighborhood of $(q^*, v(q^*))$, and *(2)* the set $\{(q, p) \mid q > q^*, p = 0\}$. This means that as the number of competitors increases, quality and price are more *negatively* correlated. Some dimensions are characterized by surplus-maximizing quality and the monopoly price, while other dimensions are characterized by excessive quality and zero price.

5 Outside options and "market exploitation"

When consumers make judgment errors, they are vulnerable to exploitation by rational firms. Therefore, it makes sense to ask what happens when they have an outside option that enables them to escape the market. Let us then extend the model of Section 2, such that the consumer's choice set is $\{0, 1, ..., n\}$, $n \geq 2$, where 0 denotes the outside option. Let F_0 be an exogenous *cdf* associated with the outside option. The support of F_0 is $T_0 \subseteq (-\infty, 1)$. In particular, F_0 does not assign an atom to $p = 1$. Extend the $S(1)$ procedure to encompass the outside option: consumers draw one sample point from each *cdf*. Faced with a sample $(p_0, p_1, ..., p_n)$, the consumer chooses $i \in \arg\min_i p_i$ (with a symmetric tie-breaking rule).

Proposition 5 *In Nash equilibrium, $Ep_i = \frac{1}{2}$ for every firm $i = 1, ..., n$.*

Thus, the presence of an outside option does not cause firms to act more competitively in terms of expected price. Recall that the restriction to symmetric equilibrium in Section 3 implied that $H_i(1) = 0$ for every firm i. The assumption that F_0 does not assign an atom to $p = 1$ has the same implication in the present model. By the "indifference principle", the expected price induced by firm i's best-reply is $\frac{1}{2}$.

One may argue that there is a difficulty in extending the $S(1)$ procedure to encompass the outside option. Consumers are naturally much more familiar with an outside option, whereas the $S(1)$ procedure reflects lack of familiarity with all alternatives. There is a simple way to resolve this difficulty without changing the model, by assuming that F_0 is a *degenerate cdf* that assigns probability one to some p_0. The interpretation is not that the outside option is genuinely deterministic, but that the consumer *knows* the expected price associated with the outside option.

Note that as long as $p_0 < 1$, Proposition 5 holds in this case. Thus, even if the consumer knows that $p_0 < \frac{1}{2}$, all firms choose $Ep = \frac{1}{2}$ in equilibrium. When the consumer chooses a firm over the outside option, he makes a decision error and loses $\frac{1}{2} - p_0$ in expectation. Since $EH_i > 0$ for every firm i in equilibrium, consumers experience "market exploitation" in equilibrium: they are worse off than if they were barred from entering the market.

If we rule out non-negative prices, this effect becomes extreme. Assume that firms are restricted to *cdf*s whose support is a subset of $[0, 1]$.

Proposition 6 *Let $n > 2$. Suppose that F_0 assigns probability one to some $p_0 \in [0, \frac{1}{2})$. Then, in symmetric Nash equilibrium, all firms play a distribution that assigns probability $\frac{1}{2}$ to each of the extreme prices, $p = 0$ and $p = 1$.*[6]

Thus, introducing a familiar, low-price outside option causes firms to raise the variance of their *cdf* to the maximal possible level. The intuition for this result is simple. When the consumer knows p_0, firms do not compete at all in the range $p > p_0$. They prefer to shift weight in this range to the extreme point $p = 1$, and this generates a large revenue for the firm, conditional on being chosen. Having secured this revenue, the firms can afford to compete fiercely in the low-price range (below p_0) in order to attract a clientele. When p_0 is sufficiently low, this causes the firms to place all remaining probability on the other extreme point $p = 0$.

According to Proposition 6, the probability that the consumer ends up choosing a firm in symmetric equilibrium is $1 - (\frac{1}{2})^n$ (as long as $p_0 > 0$). This expression increases with n, and converges to one as $n \to \infty$. Thus, if there are many firms in the market, the consumer experiences an almost certain welfare loss of $\frac{1}{2} - p_0$, relative to a world in which only the outside option is available.

The consumer is exploited in this case because he chooses a firm over the outside option whenever $p_i < p_0$ for some firm i. In other words, he behaves as if he believes that for every firm i, F_i assigns probability one to p_i. This is a "belief in the law of small numbers" writ large. The consumer

[6] When $n = 2$, the same result holds if $p_0 < \frac{3}{8}$.

treats a single sample point drawn from a firm's *cdf* as if it has the same informational content as full knowledge of p_0. Moreover, he disregards the fact that a firm's *cdf* is the result of a strategic choice, while the outside option is exogenous. For failing to draw these distinctions, the consumer suffers a welfare loss.

I find it interesting that introducing an attractive outside option has the same effect as raising n in the model without an outside option. Both interventions typically constitute effective competition policies in standard I.O. models. They continue to have a similar effect in the present model, albeit in an orthogonal direction.

6 Discussion

The premise of this paper is that consumers and firms often differ in their ability to perceive *cdf*s (or objects that can be described as *cdf*s). This asymmetry implies that firms might be able to take advantage of consumers, by raising the variance of the *cdf*s that they offer. It turns out that market competition does not protect consumers from this form of exploitation. Indeed, the firms' sole reaction to increased competition is to strengthen their obfuscation tactics. As a result, increased competition need not enhance consumer or social welfare, and may even cause them to deteriorate.

Extensions of the $S(1)$ procedure

Osborne and Rubinstein (1998) proposed a natural generalization of the $S(1)$ procedure, called $S(K)$, which in our context means that the consumer draws K independent sample points from every F_i, and chooses the alternative with the lowest *average* price in his sample. The parameter K reflects the extent to which the consumer's perception of alternatives falls short of full understanding.

There is some formal relation between the $S(K)$ procedure and the model of "inferences by believers in the law of small numbers" due to Rabin (2002). In this model, an individual decision maker observes repeated draws from an $i.i.d$ process, and tries to learn the process. He updates his belief according to Bayes' rule, as if the draws were taken from an urn with K balls without replacement. After K observations, the decision maker believes that the urn is refilled. Thus, Rabin's decision maker predicts the $(K+1)$-th observation just like an $S(K)$-agent. In other respects the two models are incomparable.

It is straightforward to modify the model of Section 2 by replacing the consumer's $S(1)$ procedure with the more general $S(K)$ procedure. The firms' payoff function continues to be $Ep_i \cdot EH_i$, with a different definition of EH_i. Define F_i^K as the *cdf* of the *average* of K independent draws from F_i. Let $H_i(p)$ be the probability that the consumer will choose firm i, conditional on the event that the realization of F_i^K is p. The expectation of H_i should be taken with respect to F_i^K.

In principle, one could adapt the equilibrium characterization technique of Section 3 to the generalized model, by analyzing the model as if the firms choose F_i^K, rather than F_i. Indeed, some arguments can be replicated using this trick. However, the "indifference principle" that plays a central role in Section 3 - namely, the payoff-equivalence between the firm's best-replying *cdf* and a two-outcome distribution - cannot be reproduced. The reason is that when $K > 1$, the support of F_i^K can never consist of exactly two outcomes.

In some cases, this model does lend itself to simple analysis. For instance, suppose that prices are forced to lie in the range $[0, 1]$, and that there is an outside option that offers the good at $p = 0$. Then, in Nash equilibrium, all firms play a probability distribution that assigns probability $\frac{K}{K+1}$ to $p = 0$ and $\frac{1}{K+1}$ to $p = 1$. When $n \to \infty$, industry profits converge to $\frac{1}{K+1}$.

Another way to enrich the $S(1)$ procedure is to endogenize the number of sample points that consumers draw from every F_i. According to the "learning" interpretation of the model proposed at the end of Section 2, the consumer's sample is a collection of "anecdotes", gathered during a brief phase of "word of mouth" learning. Given this interpretation, one could argue that the consumer is likely to hear more anecdotes about firms with a larger clientele. I conjecture that an extension of the model which addresses this concern will introduce a subtle anti-competitive effect. If consumers have a more accurate perception of a firm with a larger clientele, then it is harder for such a firm to take advantage of consumers. This creates an incentive for firms to reduce their clientele by offering less attractive *cdf*s.

Can the model be rationalized?

The modeling procedure in this paper is non-standard, in that it sets up a market model in which the two sides differ in their ability to grasp the firms' strategies. The question arises, whether one could "rationalize" the model in some sense. This question has two distinct meanings. First, we may ask whether the consumer's *individual* behavior, although non-rational in our market model, might be rational in some other market environment. The answer is affirmative. If the consumer believes that each firm offers the same terms in all contingencies, and that these terms are drawn from some common distribution, then he might as well sample one contingency only. A "rationalization" of a similar sort is to assume that the consumer is interested only in the price of the contingency that he samples, because he gets "reimbursed" by some third party in every other contingency. In the former case, we rationalize the consumer's behavior by assuming that he believes in an incorrect model. In the latter case, we rationalize it by assuming that it is the analyst who gets the model wrong.

A more interesting question is whether market equilibria in our model can be replicated as sequential equilibria in another market model, in which our imperfectly rational consumer is substituted with an imperfectly informed, expected-utility maximizing consumer. To explore this question, consider the following variant on the model of Section 2. The Consumer moves after the firms choose their strategies, but he is unable to observe their choice. However, he can condition his action on a random draw from $(F_1, ..., F_n)$. What is the relation between sequential equilibria in this game and Nash equilibria in our model?

There is a sequential equilibrium in this incomplete-information game, in which all firms play the *cdf* $F(p, n)$ given by Proposition 1, and the consumer plays the strategy induced by the $S(1)$ procedure - i.e., choosing the firm with the highest realization in his sample - both on and off the equilibrium path. In equilibrium, the consumer is indifferent among all firms, hence he does not mind playing this strategy. However, there are many *other* sequential equilibria which sustain different market outcomes. For example, we can sustain the fully monopolistic outcome (in which all firms charge $p = 1$ with probability one), using suitable out-of-equilibrium beliefs. In this sense, it is impossible to say that the incomplete-information game rationalizes the basic model. The same holds for the model of Section 4.

When we add an outside option F_0 with $Ep_0 < \frac{1}{2}$, the verdict is even harsher. Suppose that the consumer is risk-neutral (a similar argument can be devised for any other risk attitude). Then, there exists *no* sequential equilibrium that rationalizes the results of Section 5. The reason is that in sequential equilibrium, it is impossible for all firms to play a *cdf* with $Ep = \frac{1}{2}$ and for the consumer to choose a firm over the outside option. The failure to rationalize the model in this case follows from the rational-expectations aspect of sequential equilibrium: the consumer can never be systematically fooled. He will not choose an alternative which is more expensive on average.

Relation to the redistributive politics literature
At the purely formal level, there is a surprising link between the basic model and the redistributive politics model due to Myerson (1993). In Myerson's model, n political candidates play a simultaneous-move game. Each candidate chooses an income redistribution policy subject to a balanced-budget constraint. The policy can be represented as a *cdf* with an exogenously given expected value. Each voter learns his net income under each candidate's policy and votes for the candidate that promises him the highest net income. The candidates' sole motive is to be elected.

Myerson examines several electoral systems. Under proportional representation - a system Myerson does *not* analyze in his paper - candidate i's payoff function is EH_i as defined in this paper. The statement that firms in the present model prefer higher-variance price distributions is analogous to the statement that candidates in Myerson's model prefer more unequal income distributions (interpreted as cultivation of favored minorities). If we assume that the budget in Myerson's model is $\frac{1}{2}$, the symmetric equilibrium strategy is as given by Proposition 1.

Of course, the main formal difference between the above version of Myerson's model and the present model is that the latter imposes no constraint on the expected value of *cdf*s. The formal linkage emerges precisely because Ep_i turns out to be independent of n in equilibrium. The extended models of Sections 4 and 5 have an analogue neither in Myerson (1993) nor in the literature that followed it.

Myerson's model was further developed by Lizzeri (1999), who added an intertemporal element in order to incorporate budget deficits. Lizzeri and Persico (2001,2002) extended the model by assuming that candidates can choose between redistribution and investment in a public good. Although the latter is more efficient, in equilibrium candidates tend to prefer the former. This type of inefficiency is unrelated to the inefficiency result of Section 4. Rather, it is analogous to the "obfuscatory effect" of the basic model, according to which a firm may prefer a noisy *cdf* with expected price β to a degenerate *cdf* that assigns probability one to a lower price $\beta' < \beta$.

References

[1] Burdett K. and M. Coles (1997), "Steady State Price Distributions in a Noisy Search Equilibrium", *Journal of Economic Theory* **72**, 1-32.

[2] Burdett K. and K. Judd (1983), "Equilibrium Price Dispersion", *Econometrica* **51**, 955-970.

[3] Camerer C. and G. Lowenstein (2003), "Behavioral Economics: Past, Present, Future", in Advances in Behavioral Economics, C. Camerer, G. Lowenstein and M. Rabin, eds., Princeton University Press.

[4] Chen Y., G. Iyer and A. Pazgal (2003), "Limited Memory and Market Competition", mimeo.

[5] DellaVigna, S. and U. Malmendier (2004), "Contract Design and Self-Control: Theory and Evidence", *Quarterly Journal of Economics* **119**, 353-402.

[6] Eliaz K. and R. Spiegler (2004), "Contracting with Diversely Naive Agents", mimeo.

[7] Ellison G. and S. Fisher Ellison (2004), "Search, Obfuscation, and Price Elasticities on the Internet", mimeo.

[8] Erev I. and E. Haruvy (2001), "Variable Pricing: A customer Learning Perspective", mimeo.

[9] Gabaix X. and D. Laibson (2005), "Consumer Myopia, Shrouded Attributes and Information Suppression in Competitive Markets", mimeo.

[10] Kahneman D., Slovic P. and A. Tversky (1982), *Judgment under uncertainty*, Cambridge University Press.

[11] Lal, R. and C. Matutes (1994), "Retail Pricing and Advertising Strategies", *Journal of Business* **67**, 345-370.

[12] Lizzeri A. (1999), "Budget Deficits and Redistributive Politics", *Review of Economic Studies* **66**, 909-928.

[13] Lizzeri A. and N. Persico (2001), "The Provision of Public Goods Under Alternative Electoral Incentives", *American Economic Review* **91**, 225-39.

[14] Lizzeri A. and N. Persico (2002), "The Drawbacks of Electoral Competition", *Journal of the European Economic Association*, forthcoming.

[15] Mas-Colell A., M. Whinston and J. Green (1995), *Microeconomic Theory*, Oxford University Press.

[16] McAfee P. (1995), "Multi-product Equilibrium Price Dispersion", *Journal of Economic Theory* **67**, 83-105.

[17] Myerson R. (1993). "Incentives to Cultivate Favored Minorities Under Alternative Electoral Systems", *American Political Science Review* **87**, 856-869.

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

[19] Piccione M. and A. Rubinstein (2003), "Modeling the Economic Interaction of Consumers with Diverse Abilities to Recognize Equilibrium Patterns", *Journal of European Economic Association* **1**, 212-223.

[20] Rabin M. (2002), "Inference by Believers in the Law of Small Numbers", *Quarterly Journal of Economics* **117**, 775-816.

[21] Rob R. (1985), "Equilibrium Price Distributions", *The Review of Economic Studies* **52**, 487-504.

[22] Rubinstein A. (1993), "On Price Recognition and Computational Complexity in a Monopolistic Model", *Journal of Political Economy* **101**, 473-484.

[23] Salop S. (1977), "The Noisy Monopolist: Imperfect Information, Price Dispersion and Price Discrimination", *Review of Economic Studies* **44**, 393-406.

[24] Spiegler R. (2003), "The Market for Quacks", mimeo.

[25] Varian H. (1980), "A Model of Sales", *The American Economic Review* **70**, 651-659.

[26] Wilson C. (1988), "On the Optimal Pricing Policy of a Monopolist", *Journal of Political Economy* **96**, 164-176.

Appendix I: Proofs

The following notation will be employed throughout the appendices. Let $x = 1-p$. If the consumer's gross valuation of the firms' product is 1, then x represents his net payoff if he pays the price p. Redefine firm i's strategy as a *cdf* G_i over $[0,\infty)$, such that $G_i(x)$ is the probability that the consumer's net payoff is at most x. Redefine T_i as the support of G_i.

Given the profile $(G_j)_{j\neq i}$, redefine $H_i(x)$ as the probability that the consumer chooses firm i, conditional on $x_i = p$ in his sample. Let EH_i denote the expected value of H_i, where the expectation is taken with respect to G_i. When each firm plays a *cdf* with a well-defined density g, the definition of H_i and EH_i has a simple form: $H_i(x) = \Pi_{j\neq i}G_j(x)$ and $EH_i = \int_0^\infty H_i(x)g_i(x)dx$.

I begin with a lemma which holds for all versions of the model. A similar result appears in Myerson (1993).

Lemma 1 (Linearity of H) *Suppose that G_i is a best-reply to G_{-i}. Then, all the points $\{(x, H_i(x)) \mid x \in T_i\}$ lie on a straight line, except possibly for a subset which is assigned probability zero by G_i. Moreover, for every $x \in [\inf(T_i), \sup(T_i)]\backslash T_i$, $(x, H_i(x))$ cannot lie above the line.*

Proof. If T_i consists of two outcomes, the first part of the result holds trivially. Let $x_1, x_2, x_3 \in T_i$, $x_1 < x_2 < x_3$. Assume first that G_i places an atom on each of the three points. Let α_k denote the mass placed on x_k.

Suppose that $(x_2, H(x_2))$ does not lie on the line connecting $(x_1, H(x_1))$ and $(x_3, H(x_3))$. Then, firm i can deviate to a strategy G_i', which differs from G_i only in the probabilities it assigns to x_1, x_2, x_3, such that $\alpha_k' = \alpha_k + \varepsilon_k$ and $\varepsilon_1 x_1 + \varepsilon_2 x_2 + \varepsilon_3 x_3 = 0$. If $(x_2, H(x_2))$ lies *below* the line connecting $(x_1, H(x_1))$ and $(x_3, H(x_3))$, we set $\varepsilon_1, \varepsilon_3 > 0$ and $\varepsilon_2 < 0$. If $(x_2, H(x_2))$ lies *above* the line connecting $(x_1, H(x_1))$ and $(x_3, H(x_3))$, we set $\varepsilon_1, \varepsilon_3 < 0$ and $\varepsilon_2 > 0$. In both cases, Ex_i is the same according to G_i and G_i', but EH_i is higher under G_i'. Therefore, $(x_2, H(x_2))$ must lie on the line connecting $(x_1, H(x_1))$ and $(x_3, H(x_3))$. Note that the argument we applied when $(x_2, H(x_2))$ lies *above* the line connecting $(x_1, H(x_1))$ and $(x_3, H(x_3))$ does not rely on the assumption that $x_2 \in T_i$. The second part of the claim thus immediately follows.

Extending the proof to the case in which G_i does not place an atom on x_k is straightforward. Assume that G_i assigns positive probability to the neighborhood of x_k. Assume further that G_{-i} is continuous at x_k. By definition, G_i assigns zero probability to the set of points for which this is not the case. It follows that H_i is continuous at x_k. Then, we can reproduce the above deviations which involve shifting weight from x_k to other points, only now the weight is shifted from all points in an arbitrarily small neighborhood of x_k. ■

Corollary 3 (An indifference principle) *Let G_i and G_i' be a pair of cdfs that satisfy two conditions: (i) Ep_i is the same according to G and G'; (ii) $T_i' \subseteq T_i$. Then, G_i' and G_i generate the same expected payoff for firm i, given G_{-i}.*

Proof. By Lemma 1, H_i is linear on T_i (except possibly for a subset assigned zero probability by G_i). Therefore, any mean-preserving shift of weight within the support of T_i also preserves expected payoff. ■

The following pair of lemmas pertain to Propositions 1-4.

Lemma 2 *Suppose that there is no outside option. Then, in Nash equilibrium, G_i is continuous in the range $(0, \infty)$ for every firm i.*

Proof. Suppose that G_i assigns an atom to $x \in (0, \infty)$. Then, there exists $\varepsilon > 0$, such that no other firm j assigns any weight to $(x - \varepsilon, x]$ - because by shifting this weight to $x + \varepsilon$, firm j could increase EH_j by an amount that is bounded away from zero, while infinitesimally reducing Ex_j. But this means that firm i can profitably deviate by shifting the atom on x slightly downward - this will leave EH_i unaffected, while reducing Ex_j. This contradicts the assumption that $(G_i)_{i=1,\ldots,n}$ is an equilibrium. ■

Lemma 3 *Suppose that there is no outside option. Then, in Nash equilibrium, $\inf(T_i) = 0$ for every firm i.*

Proof. Assume that there exists a firm i, such that $\inf(T_i) > 0$. Denote $\inf(T_i) = x_i^*$. Then, for every $j \neq i$, $H_j(x) = 0$ for all $x < x_i^*$. It follows that if G_j assigns positive weight to $[0, x^*)$, all the weight is assigned to $x = 0$. Suppose that firm i deviates by shifting all the weight it assigns to $(x_i^*, x_i^* + \varepsilon)$ to some arbitrarily small $x > 0$. Then, it reduces Ex_i by $x_i^* \cdot G_i(x_i^* + \varepsilon)$. At the same time, it reduces EH_i by $\Pi_{j=1,\ldots,n}[G_j(x_i^* + \varepsilon) - G_j(x_i^*)]$. If $\varepsilon > 0$ is sufficiently small, the reduction in EH_i is negligible in comparison to the reduction in Ex_i, such that the deviation is profitable. ■

Proof of Proposition 1
Let G denote the symmetric equilibrium strategy. Let T denote the support of G and denote $y = \sup(T)$.

Step 1. $T = [0, y]$.
Proof. We have already shown that $\inf(T) = 0$. Suppose that G assigns zero probability to some interval $(x_1, x_2) \subset (0, y)$. Then, any firm i can deviate by shifting to x_1 all the weight it assigns to $(x_2, x_2 + \varepsilon)$. If ε is sufficiently small, the reduction in EH_i is negligible in comparison to the decrease in Ex_i, hence the deviation is profitable. □

Step 2. G is continuous over $[0, y]$.
Proof. We have already shown that G is continuous in $(0, y]$. If G assigns an atom to $x = 0$, then any firm i can deviate by shifting the weight from $x = 0$ to $x' > 0$. If x' is sufficiently small, the increase in Ex_i is negligible in comparison to the increase in EH_i, such that the deviation is profitable. □

Step 3. $Ex = \frac{1}{2}$.
Proof. By Step 2, $H(0) = 0$. By Lemma 1, $H(x) = \frac{1}{y} \cdot x$ for every $x \in [0, y]$. Let us turn to calculating y. Suppose that any firm i switches from G to a simple lottery G^* such that: *(i)* $T^* = \{0, y\}$; *(ii)* Ex is the same under G and G^*. By Corollary 3, the firm is indifferent between G and G^*. Let α denote the probability that G^* assigns to y. Then, the firm's payoff from G^* is $[1 - \alpha y - (1 - \alpha) \cdot 0] \cdot [\alpha H(y) + (1 - \alpha) H(0)]$. But since $H(0) = 0$, this expression can be simplified into $(1 - \alpha y) \cdot \alpha \cdot H(y)$.

If $y \leq \frac{1}{2}$, then the value of α that maximizes the simplified expression is $\alpha = 1$. But this means that the expected value of x according to G^* is y, whereas the expected value of x according to G is $\mu < y$. By Corollary 3, firm i is indifferent between G and a lottery G' with support $\{0, y\}$ and $Ex = \mu$. Because G' generates a lower expected payoff than G^*, it follows that G, too, is not a best reply, a contradiction. Therefore, $y > \frac{1}{2}$. In this case, the value of α that maximizes this expression is $\alpha = \frac{1}{2y}$, yielding $Ex = \frac{1}{2}$ according to G^*. But if the optimal lottery among the lotteries with support $\{0, y\}$ has an expected value $Ex = \frac{1}{2}$, then by Corollary 3, this must be the expected value according to G as well. □

The only remaining step is to derive the equilibrium strategy. In equilibrium, $EH_i = \frac{1}{n}$ for every firm i, by symmetry. Therefore, the expected payoff from G must be $\frac{1}{2n}$. By Corollary 3, this is also the firm's expected payoff from G^*. Since the expected payoff from G^* is $(1 - \alpha y) \cdot \alpha$, we obtain $y = \frac{n}{2}$. Therefore, $H(x) = \frac{2x}{n}$ for every $x \in [0, \frac{n}{2}]$. Because G is continuous over T, we obtain

$$G(x) = \sqrt[n-1]{\frac{2x}{n}}$$

When all firms play G, we have a Nash equilibrium. The reason is that since H is linear over $[0, \frac{n}{2}]$, every *cdf* over this support satisfying $Ex = \frac{1}{2}$ is a best-reply, according to Corollary 3.

Substituting back $p = 1 - x$ and $F(p) = 1 - G(x)$, we obtain the formula given by (1). ■

Proof of Corollary 1

Part *(i)* is immediate, because the construction of G relies on the result that $Ex = \frac{1}{2}$. As to part *(ii)*, a simple calculation shows that for every $x \in [0, \infty)$, $\int_0^x G(w)dw$ is increasing with n. By a well-known result (see Mas-Colell, Whinston and Green (1995), p. 198), this is equivalent to saying that $G(x, n+1)$ is a mean-preserving spread of $G(x, n)$. ■

Proof of Proposition 2

Lemmas 1-3 and Corollary 3 hold for any Nash equilibrium. Our objective is to show that there is at most one firm i for which $H_i(0) > 0$. Denote the number of firms that place an atom on $x = 0$ by k. If $k \leq n - 2$, then $H_i(0) = 0$ for every firm i. If $k = n$, then $H_i(0) > 0$ for every firm i. But this means that any firm can profitably deviate by shifting weight from $x = 0$ to an arbitrarily small $x' > 0$. Now suppose that $k = n - 1$ - i.e., there is exactly one firm j, which does not place an atom on $x = 0$. In this case, $H_j(0) > 0$ and $H_i(0) = 0$ for every $i \neq j$.

The rest of the proof follows Step 3 of the proof of Proposition 1. Denote $y = \sup(T_i)$, and recall that by lemma 3, $\inf(T_i) = 0$. Suppose that firm i switches from G_i to a simple lottery G_i^* which satisfies: *(i)* $T_i^* = \{0, y\}$; *(ii)* Ex is the same under G_i and G_i^*. By Corollary 3, the firm is indifferent between G_i and G_i^*. Let α denote the probability that G_i^* assigns to y. Then, the firm's payoff from G_i^* is $[1 - \alpha y - (1-\alpha) \cdot 0] \cdot [\alpha H_i(y) + (1-\alpha) H_i(0)]$. If $H_i(0) = 0$, then $Ex = \frac{1}{2}$, following the same arguments as in Step 3 of the proof of Proposition 1. If $H_i(0) > 0$, then the optimal value of α yields $Ex < \frac{1}{2}$, which is therefore also true under G_i. But this means that in Nash equilibrium, $Ex \leq \frac{1}{2}$ for all firms and $Ex < \frac{1}{2}$ for at most one firm. ■

Proof of Proposition 3

Let us begin by stating three properties of the supports of the firms' strategies, denoted T_1 and T_2:

- $T_1 = T_2 = T$. Assume the contrary, and suppose that firm i assigns positive probability to some interval (b, c), whereas firm j does not. Then, firm i can deviate by shifting all this weight nearer b.
- $\inf(T) = 0$. This is due to Lemma 3.
- T is connected. Assume that there is an interval $(x_1, x_2) \subset (0, \sup(T))$ which is assigned zero probability by G_1 and G_2. Then, any firm i can deviate by shifting all the weight it assigns to $(x_2, x_2 + \varepsilon)$ to x_1. If ε is sufficiently small, the reduction in EH_i is negligible in comparison to the decrease in Ex_i, hence the deviation is profitable.

The combination of these properties yields $T_1 = T_2 = T = [0, y]$.

By Lemma 2, both G_1 and G_2 are continuous in $(0, y]$. Moreover, at least one firm, say firm 2 w.l.o.g, does not place an atom on $x = 0$ (otherwise, it would be profitable for any firm to shift weight from $x = 0$ to some arbitrarily small $x' > 0$). Therefore, G_2 is continuous in $[0, y]$. By Lemma 1, both H_1 and H_2 are linear in T. By Corollary 3, each firm i is indifferent between G_i and the lottery G^* that assigns probability $\frac{1}{2y}$ to $x = y$ and probability $1 - \frac{1}{2y}$ to $x = 0$. Because G_2 does not assign an atom to $x = 0$, $G_2(x) = H_1(x) = \frac{1}{y} \cdot x$ for every $x \in [0, y]$. Thus, by Proposition 2, $Ex_1 = \frac{1}{2}$ and $Ex_2 \leq \frac{1}{2}$.

Let us show that $y \leq 1$. Because $H_1(0) = 0$, firm 1's expected payoff from G^* - and therefore, from G_1 as well - is $\frac{1}{2} \cdot \frac{1}{2y}$. It follows that $EH_1 = \frac{1}{2y}$. Firm 1's expected payoff must be at least $\frac{1}{4}$. Otherwise, the firm could deviate from G_1 to G_2 (i.e., imitate the opponent). By symmetry, both firms would have $EH = \frac{1}{2}$, and because $Ex_2 \leq \frac{1}{2}$, firm 1's payoff would be at least $\frac{1}{4}$, a contradiction. We have thus established that $\frac{1}{2} \cdot \frac{1}{2y} \geq \frac{1}{4}$, hence $y \leq 1$.

Let us now show that $y \geq 1$. Because $G_2(x) = \frac{1}{y} \cdot x$ for any $x \in [0, y]$, $Ex_2 = \frac{y}{2}$. We have shown in the preceding paragraph that $EH_1 = \frac{1}{2y}$. Since there are only two firms, $EH_2(x) = 1 - EH_1(x)$. Therefore, firm 2's payoff is $(1 - \frac{y}{2}) \cdot (1 - \frac{1}{2y})$. Firm 2's expected payoff from G^* - and therefore, from G_2 as well - is at least $\frac{1}{4y}$ (and strictly higher, if $H_2(0) > 0$, in case G_1 places an atom on $x = 0$). We have thus established that $(1 - \frac{y}{2}) \cdot (1 - \frac{1}{2y}) \geq \frac{1}{4y}$, hence $y \geq 1$.

It follows that $y = 1$, such that $Ex_2 = Ex_1 = \frac{1}{2}$. Suppose that G_1 places an atom on $x = 0$, such that $H_2(0) > 0$. By Proposition 2, $Ex_2 < \frac{1}{2}$, a contradiction. It follows that $G_1 \equiv G_2 \equiv U[0, 1]$. ■

Proof of Corollary 2

Part *(i)* is the same as in Corollary 1. Part *(ii)* is a simple consequence of the observation that the infimum of the support of $F(\pi, n)$ is $1 - \frac{n}{2}$, which is lower than $-\frac{q^*}{s}$ for $n > \frac{2v(q^*)}{s}$. Let us turn to part *(iii)*. By part *(ii)* of Corollary 1, $F(\pi, n+1)$ is a mean-preserving spread of $F(\pi, n)$. The surplus produced by the firm is by definition $w(\pi) + s\pi$. But this function is concave in π, and strictly concave in the range $\pi < -\frac{q^*}{s}$. Therefore, expected surplus decreases with n, and strictly so for $n > \frac{2v(q^*)}{s}$. ■

Proof of Proposition 5

Lemma 1 and Corollary 3 hold for any Nash equilibrium, whether or not there is an outside option. By assumption, $G_0(0) = 0$, hence $H_i(0) = 0$ for every firm i. We can now apply Step 3 in the proof of Proposition 1. ■

Proof of Proposition 6

Denote $x_0 = 1 - p_0$. Let G be the symmetric equilibrium strategy. Let T denote its support and denote $y = \sup(T)$. By definition, $H(x) = 0$ for every $x < x_0$. Therefore, G assigns zero probability to $(0, x_0]$, such that $G(x) = G(0)$ for every $x \leq x_0$. Also, because $H(0) = 0$, we can apply Step 3 in the proof of Proposition 1, hence that $Ex = \frac{1}{2}$ according to G, and firms are indifferent between G and the simple lottery that assigns probability $\frac{1}{2y}$ to $x = y$ and probability $1 - \frac{1}{2y}$ to $x = 0$. Therefore, a firm's equilibrium payoff is $\frac{1}{2} \cdot \frac{1}{2y} \cdot H(y)$. By symmetry, $EH = \frac{1}{n} \cdot [1 - G(0)]^n$. Therefore, the firm's equilibrium payoff is $\frac{1}{2} \cdot \frac{1}{n} \cdot [1 - G(0)]^n$. It follows that $H(y) = \frac{2y}{n} \cdot [1 - G(0)]^n$. Suppose that $y < 1$. By the same reasoning as in Lemma 2, G is continuous in $x \in (x_0, 1)$, hence $H(y) = 1$, but then we obtain $n = 2y \cdot [1 - G(0)]^n$, contradicting the assumption that $n > 2$. It follows that $y = 1$.

Assume that G assigns positive probability to the interval $(x_0, 1)$. Let us show that $G(x) > G(0)$ for every $x > x_0$. Assume the contrary, and let x^* denote the infimum of $T_i \cap (x_0, 1]$. Then, $H(x) = 0$ for every $x < x^*$. Suppose that firm i deviates by shifting all the probability it assigns to $(x^*, x^* + \varepsilon)$ to some $x > x_0$ arbitrarily close to x_0. Then, the firm manages to reduce Ex by $x^* \cdot [G(x^* + \varepsilon) - G(x^*)]$. At the same time, it reduces EH_i by $[G(x^* + \varepsilon) - G(x_i^*)]^n$. If ε is sufficiently small, the reduction in Ex_i more than compensates for the reduction in EH_i. Therefore, the deviation is profitable.

By Lemma 1, H is linear over T. Note that $0 \in T$ - otherwise, G would assign positive weight only to elements above x_0 (which itself exceeds $\frac{1}{2}$), contradicting our result that $Ex = \frac{1}{2}$. Because $H(0) = 0$ and because G assigns positive probability to elements above and arbitrarily close to x_0, $H(x)$ tends to $x_0 \cdot H(1)$ as x tends to x_0. At the same time, by definition, $H(x)$ tends to $[G(0)]^{n-1}$ as x approaches x_0 from above. Using our expressions for $\lim_{x \to x_0^+} H(x)$ and $H(1)$, we obtain:

$$G(0)^{n-1} = x_0 \cdot \frac{2}{n} \cdot [1 - G(0)]^n$$

Because $Ex = \frac{1}{2}$, it must be the case that $G(0) \leq \frac{1}{2}$. And since $x_0 \geq \frac{1}{2}$, we obtain $2n + 1 \geq 2^n$, a contradiction for $n > 2$.

The only remaining possibility is that G assigns probability $\frac{1}{2}$ to each of the extreme points $x = 0$ and $x = 1$. Let us verify that this is an equilibrium. Given that firm i's opponents all play G, $H_i(x) = 0$ for every $x \in (0, x_0)$, $H_i(x) = (\frac{1}{2})^{n-1}$ for every $x \in (x_0, 1)$ and $H_i(1) = \frac{2}{n} \cdot [1 - (\frac{1}{2})^n]$. For $n > 2$, $H_i(x) < x \cdot H(1)$ for every $x \in (0, 1)$. Therefore, firm i will never want to deviate to a strategy whose support is not $\{0, 1\}$. Among the lotteries whose support is $\{0, 1\}$, G is optimal. Therefore, G is a best-reply. ■

Appendix II: Non-negative prices

When the set of feasible prices is $[0, 1]$, thus ruling out negative prices, symmetric Nash equilibrium in the basic model is characterized as follows.

Proposition 7 *There is a unique symmetric Nash equilibrium in the game. Each firm plays the cdf:*

$$F(p,n) = \begin{cases} 1 - A_n & 0 \le p \le b_n \\ 1 - \sqrt[n-1]{\frac{2(1-p)}{n}} & b_n < p \le 1 \end{cases} \qquad (2)$$

where $b_n = 1 - \frac{n}{2}(A_n)^{n-1}$ *and* A_n *is the unique solution in* $[\frac{1}{2}, 1]$ *of the equation:*

$$(A_n)^n - 2A_n + 1 = 0 \qquad (3)$$

Corollary 1 holds for this result, too.

What is the shape of $F(p,n)$? For $n = 2$, $F(p,n) \equiv U[0,1]$. For every $n > 2$, the support of $F(p,n)$ is $T = \{1\} \cup [b_n, 1]$, where b_n increases with n and tends to one as $n \to \infty$; $F(p,n)$ contains an atom (whose mass is $1 - A_n$) on $p = 0$; the atom's size increases with n and tends to $\frac{1}{2}$ as $n \to \infty$; $F(p,n)$ contains no other atom. As $n \to \infty$, $F(p,n)$ approaches the distribution that assigns probability $\frac{1}{2}$ to each of the two extreme points $p = 0$ and $p = 1$. This limit distribution has greater variance than any other *cdf* over $[0,1]$. The convergence is fast: $A_6 \approx 0.51$; $b_6 \approx 0.9$.

Proof of Proposition 7

I continue to employ the notation of Appendix I: $x = 1 - p$ and $G(x) = 1 - F(1-p)$. Lemmas 1 and 3, as well as Corollary 3, hold in this model. Lemma 2 holds, except that G is continuous over $(0,1)$, rather than $(0,\infty)$. In fact, using the same reasoning as in Step 2 in the proof of Proposition 1, we can show that G is continuous over $[0,1)$. Therefore, Step 3 in the same proof holds here, too. It follows that in symmetric equilibrium, $Ex = \frac{1}{2}$, such that each firm earns a payoff of $\frac{1}{2n}$. It remains to characterize the support and exact formula of G.

Step 1: There exists a number $b \in (0,1]$, such that $T = [0, 1-b] \cup \{1\}$.

Proof: First, let us first show that $y = 1$. Suppose that $y < 1$. Because G is atomless on $[0,1)$, $G(y) = 1$ and $H(y) = 1$. Each firm can deviate to a lottery that assigns probability $\frac{1}{2y}$ to y and probability $1 - \frac{1}{2y}$ to 0. The firm's payoff would be $\frac{1}{2} \cdot \frac{1}{2y}$, which is larger than $\frac{1}{2n}$, a contradiction. It follows that $y = 1$. Suppose that $x \in T$ for some $x \in (0, 1)$. By Lemma 3, $0 \in T$. Using the same reasoning as in Step 1 in the proof of Proposition 1, we can show that every $x' \in (0, x)$ must also belong to T. $\square$

Step 2: $H(1) = \frac{2}{n}$.

Proof: By Step 3 in the proof of Proposition 1, each firm is indifferent between G and a lottery G' with support $\{0, y\}$ and $Ex = \frac{1}{2}$. By Step 1, $y = 1$. Therefore, each firm is indifferent between G and the lottery that assigns probability $\frac{1}{2}$ to each of the two extreme points, $x = 0$ and $x = 1$. The firm's payoff from this lottery is $\frac{1}{2} \cdot \frac{1}{2} H(1)$. Since this must be equal to $\frac{1}{2n}$, $H(1) = \frac{2}{n}$. $\square$

Step 2 implies that G places an atom on $x = 1$ for every $n > 2$. Otherwise, $H(1)$ would be equal to one and the firm's payoff would exceed $\frac{1}{2n}$, a contradiction. Let us denote the size of this atom by $1 - A$.

Step 3: G must be given by expressions (2)-(3).

Proof: By Step 1, $T = [0, 1 - b] \cup \{1\}$. By Lemma 1, H is linear over T. We have established that $H(0) = 0$ and $H(1) = \frac{2}{n}$. Therefore, $H(x) = 2x/n$ for every $x \in T$. Because G contains no atoms below $x = 1$, $H(x) = G^{n-1}(x)$ for every $x \in [0, 1 - b]$. Therefore, in this range:

$$G(x) = \sqrt[n-1]{\frac{2x}{n}}$$

It only remains to determine the exact values of b and A. By definition, $G(1 - b) = A$. The relation between b and A is thus given by:

$$b = 1 - \frac{n}{2} A^{n-1}$$

Let us now determine the value of A. Let g be the density function induced by G in the interval $[0, 1 - b]$. By definition:

$$Ex = (1 - G(1 - b)) \cdot 1 + \int_0^{1-b} x g(x) dx$$

Because $Ex = \frac{1}{2}$, we can retrieve $G(1 - b)$ from the expression for Ex, obtaining:

$$\frac{1}{2} = \frac{G^n(1 - b)}{2} - G(1 - b) + 1$$

which can be rewritten as:

$$A^n - 2A + 1 = 0$$

Substituting $p = 1 - x$ and $F(p) = 1 - G(x)$ for every $p > 0$, we have the desired characterization. □

Step 4: The strategy profile $(G, ..., G)$ is a Nash equilibrium.

Proof: First, let us derive H. Given the expression for G, it follows immediately that $H(x) = 2x/n$ for every $x \leq 1 - b$ and $H(x) = 2(1-b)/n$ for every $x \in (1-b, 1)$. Let us check that $H(1) = 2/n$. Denote $m = n - 1$. The precise definition of $H(1)$ in the symmetric equilibrium is:

$$H(1) = \sum_{k=0}^{m} \frac{\binom{m}{k}}{k+1} A^{m-k}(1-A)^k$$

We can rewrite:

$$\frac{\binom{m}{k}}{k+1} = \frac{m!}{(m-k)! \cdot k! \cdot (k+1)} \cdot \frac{m+1}{m+1} = \frac{1}{m+1} \cdot \binom{m+1}{k+1}$$

Denote $j = k + 1$. Then:

$$\sum_{k=0}^{m} \frac{\binom{m}{k}}{k+1} A^{m-k}(1-A)^k = \frac{1}{A(m+1)} \cdot \sum_{j=1}^{m+1} \binom{m+1}{j} A^{m+1-j}(1-A)^j$$

By a standard binomial expansion:

$$\binom{m+1}{0} \cdot A^{m+1} + \sum_{j=1}^{m+1} \binom{m+1}{j} A^{m+1-j}(1-A)^j = 1$$

Therefore, $A^{m+1} + A(m+1) \cdot H(1) = 1$. But since $A^{m+1} - 2A + 1 = 0$, it follows that $H(1) = \frac{2}{m+1} \equiv \frac{2}{n}$.

Thus, $H(x) = 2x/n$ for every $x \in T$, and $H(x) \leq 2x/n$ for every $x \notin T$. We have seen that the simple lottery that assigns probability $\frac{1}{2}$ to each of the extreme points is a best-reply. But G has the same expected value and it shifts weight from the extreme points only to points x for which $(x, H(x))$ lies on the straight line connecting $(0,0)$ and $(1, \frac{2}{n})$. Therefore, G is a best-reply, too. ■