Competition over Agents with Boundedly Rational Expectations*

Ran Spiegler†

June 27, 2005

Abstract

A market model is presented, in which firms and consumers differ in their "market understanding". In the model, rational firms compete in cumulative distribution functions over consumers with bounded ability to grasp statistical data. Increased competition causes firms to increase their effort to obfuscate, instead of increasing their effort to be more competitive. As a result, consumer welfare is not enhanced and may even deteriorate. Specifically, when firms determine the distribution of both price and quality, and the technology for producing quality is convex, 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, Barton Lipman, Michele Piccione, Ronny Razin, Ariel Rubinstein, Asher Wolinsky and numerous seminar audiences 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

In the past three decades we have seen remarkable advances in the economic analysis of market models with informational asymmetries: 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. At the same time, these market models impose a *perceptual* symmetry between firms and consumers, because they retain the meta-level assumption that "the model itself is common knowledge". In other words, all market agents, firms and consumers alike, are assumed to be perfectly able to grasp the market model and the market equilibrium.

In reality, firms and consumers often differ in their ability to understand the market. Firms have more opportunities to learn the market model and the market equilibrium than most consumers, because they interact more frequently with the market and pay undivided attention to it. In addition, many services provided in markets are multi-dimensional. Determining the value of a multi-dimensional service is a difficult computational task for most consumers. Firms, as price setters, are in a position to contribute to this difficulty, e.g. by using complex pricing schemes that exploit this multi-dimensionality.[1]

For instance, consider industries such as retail banking or healthcare. A bank or an HMO provide a large number of services. A bank may charge different fees for different services. Likewise, the quality of treatment offered by an HMO may vary across medical problems. When consumers choose a bank or an HMO, they do not know yet exactly which subset they shall need, and therefore they need to consider a large set of potentially relevant services. The task of evaluating this set is difficult, and many consumers resort to simplifying heuristics. An example of such a heuristic is sampling a small subset of services and evaluating the bank or the HMO on the basis of the sampled services.

Similarly, when a consumer purchases life insurance, he is faced with a contract that specifies many contingencies. Insurance companies may offer different coverage in different contingencies. Trading off across all contingencies is a hard computational task. Once again, a simplifying heuristic for the

[1] As Ellison and Ellison (2004) demonstrate, even in a potentially competitive environment such as internet commerce, retailers react to price search engines - whose aim is to reduce the complexity of consumer choice - with a variety of obfuscation devices, whose aim is to regain this complexity.

consumer would be to consider a small number of contingencies and evaluate insurance companies according to the terms they offer for these contingencies.

In these examples, firms' strategies have a complex multi-dimensional structure. Consumers find it difficult to grasp this structure in its entirety. This limitation may be due to inherently bounded ability to evaluate multi-dimensional objects, or due to lack of opportunities to learn the data. Instead, consumers sample a small part of the firms' strategies and extrapolate naively from their sample. This simplifying heuristic saves considerable cognitive resources, because it can be applied in many market settings.

My objective in this paper is to examine the implications of market competition in the presence of such an asymmetry, focusing on welfare implications. In particular, to the extent that firms can take advantage of consumers' "boundedly rational expectations", will competition among firms mitigate this effect?

To address this question, I study a simple market model with a finite number of firms and a continuum of consumers. Firms provide a service that consists of a continuum of dimensions. Each dimension is associated with some surplus, which the firm chooses how to divide with the consumer. For simplicity, I assume symmetry across dimensions, such that the surplus associated with each dimension is of size 1. When a dimension is interpreted as a contingency, symmetry also means that all contingencies are equally likely. Then, we can represent the firm's strategy as a *cdf* G over the interval $[0, 1]$, such that $G(x)$ is the fraction of dimensions, for which the consumer share in the surplus does not exceed x (and the firm's share is at least $1-x$).

Firms are standard profit maximizers with perfect ability to grasp their opponents' strategies. Consumers, in contrast, choose according to a procedure that reflects their limited understanding of market alternatives. The consumer focuses on one dimension at random and chooses the firm that offers the best terms in this dimension. Formally, this means that the consumer optimizes against a profile of random draws from the firms' *cdf*s (instead of optimizing against the firms' true strategy profile). When the consumer chooses a firm, he receives the expected share in the surplus induced by the firm's *cdf*. Firms take into account the consumers' choice procedure when choosing their strategy.[2]

[2] One could argue that some dimensions are more important or more likely than others, and therefore consumers will predictably check them: e.g., the quality of pediatric services at the HMO, or the amount of money the insurance policy promises your spouse when you die. See Section 2 for a detailed discussion of this point.

The consumers' choice procedure is borrowed from Osborne and Rubinstein (1998), who called it $S(1)$. This modeling device captures in a simplified way the element of "boundedly rational expectations" that I wish to analyze - namely, the "anecdotal" way in which consumers sometimes grasp a complex statistical market environment. $S(1)$ consumers are somewhat like the blind men in the famous parable of the blind men and the elephant, who base their entire conclusions about what an elephant is on the small part that they sample. Such naive extrapolation often leads to wrong conclusions. Firms are like the Buddha, who can grasp the elephant as a whole. As I demonstrate in Section 6, this consumer-firm asymmetry is *not* informational in the usual sense. Specifically, the results of this paper *cannot* be replicated by rewriting the model as a standard game with imperfectly informed consumers and applying sequential equilibrium.

The consumers' reliance on small samples brings to mind the phenomenon which Tversky and Kahneman (1971) discovered experimentally and dubbed "the law of small numbers". Decision makers tend to expect small samples to represent the distributions from which they are drawn. This expectation causes people to be over-confident when drawing inferences from small samples. Consumers who face a choice among multi-dimensional objects and "believe in the law of small numbers", will act on the false belief that even if they survey a small number of dimensions, their decision error will be negligible. The $S(1)$ procedure reflects an extreme degree of this fallacy: $S(1)$ consumers extrapolate from a *single* sample point per firm. In Section 6 I discuss a generalized model, in which consumers extrapolate from $K \geq 1$ sample points per firm.

If firms were restricted to offering degenerate distributions, consumers 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* (by offering different terms in different dimensions), in order to make it harder for consumers to perceive the true value of alternatives. Thus, the firms' strategic considerations involve two effects: competing over consumers and trying to take advantage of their inference errors. The question is how the "competitive effect" and the "obfuscation effect" interact.

The characterization of symmetric Nash equilibrium in Section 3 provides a sharp answer. There is a unique symmetric equilibrium. The equilibrium *cdf* has an expected value of $\frac{1}{2}$, independently of the number of firms. Moreover, when we add a firm to the market, the equilibrium *cdf* is a *mean-preserving spread* of the original one. As the number of firms gets larger, the equilibrium *cdf* tends to the *maximal-variance cdf* - i.e., the distribution that assigns a mass of $\frac{1}{2}$ to each of the extreme points $x = 0$ and $x = 1$. Thus, firms respond to increased competition by cultivating the "obfuscation effect" alone. As a result, increased competition does not enhance consumer welfare.[3]

Welfare considerations in the basic model are purely distributive, because the surplus is fixed. If firms could control the size of the surplus, what would be the *efficiency* implications of market competition? In Section 4 I study an extension of the model, in which firms choose both a price and a quality level for each dimension of their service. I retain the simplifying assumption of symmetry across dimensions (both in terms of the value of each dimension to the consumer and in terms of the firm's cost of producing quality). I assume that the technology for producing quality is convex, such that there is a unique efficient quality level. Consumers choose as in the basic model: they sample one dimension at random and select the firm that offers the highest net value in that dimension. When they choose a firm, they receive the expected net value induced by the firm's *cdf*.

This extended model has a unique symmetric Nash equilibrium. If the number of firms is sufficiently large, the outcome is *inefficient*: firms produce excessive quality in a positive fraction of dimensions. Moreover, the efficiency loss increases with the number of firms. At the same time, industry profits are independent of the number of firms, such that the efficiency loss is entirely borne by the consumers. This non-standard welfare effect is a result of a standard assumption, namely the convexity of technology. An increase in the number of competitors leads to a mean-preserving spread in the distribution over net surplus, due to the same "obfuscation effect" as in the basic model. But because technology is convex, this implies a decrease in expected net surplus.

Although the equilibrium outcome in the above models is not competitive, we cannot say that consumers are "exploited" by firms because they do not have any outside option. In Section 5 I examine whether $S(1)$ consumers with an outside option could be made worse off by their exposure to the market. Specifically, I modify the basic model by adding an outside option to the consumers' choice set. It turns out that as long as the *cdf* associated with the outside option is atomless, all firms choose *cdf*s with an expected

[3] If there are asymmetric equilibria, they are even *less* competitive than the symmetric equilibrium, in terms of consumers' payoffs.

value of $\frac{1}{2}$ in Nash equilibrium. Thus, when the expected value of the outside option is higher than $\frac{1}{2}$, consumers are worse off than when they are barred from entering the market.

When the outside option is associated with a degenerate *cdf* which assigns all weight to some $\mu_0 > \frac{1}{2}$ - capturing situations in which consumers are familiar with the outside option and therefore know its true value - firms choose the maximal-variance *cdf* in symmetric Nash equilibrium. This means that firms respond to an outside option (which consumers know to be attractive) in the same way that they respond to an increased number of competitors - namely by exclusively cultivating the "obfuscation" effect. As the number of competitors gets large, the welfare loss that consumers experience due to their market exposure converges to $\mu_0 - \frac{1}{2}$.

The message of this body of results is simple to state. When rational firms compete over consumers with anecdotal perception of complex data, standard notions of competitiveness have non-standard welfare implications. The reason is that firms respond to increased competition by obfuscating rather than by acting more competitively.

2 The model

A market consists of a set of profit-maximizing firms $\{1, ..., n\}$ and a continuum of consumers. The firms play a simultaneous-move, complete information game. A strategy for a firm is a cumulative distribution function (*cdf*) G_i over the interval $[0, 1]$. The G_i's are *not* required to be continuous. Let T_i denote the support of G_i, and let μ_i denote the expected value of x according to G_i. After the firms make their decisions, each consumer chooses an alternative from the set $\{1, ..., n\}$. Consumers choose according to a procedure called $S(1)$. They draw one sample point from every G_i. Given a sample $(x_1, ..., x_n)$, they choose $i^* \in \arg\max_{j=1,...,n} x_j$. In case of ties, the consumer employs the symmetric tie-breaking rule. The outcome of the consumer's choice is a new, independent draw from G_{i^*} (or any sequence of such draws).

This is a stylized model that is open to more than one interpretation. The primary interpretation that I adopt in this paper is that G_i represents a "cross section". Firms offer a service that consists of a large number (idealized as a continuum) of dimensions. All dimensions are treated symmetrically by firms and consumers alike. Each dimension is associated with a surplus of size 1, such that $G_i(x)$ is the fraction of dimensions for which the consumer's share in the surplus does not exceed x.

For a concrete illustration, consider a bank that offers many financial services. A strategy for the bank is an assignment of fees to services. At the time the consumer chooses a bank, all services have equal importance to him, because he has the vaguest idea of the services he shall actually need. Evaluating the bank's entire fee structure is a computationally difficult task. Therefore, the consumer examines one service at random and chooses the bank that charges the lowest fee for that particular service. Once the consumer chooses a bank, he will use any subset of the bank's services, and will be charged according to the bank's complete pricing strategy.[4]

For every $x \in [0, 1]$, let $H_i(x)$ be the probability that a consumer chooses alternative i, conditional on the event that the realization of G_i in his sample is x. Of course, H_i is induced by $(G_j)_{j \neq i}$. Let EH_i denote the expected value of H_i, where the expectation is taken with respect to G_i.[5] Then, firm i's payoff function is:

$$u_i(G_1, ..., G_n) = [1 - \mu_i] \cdot EH_i(x) \tag{1}$$

In other words, the firm's payoff is equal to its expected share in the surplus conditional on being chosen, multiplied by the fraction of consumers who choose it.

As expression (1) demonstrates, the role of the sampling procedure in the present model is fundamentally different from other models which involve sampling, e.g. search models. In both cases, the consumer chooses the alternative with the highest realization x_i in his sample. However, in a search model the consumer ends up receiving x_i. In contrast, in the present model the consumer receives a new, independent draw from G_i when he chooses firm i.

[4] In the long run, it may well be that the consumer will learn the entire fee structure of all banks. However, it is reasonable to assume that for a long period of time, the consumer will have to rely on a partial picture. As Camerer and Lowenstein (2003, p. 8-9) point out, "many important aspects of economic life are like the first few periods of an experiment rather than the last". I believe that this statement holds rather well for consumption problems such as choosing a bank or an HMO.

[5] When every G_i is atomless and has a well-defined density g_i, then $H_i(x) = \Pi_{j \neq i} G_j(x)$ and $EH_i = \int_0^1 \Pi_{j \neq i} G_j(x) g_i(x) dx$.

The firms' strategies in this model are *cdf*s. However, they are *not* mixed strategies: the payoff function is *quadratic*, rather than linear, in own probabilities. Therefore, we should not expect the familiar indifference property of mixed-strategy equilibria to hold in this model. 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 $G(x) = x$. If firm 3 assigns probability one to some $a \in (0, 1)$, then its payoff is $(1 - a) \cdot a^2$. If the firm switches to a distribution that assigns mass a to $x = 1$ and mass $1 - a$ to $x = 0$, then its payoff is $(1 - a) \cdot a$, which is a strict improvement.

Correlation between the sample and the outcome. The model assumes statistical independence between the consumer's sample and the outcome of his choice. Under the "cross section" interpretation of G_i, this corresponds to the assumption that all dimensions have equal value (and equal likelihood, if different dimensions correspond to different contingencies). This is of course a great simplification. When a consumer chooses a bank, he may have a rough idea of the services that he is more likely to use. Therefore, he will not sample any random service, but rather the service that he is more likely to use. This implies positive correlation between the sample and the choice outcome.

When the consumer knows that he will use exactly one service (and he knows what it is), then the sample and the choice outcome are perfectly correlated, and the model is reduced to standard Bertrand competition. Conversely, if firms are able to control the dimension that the consumer samples (say, through advertising), they can induce (almost) perfectly *negative* correlation between the consumer's sample and the outcome of his choice, and the market outcome will be monopolistic. Exploring a generalized model which allows for arbitrary correlation is left for future research.

Suppose that there is a small number of dimensions which all consumers check because they find them relatively important or likely (e.g., the main account interest rate). Then, firms will behave competitively ($x = 1$) in these dimensions. Still, the firm's services may contain many other dimensions with big overall importance, but consumers may have the faintest idea of which of these dimensions will turn out to be important. If consumers ignore these other dimensions altogether, firms will behave monopolistically ($x = 0$) in them. If consumers evaluate the other dimensions according to the $S(1)$ procedure , the model of this section shall describe the firms' behavior along these dimensions.

An alternative interpretation of G_i. We may interpret G_i not as a "cross section" but as a genuinely random strategy. A realization $x \in [0,1]$ of G_i represents a division of a unit surplus between the firm and the consumer. After the consumer chooses alternative i, the outcome of his choice is a *new, independent draw* from G_i, or any sequence of such draws. Under this interpretation, the $S(1)$ procedure does not reflect inherent bounded rationality, but rather limited familiarity with the market. The asymmetry between the two sides of the market is that firms have already learned the market model and market equilibrium, whereas consumers have just begun to learn, at the time they are required to reach a decision.

For illustration, consider the problem of choosing a professional service such as litigation or plumbing. This is a relatively non-recurring problem. The price or quality of a plumber's visit is a random variable, which is controlled by the plumber, because he may discriminate among visits on an arbitrary basis. How do we choose a plumber in case of a leakage? Lacking repeated experience in the plumbing industry, many of us rely on anecdotes: one friend of ours has had a good experience with one plumber, another friend has had a bad experience with another plumber, and so forth. As a result, we form a very crude perception of the true distribution that characterizes each plumber.

Of course, there may be other reasons for the plumber to discriminate among visits. What may appear like randomization could in fact be a deterministic strategy that conditions on some payoff-dependent variable. However, given the consumers' "anecdotal" assessment of the plumber's strategy, he may have a strict incentive to condition his behavior on the variable even if it is essentially a "sun spot".

The meaning of n. We shall interpret n as the physical number of firms in the market. Alternatively, n may be viewed as the number of alternatives that consumers are aware of. According to this interpretation, the consumer may have actively gathered a sample of size n. When a firm chooses its strategy, it is uncertain whether the consumer will be aware of the firm's existence. However, the firm knows that if it enters the consumer's consciousness, it will face $n-1$ competitors. Under both interpretations, n is a natural indicator of the intensity of competition in the market.[6]

[6] The sample size may be due to implicit search costs. For instance, the consumer may (erroneously) believe that each firm offers a deterministic x, which is drawn from some common distribution F. The value of n is selected optimally, given F and the search cost.

3 Analysis

I begin with a characterization of symmetric Nash equilibrium in this model. Let $G(x,n)$ denote the equilibrium strategy. I sometimes use the abbreviated notation G. Let T denote the support of G.

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

$$G(x,n) = \begin{matrix} \sqrt[n-1]{2x/n} & 0 \le x \le b_n \\ \sqrt[n-1]{2b_n/n} & b_n < x < 1 \\ 1 & x = 1 \end{matrix} \quad (2)$$

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

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

What is the shape of $G(x,n)$? For $n = 2$, $G(x,n)$ is the uniform distribution $U[0,1]$. For every $n > 2$, the support of $G(x,n)$ is $T = [0, b_n] \cup \{1\}$, where b_n decreases with n and tends to zero as $n \to \infty$; $G(x,n)$ contains an atom (whose mass is $1 - A_n$) on $x = 1$; the atom's size increases with n and tends to $\frac{1}{2}$ as $n \to \infty$; $G(x,n)$ contains no other atom. As $n \to \infty$, $G(x,n)$ approaches the distribution that assigns mass $\frac{1}{2}$ to $x = 0$ and mass $\frac{1}{2}$ to $x = 1$. The convergence is fast: $A_6 \approx 0.51$; $b_6 \approx 0.1$. I call the limit distribution "*the maximal-variance cdf*", because it has greater variance than any other *cdf* over $[0,1]$.

A simple calculation establishes the following implication of Proposition 1. Let $\mu(n)$ denote the expected value of x according to $G(x,n)$.

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

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

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

Corollary 1 demonstrates that *an increase in n results in an increase in the variance of the equilibrium cdf, without affecting its mean.* Thus, the number of competitors, normally an indicator of the market's competitiveness, has an orthogonal effect when consumers form beliefs according to the $S(1)$ procedure. Industry profits are $\frac{1}{2}$, independently of n.

Firms in the model have two strategic considerations. First, there is the usual competitive motive, which induces firms to offer attractive distributions: if G_i first-order stochastically dominates G_j, firm i will attract a higher clientele than firm j. This is what I call the "*competitive effect*". Second, there is an incentive to confuse the consumer by introducing greater variance. In this way, a low-μ firm may increase the probability that the consumer will choose it over a high-μ firm. This is what I call the "*obfuscation effect*". A priori, one might expect that increased competition would heighten both effects. The surprising feature of Corollary 1 is that firms respond to greater competition by cultivating the obfuscation effect alone.

Indeed, firms strictly prefer to obfuscate in equilibrium. For $n = 2$, it is easy to verify that although $G(x, 2)$ has full support, the only payoff-equivalent degenerate *cdf* is the one that assigns all weight to $x = \frac{1}{2}$. For $n > 3$, there exists *no* degenerate *cdf* which is payoff-equivalent to $G(x, n)$. This shows once again that randomization in this model is different from ordinary "mixing". Firms in the present model randomize in order to confuse consumers, rather than to surprise competitors.

To illustrate the rough intuition behind part *(ii)* of Corollary 1, consider once again the case of a bank providing multiple financial services. An individual service plays two roles: it attracts clients and it generates revenues. The two roles are independent: the service generates revenues from clients who chose the bank because it offers good terms for the service that the consumer sampled. As the number of competing banks increases, it becomes harder to generate 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.

This argument brings to mind the phenomenon of "loss-leader" pricing by multi-product firms. Proposition 1 suggests an interpretation of this marketing tool as an obfuscation device. As competition becomes more intense, loss-leader pricing is more ubiquitous, but that is not evidence for more competitive behavior, but rather for a greater effort to obfuscate. 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 when physically at the store. However, at that point consumers face a "hold-up" problem because the firm can exploit their search costs. Consumers have rational expectations: they anticipate the hold-up and therefore reduce their willingness to shop. Advertising serves as a commitment device that partially resolves the hold-up problem. This results in a high-variance price distribution inside each store, because firms compete fiercely over advertised items and sell unadvertised items at the monopoly price. One possible lesson from Corollary 1 is that loss-leader pricing may have another rationale when consumers have "boundedly rational expectations". Pursuing this logic in a more concrete I.O. setting is left for future work.

The reasoning behind part *(i)* of Corollary 1 is more subtle. Recall that firm i typically prefers its best-replying strategy G_i to any distribution whose support consists of a single element. However, a lemma that plays a central role in the proof of Proposition 1 (see Corollary 1) establishes another "*indifference principle*": firm i is always indifferent between G_i and a distribution whose support consists of *two* elements, namely the extreme points in T_i, $x_* = \inf(T_i)$ and $x^* = \sup(T_i)$. Moreover, this simple distribution has the expected value as G_i.[7]

Let α denote the mass that such a two-outcome distribution assigns to x^*. The firm's payoff is:

$$[1 - \alpha(x^* - x_*)] \cdot [\alpha H_i(x^*) + (1 - \alpha)H_i(x_*)] \tag{4}$$

In a series of straightforward steps, I show that $x_* = 0$ and $H_i(0) = 0$ for each firm i. It then immediately follows from expression 4 that the expected value of the optimal two-outcome distribution is independent of $H_i(x^*)$, hence of n. But by the "indifference principle", this must also be a property of the expected value of G_i.

Welfare implications

A natural measure of consumers' welfare in this model is their expected payoffs:

$$\sum_{i=1}^{n} \mu_i \cdot EH_i \tag{5}$$

[7]The lemma itself is a consequence of a technical argument originally due to Myerson (1993) - see Appendix.

because μ_i is their expected share in the surplus conditional on choosing firm i and EH_i is the probability that they choose firm i. One could argue that this is not an obvious measure, because $S(1)$ consumers do not maximize their utility against $(G_1, ..., G_n)$. However, the reason they do not maximize is that they do not hold correct beliefs. If we informed the consumers ex-post of the true strategy profile, then presumably they would use expression (5) to evaluate their welfare. Given this criterion, the implication of Corollary 1 is that "greater competition" (in the sense of larger n) has no effect on consumer welfare, because $\Sigma\mu_i EH_i = \frac{1}{2}$ for every $n \geq 2$ in symmetric equilibrium.

Expression (5) does not take into account the consumers' risk attitudes. The justification for welfare judgments that involve risk attitudes is not clear in the present context. The reason is that $S(1)$ consumers behave as if they believe that firms play degenerate *cdf*s. If this is truly what consumers believe, then they experience no subjective uncertainty at the time they make their choices. Hence, it is not clear why their risk attitudes should be relevant to our welfare analysis. At any rate, if we used risk aversion as a criterion for welfare comparisons, then an increase in n would be welfare reducing because $G(x, n+1)$ is a mean-preserving spread of $G(x, n)$.

Asymmetric equilibria

Let us turn to asymmetric equilibria, starting with the case of $n = 2$.

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

The uniqueness result relies on an imitation argument, which is unavailable when $n > 2$. Therefore, I do not know whether asymmetric equilibria exist when $n > 2$. However, the next result shows that *if* asymmetric equilibria exist, they are *less competitive* (in terms of consumer welfare) than the symmetric equilibrium. Thus, the symmetric equilibrium has a special status in the model, because it is the *most competitive equilibrium.*

Proposition 3 *In Nash equilibrium, $\mu_i \leq \frac{1}{2}$ for every firm i, and $\mu_i = \frac{1}{2}$ for at least $n - 1$ firms.*

The reasoning behind this result is as follows. The "indifference principle" referred to earlier holds in asymmetric equilibrium. However, now I cannot rule out the possibility that $H_i(0) > 0$ for some firm i because all its competitors place an atom on $x = 0$. For this firm i, the expected value of the optimal distribution within the restricted class of two-outcome distributions falls below $\frac{1}{2}$. By the "indifference principle", this is also a property of G_i.

The role of the bounds on x

Our model assumes that $x \in [0, 1]$. The motivation for imposing bounds on x is that x represents a division of some fixed surplus between the consumer and the firm. In particular, if we want to interpret $1 - x$ as a price, then we should place an upper bound on x, in order to rule out negative prices.

The main result in this section - namely, Corollary 1 - is independent of the existence or value of an upper bound on x. In contrast, a lower bound on x is crucial for equilibrium existence. If there were no lower bound on x, firms would be able to attract some clientele by assigning positive mass to a sufficiently high x, and then extract an arbitrarily large amount from this clientele by assigning positive mass to arbitrarily low values of x. In general, if the lower bound on x is some $a < 0$, then in symmetric equilibrium consumer welfare is $\frac{1}{2} - a$.

4 Competition and inefficiency

In the model of Section 2, firms choose how to divide a fixed surplus. Therefore, efficiency considerations are irrelevant. In this section, I enrich the model by allowing firms to compete in (distributions over) both quality and prices, thereby introducing efficiency considerations.

In the extended model, the firm's service is characterized by its quality and price. The firm produces quality by incurring a cost $c \in [0, \infty)$, such that the resulting quality is $f(c)$, where f is a strictly increasing and strictly concave production function. Assume that $f(c) - c$ attains a unique maximum at some $c^* \in (0, \infty)$. Denote $s^* = f(c^*) - c^*$. This is the maximal surplus that firms can generate in this model. Given a cost level c, the firm's price p is restricted to lie in $[0, f(c)]$. Consumers' payoff from a pair (c, p) is $x = f(c) - p$. The firm's payoff from (c, p), when chosen by the consumer, is

$p - c$.

A strategy for a firm in this model is a probability distribution over all pairs (c, p) satisfying $c \in [0, \infty)$ and $p \in [0, f(c)]$. Consumers follow the $S(1)$ procedure: they draw one sample point from each firm, and then choose the firm with the highest value of x in their sample. The outcome of the consumer's decision is a new, independent draw from his chosen distribution (or any sequence of such draws).

To illustrate the model, recall the retail banking example of Section 2, and suppose now that a financial service is characterized by a price as well as a level of quality. A bank chooses a price and a quality level for each service. The consumer samples a service at random and chooses the bank that offers the highest net value for that particular service.

At first glance, this model looks like a considerable complication of the basic model, because now a strategy is a distribution over *pairs*. A simplification is immediately made possible thanks to the following observation. If a firm assigns weight to an outcome that generates a consumer payoff of x, it will set c to be the most efficient cost level that is compatible with x. In other words, the firm assigns weight only to outcomes that lie on the Pareto frontier. Thus, for every (c, p) in the support of the firm's strategy: *(i)* if $f(c) - p \leq f(c^*)$, then $c = c^*$; *(ii)* if $f(c) - p \geq f(c^*)$, then $p = 0$. It follows that any inefficiency necessarily takes the form of *over-investment in quality*.

We can now redefine firm i's strategy as a *cdf* G_i over the unidimensional variable $z = c - p$, which gets values in the interval $[-s^*, +\infty)$. As in the basic model, a higher z is more favorable for the consumer and less favorable for the firm. Every $z \leq c^*$ represents an efficient outcome, such that the consumer's payoff is $x = z - [f(c^*) - c^*]$. Every $z > c^*$ represents an inefficient outcome, such that the consumer's payoff is $f(z)$. Firm i's payoff function is thus $\beta_i \cdot EH_i$, where β_i is the expected value of $-z$ according to G_i, and EH_i is the fraction of consumers who choose firm i, given $(G_1, ..., G_n)$. The simplification of the extended model makes it amenable to the same characterization techniques that served us in Section 3.

Proposition 4 *There is a unique symmetric Nash equilibrium in the game. Each firm plays the cdf:*

$$G(z, n) = \sqrt[n-1]{\frac{2(s^* + z)}{s^* n}} \tag{6}$$

defined over the support $[-s^*, s^*(\frac{n}{2} - 1)]$.

Note that in contrast to Proposition 1, the equilibrium strategy in Proposition 4 is atomless. The reason is that c is unbounded from above. This simple formula has the following implications:

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

(i) $\beta = s^*/2$ *for every* $n \geq 2$.

(ii) For every $n > 2(\frac{c^*}{s^*} + 1)$*, the equilibrium outcome is inefficient, such that* $c > c^*$ *with positive probability.*

(iii) The expected value of $f(c) - c$ *according to* $G(z, n)$ *is strictly decreasing with* n *for* $n > 2(\frac{c^*}{s^*} + 1)$.

The number of competitors does not affect equilibrium industry profits. However, starting at some critical number of firms, market competition results in an inefficient outcome. Therefore, all the inefficiency is incurred by the consumers. Moreover, the inefficiency grows with n. Thus, *increased competition in this model reduces total welfare, and the entire welfare loss is borne by the consumers.*

As n tends to infinity, firms concentrate most of their weight near the point $(c, p) = (c^*, f(c^*))$. The remaining weight is concentrated in the range $\{(c, p) \mid c > c^*, p = 0\}$. Thus, for each service, the firm either make an efficient investment in quality and charge the monopoly price, or it makes an inefficiently large investment in quality and charges nothing. The latter services are "loss leaders", designed to attract customers.[8]

Let us sketch the reasoning behind Corollary 2. Expression 6 implies that $G(z, n+1)$ is a mean preserving spread of $G(z, n)$. Therefore, the expected value of any concave function defined on z decreases with n. Consider the function v, defined as follows: $v(z) = s^*$ for $z \leq c^*$ and $v(z) = f(z) - z$ for $z > c^*$. As we noted above, for every z in the support of $G(z, n)$, $z \leq c^*$ if and only if $c = c^*$, and $z > c^*$ if and only if $z = c$. Therefore, the expectation

[8] If we ruled out negative-profit outcomes, we would return to the model of Section 2, except that the size of divisble surplus would be s^*.

of v according to $G(z, n)$ is equal to the expectation of $f(c) - c$ according to $G(z, n)$. By the concavity of f, the expectation of $f(c) - c$ decreases with n.

Thus, two forces are at play in the inefficiency result: the mean-preserving-spread property of $G(z, n)$ and the convexity of the production technology. I find it interesting that convexity of technology - a property that often leads to normatively attractive results in market models - plays a detrimental role in terms of efficiency in the present model.

5 Outside options and "market exploitation"

When consumers make judgment errors, they are vulnerable to being exploited by rational firms. Although the models examined so far yield non-competitive outcomes, we cannot speak of "market exploitation" because consumers have no outside option. In order for the notion of market exploitation to have any meaning, we must endow our consumers with an outside option.

Formally, extend the consumers' choice set to $\{0, 1, ..., n\}$. Let G_0 be a *cdf* associated with the outside option. Extend the $S(1)$ procedure to encompass the outside option: consumers draw one sample point from every G_i, $i = 0, 1, ..., n$, and select the alternative with the highest realization in their sample. Thus, the only difference between G_0 and $G_1, ..., G_n$ is that the former is exogenously given whereas the latter are determined endogenously.

Proposition 5 *Suppose that G_0 is atomless. Then, in Nash equilibrium $\mu_i = \frac{1}{2}$ for every firm i.*

Thus, the existence of an outside option does not affect the expected value of the firms' strategies. The reasoning behind this result is essentially the same as in the case of Corollary 1. Recall that the restriction to symmetric equilibrium in the model of Section 2 implies that $H_i(0) = 0$ for each firm i. The assumption that G_0 is atomless has the same implication in the present model. This in turn implies - invoking the "indifference principle" - that the expected value of firm i's best-replying strategy 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 consumers' choice procedure, by assuming that G_0 is a *degenerate cdf* that assigns probability one to μ_0. This is as if the consumer *knows* the value of his outside option.

Proposition 6 *Let $n > 2$. Suppose that G_0 assigns probability one to some $\mu_0 > \frac{1}{2}$. Then, in symmetric Nash equilibrium firms play the maximal variance cdf.*[9]

When $\mu_0 > \frac{1}{2}$, the outside option is "good", in the sense that its expected value μ_0 exceeds the consumer's true expected payoff in the model without an outside option. However, according to Proposition 6, this does not cause firms to act more competitively. Instead, they raise the variance of their *cdf* to the maximal possible level. The intuition for this result is simple. When consumers know μ_0, firms do not compete at all in the range $x < \mu_0$. They prefer to shift weight in this range to the extreme point $x = 0$, and this generates a large revenue from their clients. Having secured this revenue, the firms can afford to compete fiercely over high realizations in order to attract clients, and this causes them to place all remaining weight on the other extreme point $x = 1$.

It is significant that introducing an attractive outside option (with known value) has the same effect as raising n in the model without an outside option. Both interventions would normally constitute an improvement in the market environment from the point of view of a consumer with rational expectations. The two interventions continue to have a similar effect when consumers have boundedly rational expectations, albeit in an orthogonal direction.

Let us turn to welfare analysis. I continue to use expression (5) to measure consumer welfare, except that the summation is now over $i = 0, 1, ..., n$. According to Propositions 5 and 6, when a consumer ends up choosing a firm, he experiences a welfare loss of $\mu_0 - \frac{1}{2}$ relative to a situation in which he only possesses the outside option. Since $EH_i > 0$ for every firm i, consumers experience "market exploitation" in equilibrium: they are worse off than if they were barred from entering the market.

In the case of a degenerate G_0, we may take this conclusion further. According to Proposition 6, the probability that consumers end up choosing a

[9] When $n = 2$, the same result holds for $\mu_0 > \frac{5}{8}$.

firm in symmetric equilibrium is $1-(\frac{1}{2})^n$ (as long as $\mu_0 < 1$). This expression is increasing in 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 $\mu_0 - \frac{1}{2}$, relative to a world in which only the outside option is available.

Consumers are exploited in this case because they choose a firm over the outside option whenever $\arg\max_{i=1,...,n} x_i > \mu_0$. In other words, they behave as if they believe that for every firm i, G_i assigns probability one to x_i. This is a "belief in the law of small numbers" writ large: consumers behave as if they believe that one sample point drawn from a firm's *cdf* has the same informational content as full knowledge of μ_0. Moreover, they disregard the fact that a firm's *cdf* is a strategic choice that takes into account the consumers' inference procedure, while the outside option is exogenous. For failing to draw these distinctions, consumers suffer a welfare loss.

6 Discussion

The premise of this paper was that consumers and firms often differ in their ability to perceive statistical market regularities. This asymmetry implies that firms might be able to take advantage of consumers, by introducing statistical complexity into the market. 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, consumer welfare and total surplus are not enhanced - and are sometimes diminished - by increased competition.

6.1 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 consumers draw K independent sample points from every G_i, and choose the alternative with the highest *average* realization in their sample. The parameter K reflects the extent to which the consumer's belief formation process departs from full understanding of market alternatives. The larger K, the smaller the departure.

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 represented by expression (1). However, the definition of EH_i is different. Define G_i^K as the *cdf* of the *average* of K independent draws from G_i. Let $H_i(x)$ be the probability that a consumer will choose firm i, conditional on the event that the realization of G_i^K is x. The expectation of H_i in expression (1) should now be taken with respect to G_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 chose G_i^K, rather than G_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. To see why, assume that the points $0, \frac{1}{K}, 1$ all belong to the support of G_i. Assume further that $H_i(\frac{1}{K}) < \frac{1}{K} \cdot H_i(1) + \frac{K-1}{K} \cdot H_i(0)$. Ideally, firm i would like to deviate by shifting all the weight from $x = \frac{1}{K}$ to $x = 0$ and $x = 1$. However, such a deviation is impossible: by the definition of G_i^K, if 0 and 1 belong to the support of G_i, then $\frac{1}{K}$ necessarily belongs to the support of G_i^K.

Nevertheless, some lower bounds on industry equilibrium profits can be obtained. It can be shown that an individual firm cannot do worse than when it plays a simple distribution G^* that assigns mass $\frac{K}{K+1}$ to $x = 1$ and mass $\frac{1}{K+1}$ to $x = 0$. This implies that industry profits are bounded from below by $K^K/(K+1)^{K+1}$. In some cases the lower bound can be improved. For example, suppose that we add an outside option G_0 that assigns all weight to $x_0 = 1$. Then, $H_i(x) = 0$ for every $x < 1$. Therefore, all firms play G^* in Nash equilibrium. When $n \to \infty$, industry profits converge to $\frac{1}{K+1}$. Obtaining tight general lower bounds on industry profits as a function of K is left for future research.

Another way of enriching the $S(1)$ procedure is to endogenize the number of sample points that consumers draw from every G_i. Recall that when G_i is a genuinely random strategy, a sample point in the description of the $S(1)$

procedure is interpreted as a random anecdote. In this case, one could argue that if firm i has a larger clientele than firm j, the consumer would be able to gather more anecdotes about firm i. This extension may have a subtle effect on the market outcome. Recall that firms in this model have an incentive to confuse consumers. If consumers have more accurate perceptions of firms with a larger clientele, it is harder for these firms to take advantage of their clients. This creates an incentive for firms to reduce their clientele, say by offering an unattractive distribution. Thus, firms in this extended model may have a strictly anti-competitive motive.

6.2 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 consumers' *individual* behavior, although non-rational in our market model, might be rational in some other market environment. The answer is clearly affirmative, if consumers believe that each firm offers the same terms in all dimensions, where these terms are drawn from some common distribution. Thus, we can "rationalize" the consumers' individual behavior in the sense that they behave optimally with respect to an incorrect market model.

The 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 consumers are substituted with rational, imperfectly informed consumers. To explore this question, consider the following variant on the model of Section 2. Consumers move after the firms choose their strategies, but they are unable to observe them. However, they can condition their action on a random draw from $(G_1, ..., G_n)$. What is the relation between sequential equilibria in this game and our analysis of the model of Section 2?

There is a sequential equilibrium in this incomplete-information game, in which all firms play the *cdf* given by Proposition 1, and consumers play the strategy induced by the $S(1)$ procedure - i.e., they choose the firm with the highest realization in their sample. For $n > 2$, we need to sustain this equilibrium with suitable out-of-equilibrium beliefs, because the equilibrium distribution does not have a full support. In equilibrium, consumers are indifferent among all firms, hence they do not mind using the decision rule prescribed by $S(1)$.[10] However, there is a huge set of *other* sequential equilibria which sustain many other market outcomes. For example, we can sustain the fully monopolistic outcome (in which all firms assign all weight to $x = 0$), using suitable out-of-equilibrium beliefs. Therefore, we cannot say that the incomplete-information game rationalizes the model of Section 2. A similar argument applies for the model of Section 4.

When we add an outside option with $\mu_0 > \frac{1}{2}$, the mismatch is worse: there exists *no* sequential equilibrium that rationalizes the results of Section 5. The reason is as follows. In sequential equilibrium, it is impossible for all firms to play a *cdf* with $\mu = \frac{1}{2}$ and for consumers to choose firms over the outside option with positive probability. The failure to rationalize the model in this case follows from the rational-expectations aspect of sequential equilibrium. Consumers can never be systematically fooled in sequential equilibrium. They will not choose an alternative that gives them less in expectation than the outside option.

Even if we could rationalize the model, this would not undermine our welfare analysis, which occupies a central place in this paper. A standard model based on full rationality and informational asymmetries might be observationally equivalent to the present model, but it would lead to different welfare implications. As long as one believes in the economic relevance of the premise underlying our analysis, one should find its welfare implications economically relevant.

6.3 Related literature

The $S(1)$ procedure was introduced by Osborne and Rubinstein (1998), who studied a solution concept for games in which *all* players behave according to this procedure. Their focus was therefore on devising a novel equilibrium concept for such a situation. In contrast, the present paper studies a market model, in which only non-strategic agents (the consumers) follow the procedure. Therefore, there is no need to tamper with standard equilibrium concepts.

[10] The out-of-equilibrium beliefs must be *inconsistent* with $S(1)$ in order to sustain the sequential equilibrium. Suppose that one firm deviates from G to a *c.d.f* that assigns probability one to some $x > \frac{1}{2}$, $x \notin T$. If the agent's out-of-equilibrium belief is that this is indeed the firm's strategy, then there exists such a deviation which is profitable for the firm.

Recall that when G_i represents a genuinely random strategy, the meaning of the $S(1)$ procedure is that consumers try to draw inferences about G_i from anecdotes. In this regard, the $S(1)$ procedure is not far in spirit from the theory of case-based decision making, due to Gilboa and Schmeidler (2001), which also aims to describe choice under uncertainty when insufficient familiarity with the environment makes probabilistic reasoning impractical. As in the $S(1)$ model, decision makers in case-based theory base their assessments of actions on random cases. However, case-based theory focuses on the similarity judgments that people make when drawing inferences from anecdotes, whereas the $S(1)$ procedure ignores them.

A companion of the present paper is Spiegler (2003), which studies a "market for quacks", in which firms provide a worthless treatment (in the sense that it has the same probability of success as the consumers' default option). Firms play an ordinary price competition game and consumers choose according to the $S(1)$ procedure. My objectives in that paper are to demonstrate how "the market for quacks" can be active and to study its welfare and comparative-statics properties. The two models differ in several ways. While the consumers' uncertainty is entirely endogenous in the present paper, in Spiegler (2003) it is entirely exogenous. Consequently, the issue of obfuscation is irrelevant and standard mixed-strategy equilibrium analysis applies. More importantly, the fact that consumer uncertainty is exogenous leads to very different welfare implications: the quacks' adverse welfare effects disappear as $n \to \infty$.

The present paper belongs to a small literature which studies market interaction between rational firms and consumers with imperfect perceptions: bounded ability to grasp intertemporal patterns in Piccione and Rubinstein (2003); limited memory in Chen, Iyer and Pazgal (2003); imperfect awareness of contingencies that may arise after the good is purchased in Gabaix and Laibson (2004); and biased beliefs concerning future tastes in DellaVigna and Malmendier (2004) and Eliaz and Spiegler (2004); to mention a few instances.

Within this literature, some works have examined what is also a theme in the present paper, namely rational firms' incentive to randomize when consumers have a limited perception of stochastic environments. Rubinstein (1993) demonstrates that a monopolist may find it optimal to use a probabilistic pricing strategy, in order to discriminate between consumers with diverse abilities to categorize the realization of a random variable. Erev and Haruvy (2003) 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. However, 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.

References

[1] 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.

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

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

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

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

[6] Erev I. and E. Haruvy (2003), "On the Strategic Use of Price Variability", mimeo.

[7] Gabaix X. and D. Laibson (2004), "Shrouded Attributes and Information Suppression in Competitive Markets", mimeo.

[8] Gilboa I. and D. Schmeidler (2001), *A Theory of Case-Based Decisions*, Cambridge University Press.

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

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

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

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

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

[14] 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.

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

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

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

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

[19] Tversky A. and D. Kahneman (1974), "Judgment under Uncertainty: Heuristics and Biases", *Science* **185**, 1124-1131.

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

Appendix: Proofs

Proof of Proposition 1

Lemma 1 *Let* $(G_1, ..., G_n)$ *be a Nash equilibrium. Then, for every firm* i, G_i *is continuous in* $(0, 1)$.

Proof. Assume the contrary - i.e., that there is some firm i, such that G_i assigns an atom to some $x \in (0, 1)$. Then, there exists $\varepsilon > 0$, such that no other firm j assigns any weight to $(x - \varepsilon, x]$ - because by shifting this weight to $x + \varepsilon'$, for some $\varepsilon' < \varepsilon$, firm j could increase EH_j by an amount that is bounded from zero, while infinitesimally reducing μ_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 μ_j. This contradicts the assumption that $(G_i)_{i=1,...,n}$ is an equilibrium. ■

The following lemma involves a technical argument originally found in Myerson (1993). That paper studies a model of electoral competition, in which candidates choose income-redistribution policies, under a fixed-budget condition. Myerson analyzes the implications of this type of competition under various electoral rules. The function H_i in the present model has an analogue in Myerson's model, which he proves to be linear, using essentially the same argument as the lemma below.[11]

Lemma 2 (Linearity of H**)** *Suppose that* G *is a best-reply for firm* i *to* G_{-i}. *Then, there exist numbers* a, c $(a > 0)$, *such that* $H_i(x) = ax + c$ *for every* $x \in T$ *and* $H_i(x) \leq ax + c$ *for every* $x \notin T \cap [\inf(T), \sup(T)]$.

Proof. For expositional convenience, let us first present the argument for the case of G with a finite support $\{x_1, ..., x_K\}$, $K > 2$.

Firm i's payoff is given by $u_i(\cdot) = (1 - \Sigma p(x_k)x_k) \cdot (\Sigma p(x_k)H_i(x_k))$. Suppose that $H_i(\cdot)$ is not linear in the x_k's. Then, there are three outcomes in T, $x_k < x_l < x_m$, such that:

$$\frac{H_m - H_l}{H_m - H_k} \neq \frac{x_m - x_l}{x_l - x_k} \tag{7}$$

Let F_i differ from G only in the probabilities it assigns to x_k, x_l, x_m, such that $p'_l = p_l - \varepsilon - \delta$, $p'_k = p_k + \varepsilon$ and $p'_m = p_m + \delta$, and set ε and δ such that the expected value of x is the same under F_i and G. Suppose that the L.H.S

[11] I thank Ronny Razin for referring me to Myerson's paper.

is higher (lower) than the R.H.S in expression (7). Then, by setting $\varepsilon, \delta > 0$ (< 0), we guarantee that $EH_i(x)$ increases as a result of the deviation.

The extension to the case of G with an infinite support is straightforward. Let x, y, z be three outcomes in T, with $0 \leq x_1 < x_2 < x_3 \leq 1$. By Lemma 1, none of the G_j's contains an atom in $(0, 1)$. By the definition of H, $H_i(x)$ is continuous in $(0, 1)$. Therefore, if G does not assign an atom to x_i, and expression (7) holds for x_1, x_2, x_3, the expression continues to hold if we substitute x_i with any y in a sufficiently small neighborhood of x_i. Therefore, we substitute weight shift from x_i with weight shift from the neighborhood of x_i.

This establishes that there exist $a > 0$ and c, such that $H_i(x) = ax + c$ for every $x \in T$. It remains to be shown that for every $x \notin T$, $\inf(T) < x < \sup(T)$, $H_i(x) \leq ax + c$. Assume the contrary - i.e., that there exists $x \notin T$, $\inf(T) < x < \sup(T)$, such that $H_i(x) > ax+c$. Then, we can find two points in T, y and z, such that by the same reasoning as in the previous paragraphs, we can shift weight in a mean-preserving fashion from one of these two points (or from their close neighborhoods) to x and the third point, and this would increase EH_i.

Note that the above argument presumes that T contains at least two points. Let us prove that this must be the case. Lemma 1 establishes that there are no atoms in $(0, 1)$. An atom of measure one on $x = 1$ yields zero payoffs to the firm. This is inconsistent with profit maximization: if the firm shifts some of the weight away from $x = 1$, it guarantees that both EH_i and μ_i are positive. An atom of measure one on $x = 0$ would induce $EH_i > 0$ only if all other firms also place an atom of measure one on $x = 0$. But then, G is not a best-reply to G_{-i}, by a standard "Bertrand" argument. ■

Corollary 3 (An indifference principle) *Suppose that G is a best-reply for firm i to G_{-i}. Let F be a cdf that satisfies two conditions: (i) $\mu_F = \mu_G$; (ii) $T_F \subseteq T_G$. Then, F is also a best-reply to G_{-i}.*

Proof. Assume the contrary. Then, there must be three outcomes in the support of G, $x_k < x_l < x_m$, such that expression (7) holds. But this contradicts Lemma 2. ■

Lemma 3 *Let $(G_1, ..., G_n)$ be a Nash equilibrium. Then, $\inf(T_i) = 0$ for every firm i.*

Proof. Assume the contrary, and suppose that there exists a player i, such that $\inf(T_i)$ is strictly positive. Denote $\inf(T_i) = x_i^*$. Then, for every $j \neq i$, $H_j(x) = 0$ for every $x < x_i^*$. 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 μ_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 μ_i more than compensates the reduction in EH_i. Therefore, the deviation is profitable.

Note that the argument would fail if $x_i^* = 1$. However, as we saw in the proof of Lemma 2, this can never be the case if G_i is a best-reply to G_{-i}. ■

The above results hold for any Nash equilibrium. From this point, I focus on symmetric equilibria. Let G denote the equilibrium strategy. Then, all firms share the same $H(x)$. Lemma 2 has already established that $H(x)$ is linear over T. Our objective in this proof is to characterize T and H, and this will immediately give us the expression for G. Denote $y = \sup(T)$. Lemma 3 establishes that $\inf(T) = 0$. Without loss of generality, assume that $0, y \in T$.

Step 1: $H(0) = 0$.

Proof: This could be violated only if G placed an atom on $x = 0$. But then, any firm i can shift the atom slightly upwards, and this will increase EH_i by an amount that is bounded away from zero, while increasing μ_i by an infinitesimal amount. Therefore, the deviation is profitable.

Step 2: $\mu = \frac{1}{2}$ and the firms' equilibrium payoff is $\frac{1}{2n}$.

Proof: Consider the simple lottery F, which satisfies: *(i)* $T_F = \{0, y\}$; (ii) $\mu_F = \mu_G$. Because 0 and y are the infimum and supremum of T, such a lottery F exists. By Corollary 3, any firm i is indifferent between G and F. Let α denote the probability that F assigns to y. Then, firm i's payoff from F 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 $[1 - \alpha y] \cdot \alpha H(y)$. The value of α that maximizes this expression is $\alpha = \frac{1}{2y}$, yielding $\mu_F = \frac{1}{2}$. If $\mu_G \neq \frac{1}{2}$, then G is not a best-reply, a contradiction. Therefore, $\mu_G = \frac{1}{2}$. By symmetry, $EH = \frac{1}{n}$. Therefore, the firms' equilibrium payoff is $\frac{1}{2n}$.

Step 3: There exists a number $b \in (0, 1]$, such that $T = [0, b] \cup \{1\}$.

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 $y = 1$.

Second, let us show that G is strictly increasing in $[0, b]$, where $b = \sup(T \backslash \{1\})$. In other words, we need to show that G contains no "holes" below b. The proof is identical to the proof of Lemma 3, and therefore the details are omitted.

Step 4: $H(1) = \frac{2}{n}$.

Proof: As we have seen in the proof of Step 2, each firm is indifferent between G and a lottery F which satisfies: *(i)* $Supp(F) = \{0, y\}$; *(ii)* $\mu_F = \frac{1}{2}$. By Step 3, $y = 1$. Therefore, each firm is indifferent between G and the maximal variance *cdf*. The firm's payoff from the latter is $\frac{1}{2} \cdot \frac{1}{2} H(1)$. Since this expression must be equal to $\frac{1}{2n}$, $H(1) = \frac{2}{n}$.

Step 4 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 5: G must be given by expressions (2)-(3).

Proof: By Step 3, $T = [0, b] \cup \{1\}$. By Lemma 2, 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 $x \in [0, b]$. Therefore, in this domain:

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

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

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

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

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

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

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

which can be rewritten as:

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

and thus we have the desired characterization.

Step 6: The strategy profile $(G, ..., G)$ given by expressions (2)-(3) is a Nash equilibrium.

Proof: First, let us verify that G induces the function H which is given by Figure 1. Given the expression for G, it follows immediately that $H(x) = 2x/n$ for every $x \leq b$ and $H(x) = 2b/n$ for every $x \in (b, 1)$. Let us check that $H(1) = 2/n$. For the sake of convenience, 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, $A^{m+1} + A(m+1) \cdot H(1) = 1$. By expression (3), $A^{m+1} - 2A + 1 = 0$. 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 $x \notin T$. Suppose that there is a *cdf* F which is better than G. Modify F by shifting weight from $(0,1)$ to $\{0,1\}$ in a mean preserving fashion. By the structure of H, this modification cannot be payoff-reducing. Therefore, when looking for profitable deviations, we need only consider deviations to lotteries whose support is $\{0,1\}$. Now, we have already seen that each firm is indifferent between G and the maximal variance *cdf*, which in turn is the optimal lottery within the class of lotteries whose support is $\{0,1\}$. Therefore, G is a best-reply.

Proof of Corollary 1

Part *(i)* is immediate, because the construction of G relies on the result that $\mu = \frac{1}{2}$. As to part *(ii)*, a simple calculation shows that for every $x \leq b_{n+1}$:

$$\sqrt[n]{\frac{2x}{n+1}} > \sqrt[n-1]{\frac{2x}{n}}$$

At the same time, $A_{n+1} < A_n$. Thus, as we move from n to $n+1$, the weight in (b_n, b_{n+1}) is shifted both leftward to $(0, b_n)$ and rightward to $x = 1$, and the weight in $(0, b_n)$ is shifted leftward.

Proof of Proposition 2

Let us first state several properties of the supports of G_1 and G_2. First, note that $T_1 = T_2 = T$, up to zero-measure differences. Assume the contrary, and suppose that firm i assigns weight to some interval (b, c), whereas firm j does not. Then, firm i can profitably deviate by concentrating all this weight nearer b. Second, denote $b = \sup(T \backslash \{1\})$. Then, both G_1 and G_2 are strictly increasing in $(0, b)$. That is, the support does not contain "holes" between 0 and b. The proof is the same as in Lemma 3, and therefore omitted. Let y denote $\sup(T)$.

By Lemma 1, both G_1 and G_2 are continuous in $(0,1)$. Moreover, at least one firm, say firm 2 without loss of generality, does not place an atom on $x = 0$. Therefore, G_2 is continuous in $[0,1)$, such that for every $x \in [0, b)$, $H_1(x) = G_2(x)$. By Lemma 2, $H_1(x)$ is linear over the support of G. It follows that for every $x \in [0, b]$, $G_2(x) = x \cdot \frac{H_1(y)}{y}$.

By the proof of Proposition 2, $\mu_1, \mu_2 \leq \frac{1}{2}$. Therefore, for every $i = 1, 2$, firm i's payoff is at least $\frac{1}{4}$. Otherwise, the firm could deviate from G_i to

G_j (i.e., by imitating its opponent). By symmetry, both firms would have $EH = \frac{1}{2}$, and because $\mu_j \leq \frac{1}{2}$, the firm's payoff would be no lower than $\frac{1}{4}$, a contradiction.

Because G_2 does not place an atom on $x = 0$, $H_1(0) = 0$, and therefore, by the proof of Proposition 2, $\mu_1 = \frac{1}{2}$. Moreover, by Corollary 3, firm 1 is indifferent between G_1 and the lottery that assigns probability $\frac{1}{2y}$ to $x = y$ and probability $1 - \frac{1}{2y}$ to $x = 0$. Thus, firm 1's payoff is $\frac{1}{2} \cdot \frac{1}{2y} H_1(y)$, and $EH_1(x) = \frac{1}{2y} \cdot H_1(y)$. Because firm 1's payoff is at least $\frac{1}{4}$, $\frac{1}{2} \cdot \frac{1}{2y} H_1(y) \geq \frac{1}{4}$. Therefore, $H_1(y) \geq y$.

Suppose that $y < 1$. Then, because G_2 contains no atom in $[0, 1)$, $H_1(y) = 1$ and G_2 is uniform on $[0, y]$. Therefore, $EH_1(x) = \frac{1}{2y}$ and $\mu_2 = \frac{y}{2}$. Because there are two firms in the market, $EH_2(x) = 1 - EH_1(x)$. Therefore, firm 2's payoff is $\frac{y}{2} \cdot (1 - \frac{1}{2y})$. Because $\frac{y}{2} \cdot (1 - \frac{1}{2y}) \geq \frac{1}{4}$, $y \geq 1$, a contradiction.

Now suppose that $y = 1$. We have shown that $H_1(1) \geq 1$. Therefore, $H_1(1) = 1$, which means that G_2 places no atom $x = 1$. We have shown that $G_2(x) = x \cdot \frac{H_1(y)}{y}$ for every $x \in [0, b]$. Therefore, $b = y$ and $G_2(x)$ is the uniform distribution over $[0, 1]$. In particular, $\mu_2 = \mu_1 = \frac{1}{2}$. Therefore, it must be the case that $H_2(0) = 0$ - otherwise, by the proof of Proposition 2, μ_2 would have been strictly lower than $\frac{1}{2}$ - such that G_1 does not place an atom on $x = 0$. Note that since G_2 assigns weight to any $x < 1$, G_1 cannot place an atom on $x = 1$ (because then it would not be optimal for firm 2 to assign weight to values of x that are close to 1). Therefore, G_1 must be the uniform distribution over $[0, 1]$.

Proof of Proposition 3

Lemmas 1-3 and Corollary 3 hold for arbitrary equilibria. Let us now show that $H_i(0) = 0$ for at least $n - 1$ firms. 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, and any firm can profitably deviate by shifting the atom slightly upwards. 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 replicates Step 2 in the proof of Proposition 1. Let y denote the supremum of T_i. Consider the following mean-preserving modification of G_i. By Lemma 3, $\inf(T_i) = 0$. Without loss of generality, let $0 \in T_i$. Let F satisfy $\mu_F = \mu_i$ and $T_F = \{0, y\}$. By Corollary 3, firm i is indifferent between G_i and F. Let α denote the probability that F assigns to y. Then, firm i's payoff from F is

$$[1 - \alpha y - (1 - \alpha) \cdot 0] \cdot [\alpha H_i(y) + (1 - \alpha) H_i(0)]$$

If $H_i(0) = 0$, then the value of α that maximizes this expression is $\alpha = \frac{1}{2y}$, yielding $\mu_F = \frac{1}{2}$, such that $\mu_i = \frac{1}{2}$. However, it is possible that for exactly one firm j, $H_j(0) > 0$, such that the optimal α yields $\mu_F < \frac{1}{2}$, such that $\mu_i < \frac{1}{2}$.

Proof of Proposition 4

Let us use the abbreviated notation G for the equilibrium strategy. Let T denote the support of G. Denote $y = \sup(T)$. Recall that $-z$ denotes the firm's profit. Thus, $z = -s^*$ is the point of the firm's maximal possible profit, whereas any $z > 0$ represents a loss for the firm. It follows that we can replicate many of the steps in the proof of Proposition 1. Moreover, because the firms' strategy space places no upper bound on c (unlike the basic model of Section 2), G is atomless. Therefore, G is continuous and strictly increasing over $T = [-s^*, y]$. Also, $H(\cdot)$ is linear over T. The proofs mirror those given in the proof of Proposition 1.

By the above properties of G, $H(-s^*) = 0$ and $H(y) = 1$. By the linearity of H over T:

$$H(z) = \frac{z + s^*}{y + s^*} \tag{8}$$

for every $z \in [-s^*, y]$.

The same indifference principle stated in Corollary 3 applies in the present context. Therefore, there exists a simple lottery F with support $\{-s^*, y\}$, such that: *(i)* The expected value of $-z$ according to F is equal to the expected value of $-z$ according to G; *(ii)* the firm is indifferent between G and F; *(iii)* F is the optimal lottery within the restricted family of lotteries with support $\{-s^*, y\}$. The proof is identical to the analogous part in the proof of Proposition 1.

The firm's payoff from a lottery F that assigns probability α to $z = y$ and probability $1 - \alpha$ to $z = -s^*$ is:

$$[-\alpha y + (1 - \alpha) \cdot s^*] \cdot [\alpha H(y) + (1 - \alpha) H(-s^*)] \tag{9}$$

The optimal such lottery satisfies $\alpha = s^*/2(s^* + y)$. It is easy to show that the expected value of $-z$ according to this lottery is $s^*/2$. By symmetry,

$EH = 1/n$, hence the firm's equilibrium payoff is $s^*/2n$. Substituting back in Expression 9 and using our knowledge that $H(y) = 1$, we obtain:

$$y = s^*(\frac{n}{2} - 1)$$

Substituting in Expression 8, we obtain:

$$H(z) = \frac{2(z + s^*)}{s^* n}$$

and because $H(z) = G^{n-1}(z)$ for every $z \in T$, we obtain the desired expression for G.

Proof of Corollary 2

Part *(i)* is immediate, because G is constructed on the basis of the result $\beta = \frac{s^*}{2}$. Part *(ii)* is a simple consequence of the observation that $y > c^*$ for every $n > 2(\frac{c^*}{s^*} + 1)$. Part *(iii)* is more involved. A simple calculation establishes that for every z, $\int_0^z G(w)dw$ is increasing with n. By a well-known result (see Mas-Colell, Whinston and Green (1995), p. 198), this is equivalent to $G(z, n+1)$ being a mean preserving spread of $G(z, n)$. Define the function v as follows: $v(z) = s^*$ for $z \leq c^*$ and $v(z) = f(z) - z$ for $z > c^*$. Because f is strictly concave, v is concave, and strictly concave in the range $z > c^*$. Therefore, because $G(z, n + 1)$ is a mean preserving spread of $G(z, n)$, the expectation of $v(z)$ is decreasing with n. Moreover, the expectation of $v(z)$ is strictly decreasing with n, in the range $n > 2(\frac{c^*}{s^*} + 1)$. Recall that $c = c^*$ for every $z \leq c^*$ and $c = z$ for every $z > c^*$. Therefore, the expectation of $f(c) - c$ according to G is equal to the expectation of $v(z)$ according to G. It follows that the expectation of $f(c) - c$ is strictly decreasing with n, for $n > 2(\frac{c^*}{s^*} + 1)$.

Proof of Proposition 5

Lemmas 1-3 and Corollary 3 hold for arbitrary equilibria in the model without an outside option. Because G_0 is assumed to be atomless, these results continue to hold in the present context. The same assumption implies that $H_i(0) = 0$ for every firm i. Therefore, using the same reasoning as in Step 2 in the proof of Proposition 1, it follows that in Nash equilibrium, $\mu_i = \frac{1}{2}$ for every firm i.

Proof of Proposition 6

Let G be the symmetric equilibrium strategy. By definition, $H(x) = 0$ for every $x < x_0$. Therefore, $G(x) = G(0)$ for every $x \leq x_0$. By the proof of Proposition 2, $\mu_G = \frac{1}{2}$, and each firm is indifferent between G and the maximal variance *cdf* (which assigns probability $\frac{1}{2}$ to each of the extreme points $x = 0$ and $x = 1$). Therefore, the firm's equilibrium payoff is $\frac{1}{2} \cdot \frac{1}{2} H(1)$. By symmetry, $EH = \frac{1}{n} \cdot [1 - G(0)]^n$. Therefore, the firm's equilibrium payoff is equal to $\frac{1}{2} \cdot \frac{1}{n} \cdot [1 - G(0)]^n$. It follows that $H(1) = \frac{2}{n} \cdot [1 - G(0)]^n$.

Assume that G assigns positive weight 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 weight it assigns to $(x^*, x^* + \varepsilon)$ to some $x > x_0$ arbitrarily close to x_0. Then, it reduces μ_i 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 $\varepsilon > 0$ is sufficiently small, the reduction in μ_i more than compensates the reduction in EH_i. Therefore, the deviation is profitable.

By Lemma 2, H is linear in the support of G. It must be the case that $0 \in T$ - otherwise, G would assign positive weight only to elements above x_0, and since $x_0 > \frac{1}{2}$, we would have a contradiction with $\mu_G = \frac{1}{2}$. Because $H(0) = 0$ and G assigns positive weight 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, $H(x) = [G(0)]^{n-1}$ as x approaches x_0 from above.

We are now able to use our expressions for $\lim_{x \to x_0^+} H(x)$ and $H(1)$, to derive the following identity:

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

But $G(0) \leq \frac{1}{2}$, for otherwise μ_G would not be equal to $\frac{1}{2}$. Moreover, we assumed that $x_0 \geq \frac{1}{2}$. Combining the two inequalities yields $2n + 1 \geq 2^n$, a contradiction for $n > 2$. The only remaining possibility is that G assigns probability $\frac{1}{2}$ to each of the extreme points. 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.