THE MARKET FOR QUACKS

by

Ran Spiegler

Working Paper No.21-2003

September, 2003

The Foerder Institute for Economic Research
and
The Sackler Institute of Economic Studies

The Market for Quacks*

Ran Spiegler†

Previous version: March 9, 2003
This version: July 13, 2003

Abstract

A group of "quacks" plays a price-competition game, facing a continuum of "patients", who recover with probability α whether or not they acquire a quack's treatment. If patients were rational, the market would be inactive. I assume, however, that patients choose according to a boundedly rational procedure due to Osborne and Rubinstein (1998): they sample every alternative and choose the best alternative in their sample. This procedure captures an aspect of Tversky and Kahneman's "representativeness" heuristic, namely people's tendency to use limited experience to form quality assessments, ignoring base rates and sample size.

I show that this element of bounded rationality has significant implications. The market for quacks is active. Quacks charge positive prices and inflict a welfare loss on patients, which may increase with the number of quacks. As α decreases - i.e., when the situation becomes more "hopeless" - expected price *increases* and demand decreases. As $\alpha \to 0$, market equilibrium converges to a state of monopolistic competition. Replacing a quack with an expert leaves the other quacks' performance unaffected.

*I am grateful to Eddie Dekel, Kfir Eliaz, Volker Nocke and Ariel Rubinstein, for their helpful comments.

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

1 Introduction

In economic models, decision makers are traditionally viewed as Bayesian rational agents, who make statistical inferences in accordance with the laws of probability. By now, it has become widely accepted (for instance, see Rabin (1998)) that the systematic departures from Bayesian inference documented by psychologists (e.g., Kahneman, Slovic and Tversky (1982)) are in principle relevant to economic behavior. Given economists' traditional focus on market behavior, it is particularly interesting to examine how agents' departures from Bayesian inference would affect our conclusions regarding their market performance.

This paper studies market interaction between Bayesian rational firms on one side of the market, and non-Bayesian, boundedly rational consumers on the other. While firms are standard profit maximizers, consumers follow a decision rule in the spirit of the Tversky-Kahneman descriptive theories of judgment under uncertainty. My objective in this paper is to explore how consumers' boundedly rational statistical inferences affect the market performance of firms and consumer welfare. On one hand, bounded rationality may expose consumers to exploitation by rational firms. On the other hand, competition among firms may mitigate this exploitative effect. The interplay between these two forces is the subject of this paper.

I conduct this investigation in the context of the simplest possible market model. A group of n identical firms plays a standard price-competition game over a continuum of identical consumers. Every consumer enters the market with some problem. If he acquires the "treatment" offered by firm i, he recovers with some positive probability $\alpha < 1$. However, even if the consumer defaults and acquires none of the treatments offered in the market, he recovers with some positive probability $\alpha_0 \leq \alpha$. A consumer's utility from an alternative is equal to the rate of recovery (I also use the term "success rate") that he assigns to the alternative, minus its price.

Throughout the paper, I will refer to the firms as "*healers*" and to the consumers as "*patients*". As the nicknames suggest, I have in mind industries such as unconventional medicine, psychotherapy, self-help, business consulting, etc. The relevant characteristic of these markets is that with some probability, the goods and services traded fail to fix the consumer's problem, and with some probability, the consumer's problem gets fixed even if he does not acquire any of these goods and services.

We will focus on the special case, in which *the patients' rate of recovery* *is independent of their decision*, i.e., $\alpha_0 = \alpha$. In this case, I shall refer to the healers as "*quacks*", because they have absolutely no advantage over the default. Figuratively speaking, they bottle water and sell it as a medicine. Of course, this is an idealization. Even in the most blatant cases of quack medicine, placebo effects endow healers with some advantage over the default. I focus on this idealized case for expositional purposes. As we shall see, the results in this paper extend easily to the case in which the healers have genuine healing powers relative to the default.

If patients were rational, they would regard the entire industry as providing a worthless treatment, and so the "market for quacks" would be inactive. This is where I introduce a modeling innovation: patients choose according to the following choice procedure. Each patient independently samples each of the $n+1$ alternatives K times. A patient's sample assigns some empirical success rate a_i to alternative i. The patient chooses the alternative i that maximizes $a_i - p_i$ in his sample. The healers take into account the patients' choice procedure when calculating their profits.

The patients' choice procedure is borrowed from Osborne and Rubinstein (1998), who called it $S(K)$ and studied normal-form games with players who choose their strategy according to this procedure. In the present context, the $S(K)$ procedure serves as a model of the element of bounded rationality that I wish to focus on, namely the tendency of consumers to form quality assessments on the basis of limited experience, ignoring the fact that this experience was randomly generated.

This tendency is in fact an important aspect of the Tversky-Kahneman "*representativeness*" heuristic. (See Kahneman, Slovic and Tversky (1982), Part II.) In the context of inferences from random samples, representativeness means that people expect a small sample to have the same shape as the underlying probability distribution from which it is drawn. Tversky and Kahneman used representativeness to explain behavioral anomalies such as "insensitivity to base rates" and "the law of small numbers". (Tversky and Kahneman (1974)). The $S(K)$ procedure mimics these features. Our patient maximizes expected utility against the empirical distribution of success rates given by his sample, as if it were the true distribution. His judgments end up being insensitive to the prior rate of recovery and the sample size. Thus, one merit of our model of the patients' behavior is that it is grounded in the Tversky-Kahneman theories of judgment under uncertainty.

Note the asymmetry between healers and patients: the former are rational, whereas the latter employ the $S(K)$ procedure. It is not that I find patients inherently irrational comparing to healers. Rather, healers have much more scope for learning than patients. Using the terminology of the model, healers have $K = \infty$ whereas patients have $K < \infty$. Many industries that fall under the model's scope are not learning-friendly. Choosing a treatment for an illness for which there exists no standard treatment, a therapist at a time of emotional crisis, or a consultant to save a business in trouble, are choice problems that a single decision maker rarely encounters. Moreover, it is hard to generalize from other people's experience, either because they may be reluctant to share them candidly, or because personal circumstances are too idiosyncratic. At any rate, K "parameterizes" the patients' departure from Bayesian rationality: as K increases, our $S(K)$-patient and a standard rational patient become more likely to reach the same decisions.

I devote most of the analysis to the simple case of $K = 1$. In this case, patients form quality assessments on the basis of a *single* observation per alternative. The price-competition game has a unique Nash equilibrium, which is symmetric and mixed. Equilibrium strategies are given by a simple formula. For every α, the "market for quacks" is active. Healers charge positive prices. This pricing behavior reflects an implicit claim to skills that the healers do not really possess. Therefore, on account of their equilibrium behavior, I refer to the healers as "*charlatans*".

The healers' charlatanry inflicts a welfare loss on the patients: those who end up acquiring the healers' treatments are worse off than those who end up choosing the default. The size of the welfare loss is not monotonically decreasing in n: it attains a maximum at $n^* \geq 2$ whenever $\alpha \lessapprox 0.39$. It follows that *greater competition may actually increase the welfare loss inflicted on patients*. The reason is that the patients' choice procedure induces an aggregate demand function which is increasing in n, and this force may outweigh the competitive pressures generated by a larger number of healers. The patients' welfare loss can be substantial: for every $\alpha < \frac{1}{2}$, we can find a number of healers $n \geq 2$ such that the welfare loss exceeds $\frac{1}{4}$; and as α approaches zero, the maximal possible welfare loss given α tends to $\frac{1}{e}$.

Turning to comparative statics, the expected equilibrium price is *decreasing* in α. In particular, as α tends to zero, the market equilibrium converges to a state of *monopolistic competition*: every healer is facing a demand which is insensitive to competitors' prices, earns zero profits, and charges the "monopoly price" $p = 1$. The intuition for this result is simple. As α decreases, multiple successes are less likely to occur in a patient's sample. This weakens competitive pressures and causes prices to increase. At the same time, a lower rate of recovery decreases demand because fewer patients have a good experience with healers. As α approaches zero, the market is almost totally differentiated because the chances that a patient has more than one successful healer in his sample is negligible. Thus, every healer has a monopolistic power over a small group of patients.

This characterization ties together several intuitions about real-life charlatanry. As α approaches zero, the healers over-pricing gets worse. This is consistent with the intuition that charlatanry becomes more extreme as the underlying situation becomes more "hopeless". Also, as α approaches zero, every healer acts more and more like a "*guru*": he attracts a small crowd of admirers, who believe that he has strong healing powers and that he has no substitute. Thus, the model demonstrates the phenomena of charlatanry and "guruism" may emerge naturally in the market for quacks. Moreover, they are exacerbated when the rate of recovery in the industry decreases. These linkages are a consequence of our simple behavioral assumption, in the spirit of the Tversky-Kahneman theories of judgment, namely that patients use limited experience to evaluate healers, without taking into account the randomness of their experience.

What happens to our results when healers have genuine healing powers relative to the default - i.e., when $\alpha_0 < \alpha < 1$? In this case, the rational-patients benchmark collapses into standard Bertrand competition, such that equilibrium prices and profits are zero. In our model with $S(1)$-patients, equilibrium behavior is exactly the same as in the case of $\alpha = \alpha_0$. Market charlatanry persists: healers charge prices higher than $\alpha - \alpha_0$ with positive probability, which decreases in α. Thus, our substantive conclusions do not rely on the assumption that healers are quacks. The case of absolute quackery is interesting, however, because it demonstrates that a market which fails to exist when consumers are rational, turns out to have a rich I.O. characterization and substantial welfare implications when consumers make boundedly rational inferences.

In the remainder of the paper, I examine the robustness of quack behavior to a couple of perturbations. First, I raise the recovery rate associated with a single healer, so as to turn him from a "quack" into an "expert". This turns out to have no impact whatsoever on the other healers' equilibrium performance: they play the same strategy and earn the same payoffs as in the basic model. In order to crowd quacks out of the market, there must be multiple experts in the market. Second, I assume that healers can credibly disclose their success rates. (I allow for heterogeneity among healers and patients.) Given the patients' choice procedure, *disclosure of success rates turns out to be a dominated strategy*, even for high-quality healers. In both variants of the model, the message is similar: high-quality healers do *not* necessarily crowd low-quality healers out of the market.

Finally, I explore the general case of $K > 1$. Existence of Nash equilibrium is guaranteed if the distribution of the patients' willingness to pay is atomless. I provide two asymptotic results that illustrate the effect of patients' bounded rationality on their market performance. As $\alpha \to 0$, equilibrium prices converge to $\frac{p^*}{K}$, where p^* is the price that a monopolistic healer with $\alpha = 1$ would charge. For every α, expected equilibrium price converges to zero as $K \to \infty$. Both results demonstrate that greater rationality on the consumers' part can mitigate charlatanry and lead to a more competitive outcome.

To summarize, the paper makes two contributions to the study of competition among rational firms in the face of boundedly rational consumers. First, the $S(K)$ procedure is a modeling tool that mimics an aspect of the Tversky-Kahneman "representativeness" theory of judgment, and as such, it enables us to analyze the extent to which consumers who make boundedly rational statistical inferences are exposed to exploitation by rational firms in an otherwise-competitive market. Second, the results help clarifying the nature of "institutions" such as charlatans and gurus.

The paper proceeds as follows. Section 2 presents the basic model for the case of $K = 1$. Section 3 analyzes market equilibrium. Section 4 perturbs the basic model by replacing a quack with an expert. Section 5 perturbs the model by allowing type disclosure. Section 6 analyzes the general $K > 1$ case. Section 7 offers concluding remarks and a related-literature survey. In particular, it contrasts the bounded-rationality modeling approach to the phenomenon of charlatanry with a more orthodox approach based on adverse selection.

2 A Basic Model

A two-sided market consists of a continuum of measure one of identical consumers ("*patients*" henceforth) on one side and n identical firms ("*healers*" henceforth) on the other. When a patient in the model acquires the treatment of a healer $i \in \{1, ..., n\}$, he "recovers" with probability $\alpha \in (0, 1)$. I use the terms "recovery" and "success" interchangeably. The patient can also choose the default option, denoted $i = 0$, in which case he recovers with probability $\alpha_0 \in (0, 1)$. Every patient is willing to pay 1 for sure recovery. (As shall become clear, we need not address the patients' risk attitudes at this stage.) Healers are standard profit maximizers. They compete by choosing prices simultaneously. Denote healer i's price by p_i. Of course, $p_0 = 0$. I assume that the healers' activity entails no cost, and I abstract from moral-hazard considerations.

Let us focus on the special case of $\alpha_0 = \alpha$. *healers have no advantage over the default.* (This extreme assumption will be relaxed towards the end of Section 3.) In this case, I find it apt to refer to the healers as "*quacks*", because their treatments are equivalent to the default. If patients were standard rational agents, they would not be willing to pay anything to the quacks, and the market would be inactive.

This is where I introduce a modeling innovation: patients choose according to a procedure called $S(1)$, due to Osborne and Rubinstein (1998). Each patient samples every alternative (including the default) once. For every $i = 0, 1, ..., n$, let x_i denote the outcome of the patient's draw of alternative i: $x_i = 1$ (recovery) with probability α and $x_i = 0$ (no recovery) with probability $1 - \alpha$. The x_i's are independently drawn. Given the realization of his sample, the patient chooses an alternative $i \in \arg\max_{i=0,1,...,n} x_i - p_i$. In case of ties, assume that the patient chooses the alternative with the highest p_i. If a tie remains, apply the usual symmetric probabilistic tie-breaking rule.

The healers take into account the patients' choice procedure when calculating their profits. For example, if $p_1 > p_j$ for every $j > 1$, then healer 1's profits are equal to $p_1 \cdot \alpha \cdot (1 - \alpha)^n$, because the healer's clientele consists of all the patients who had a good experience with him and a bad experience elsewhere. On the other hand, if $0 < p_1 < p_j$ for every $j > 1$, then healer 1's profits are equal to $p_1 \cdot \alpha \cdot (1 - \alpha)$, because the healer's clientele consists of all the patients who had a good experience with him and a bad experience with the default.

The $S(1)$ procedure captures an extreme form of a common behavioral regularity, namely people's tendency to draw sweeping inferences from limited experience, ignoring the fact that this experience was randomly generated. Specifically, they tend to neglect base rates and sample size. Tversky and Kahneman (1974) interpreted this tendency as a manifestation of the "representativeness" heuristic. Thus, a merit of the $S(1)$ procedure is that it is grounded in the Tversky-Kahneman theories of judgment under uncertainty.

To see the relevance of this behavioral assumption, recall that we are interested in markets for treatments that are mainly sought at times of crisis. It is in the nature of such markets that they provide consumers with rare opportunities to learn the rate of recovery associated with each alternative. Thus, patients are typically forced to rely on small samples when they form quality judgments. The $S(1)$ procedure takes an extreme case, in which the patient's choice is based on a single observation per alternative. However, the small sample size is only part of the story. More importantly, the $S(1)$-patient errs in viewing the small sample as being fully informative of the healers' quality. He forms a deterministic action-consequence correspondence on the basis of one observation per healer, ignoring the probability distribution from which the x_i's were drawn.

The sampling procedure should not be taken literally. I do not describe how patients ended up having samples at their disposal, and sampling costs are irrelevant. The $S(1)$ procedure is not a model of the consumers' data gathering process, but a model of how they make inferences from gathered data. This is in sharp contrast to the rational-search literature, which assumes Bayesian inference and focuses attention on the way consumers optimally design samples when there are sampling costs. (Section 7 contains a discussion of related search literature.)

I conclude this section with a comment on the interpretation of mixed strategies in this model. Healers are allowed to use mixed strategies. However, once a price p_i has been realized, healer i is committed to it as far as the patients are concerned. In particular, the value of p_i is fixed across the patients' samples. The patients know the exact prices; the only source of variance in their samples is the healers' imperfect success rate, which is exogenously given. Thus, when a healer employs a mixed strategy, he introduces uncertainty into his opponents' environment, but not into his patients'.

3 Equilibrium

This section is devoted to analyzing Nash equilibrium in the price-competition game with $S(1)$-patients. The following proposition is the basic result of this paper. (The proof appears in the appendix.)

Proposition 1 *There is a unique Nash equilibrium in the game. Every healer plays the same mixed strategy, given by the c.d.f.*

$$G(p) = \frac{1}{\alpha} \cdot [1 - \frac{1-\alpha}{\sqrt[n-1]{p}}] \tag{1}$$

over the support $[(1-\alpha)^{n-1}, 1]$.

The price-competition game has a unique Nash equilibrium, which is symmetric and mixed. To use the terminology of search literature, the game results in equilibrium price dispersion. The healers' equilibrium strategy has a simple functional form. For instance, when $n = 2$ and $\alpha = \frac{1}{2}$, the induced density function is $g(p) = p^{-2}$, defined over the support $[\frac{1}{2}, 1]$.

Let us sketch the reasoning behind the uniqueness and symmetry properties. For this purpose, let us take it for granted that equilibrium strategies are atomless. Consider a healer who charges the highest prices in the market with some probability. The only patients that he attracts are those having a good experience with him and a bad experience elsewhere. Thus, he exercises monopoly power over them. Therefore, the highest price in the market must be the "monopoly price" $p = 1$. This yields the expression $\alpha(1-\alpha)^n$ for this healer's equilibrium payoff. But this must also be the equilibrium payoff for all other healers. Otherwise, a less profitable healer would be able to deviate profitably by imitating the behavior of a more profitable healer. By a similar "imitation" argument, the support of the healers' mixed strategies must be identical (up to measure-zero differences). Because the game is symmetric and every healer's payoff is monotonically increasing in the opponents' prices, the fact that equilibrium payoffs and supports are identical across healers implies that we can apply a symmetry argument to obtain the desired result.

The following corollary derives equilibrium expected price as a function of the rate of recovery α. (The proof merely involves taking expectations and therefore omitted.)

Corollary 1 *Healers' expected equilibrium price is given by:*

$$E(p) = \begin{array}{ll} -\frac{1-\alpha}{\alpha}\ln(1-\alpha) & , n = 2 \\ \frac{1-\alpha}{\alpha(n-2)}[1-(1-\alpha)^{n-2}] & , n > 2 \end{array}$$

In both cases, $E(p)$ is strictly decreasing in α. In particular:

$$E(p) \underset{\alpha\to 1}{\to} 0$$
$$E(p) \underset{\alpha\to 0}{\to} 1$$

The intuition behind the comparative statics is simple. As α decreases, the probability that a patient's sample will contain multiple successes decreases. Thus, a decrease in α leads to weaker competitive pressures, such that expected prices approaches the monopoly price 1. At the same time, a decrease in α causes the demand for healers to shrink. (The fraction of patients who acquire the treatment of a healer is $1 - \alpha - (1 - \alpha)^{n+1}$: this is the probability that a patient has a bad experience with the default and at least one good experience with a healer.) Thus, the comparative statics combine two effects of decreasing rates of recovery: shrinking demand and higher prices.

As α approaches zero, market equilibrium tends to a state of *monopolistic competition*: every healer faces a demand which is virtually insensitive to competitors' prices, and his profits approach zero. At the other extreme, as α approaches one, market equilibrium converges to the rational-patients benchmark of zero prices expected prices and profits converge to zero. For every α, healers have a clientele of positive size. That is, the "market for quacks" is always active in equilibrium.

The healers' equilibrium behavior marks them as "charlatans": they charge a positive price for their services, although in reality their treatments are worthless. Moreover, the false pretense implicit in their over-pricing gets worse as α decreases. This finding is consistent with the intuition that charlatanry becomes more extreme when the situation becomes more "hopeless". Corollary 1 also establishes an intuitive link between the phenomena of charlatanry and "guruism". A guru is a healer who attracts a small number of patients, who are willing to pay large amounts for his treatment and dismissive of the quality of alternatives. These two phenomena are distinct in principle: not every guru is necessarily a charlatan. However, they are linked by Corollary 1: as α decreases, the healers increasingly act as charlatans and increasingly become gurus.

This interpretation of the main result thus links together several features of "market for quacks". First, quacks are charlatans to some extent. Second, quacks are gurus to some extent. Thus, charlatanry and guruism are exacerbated as the underlying situation becomes more hopeless. The model traces these relations to the patients' tendency to draw sweeping conclusions from limited experience, without taking into account its random origin.

Welfare analysis. What is the interplay between the force of exploiting boundedly rational consumers and the force of competition, in terms of consumers' welfare? Let us now turn to this question. It follows directly from our previous analysis that the industry's equilibrium profits are given by the following expression:

$$\pi^*(\alpha, n(\alpha)) = n\alpha(1-\alpha)^n \tag{2}$$

Because healers are worthless, industry profits are equal to the welfare loss inflicted on the patients. Thus, Equation (2) also provides the expression for the welfare loss inflicted on patients by quacks in equilibrium.[1] The right-hand side of Equation (2) is not monotonic in α: it attains an optimum at $\alpha^* = \frac{1}{n+1}$. This value maximizes the probability of a unique success in the patient's sample. In other words, a^* maximizes the probability that a single healer will be able to exercise monopoly power over patients.

Note that the patients' welfare loss can be substantial. For every $\alpha < \frac{1}{2}$, there exists a number of healers $n \geq 2$, such that the patients' loss exceeds $\frac{1}{4}$. As $\alpha \to 0$, the maximal welfare loss converges to $\frac{1}{e}$.

The patients' welfare loss is not monotonically decreasing in n. For every α, the number of healers that maximizes the patients' welfare loss is $n^* = -\frac{1}{\ln(1-\alpha)}$. For every $\alpha \lessapprox 0.39$, $n^* \geq 2$. That is, *greater competition may actually increase the welfare loss inflicted on patients.* As $\alpha \to 0$, n^* tends to infinity, such that the perverse effect of greater competition holds for a larger domain of n. The reason for this result is simple. On one hand, a greater number of healers in the market implies a stronger incentive to reduce prices. This is the standard "competitive" effect. On the other hand, the $S(1)$ procedure implies that aggregate demand for healers is increasing in n.

[1] Equation (2) is also consistent with the case of $n = 1$, because the maximal payoff of a monopolistic healer is $\alpha(1-\alpha)$.

This is a non-standard, "exploitative" effect, which results from the patients' bounded rationality. As α decreases, it takes a larger n for the former effect to outweigh the latter.

For any fixed number of healers, the "exploitative" and "competitive" effects can be separated in a simple manner. It is easy to show that the max-min payoff in the game is equal to $\alpha(1-\alpha)^n$, which is exactly the expression for the healers' equilibrium payoffs. Thus, competition among healers implies that they earn no more than their max-min payoffs. However, these max-min payoffs are positive because patients err with positive probability. *The "exploitative effect" determines the max-min payoff, and the "competitive effect" does not allow healers to earn more than their max-min payoffs.*

Relaxing quackery. The basic model assumes $\alpha_0 = \alpha$. When $\alpha > \alpha_0$, the healers are not quacks: they have genuine healing powers relative to the default. If patients were rational, the model would be reduced to standard Bertrand competition, such that equilibrium prices and profits would be zero, and the patients' utility would be $\alpha - \alpha_0$. By comparison, in the case of $S(1)$-patients, it can easily be shown that Proposition 1 and Corollary 1 survive this modification: *the healers' equilibrium behavior is entirely independent of the default rate of success*. Thus, the assumption that healers are quacks is irrelevant for the basic result. The reason is simple: the default rate enters the expression for a healer's equilibrium payoff through the multiplicative term $1 - \alpha_0$, hence it does not affect the healer's strategic calculations.

Charlatanry persists to some extent in the modified model. In equilibrium, healers charge prices above $\alpha - \alpha_0$ with positive probability. Moreover, if we fix $\alpha - \alpha_0$ and let α_0 be sufficiently small, then $E(p) > \alpha - \alpha_0$, such that patients who end up choosing a healer (their fraction in the population is $(1 - \alpha_0) \cdot [1 - (1 - \alpha)^n]$) are worse off on average than patients who end up choosing the default.

In summary, the main equilibrium effects are independent of the assumption that healers are quacks. However, the case of $\alpha_0 = \alpha$ is of particular interest because it gives rise to an example for a market that would be inactive if consumers were rational, and owes its existence to their bounded rationality.

The patients' knowledge of the default. The basic model assumes that the patients' choice procedure treats the default and the healers symmetrically: it samples each of them once. It could be argued that patients are more familiar with the default than with the healers. In particular, they may actually know the value of α_0, such that $x_0 = \alpha_0$ with probability one. The essential features of our equilibrium characterization - uniqueness, symmetry, price dispersion and the qualitative comparative statics - will remain unchanged. Only fine details have to be modified: the "monopoly price" becomes $1 - \alpha_0$ instead of 1, the exact expression for G is slightly different; and the welfare analysis needs to be refined.

The patients' valuations. The results of this section are easily extendible to the case in which patients' willingness to pay for sure recovery is distributed according to some continuous *c.d.f* F. Suppose that F satisfies standard properties, which guarantee a unique solution p^* to the healer's profit maximization problem when $n = 1$ and $\alpha < 1$. In this case, equilibrium characterization is similar to Proposition 1. There is a unique Nash equilibrium, which is symmetric and mixed. Each firm's strategy is defined by the same function G as in Proposition 1, except that in the right-hand side of equation (1), p is replaced by $\frac{p[1-F(p)]}{p^*[1-F(p^*)]}$. (The support is $[p^L, p^*]$, where p^L is the price $p < p^*$ that solves the equation $\frac{p[1-F(p)]}{p^*[1-F(p^*)]} = (1-\alpha)^{n-1}$.)

4 An Expert Competing against the Quacks

Our equilibrium analysis established that when patients choose according to the $S(1)$ procedure, there will be an active market for quacks. Although healers have no advantage over the default, they attract a sizeable clientele in equilibrium. One might expect that if we switched one of the healers with a "real expert", the latter would crowd the quacks out of the market.

Let us examine this intuition in the present model. Modify the basic model of Section 2 by switching the success rate of a single healer, denoted e, from α to some $\alpha_e \in (\alpha, 1]$. Apart from this modification, the model remains intact. In particular, every healer $i \neq e$ has a success rate α $(= \alpha_0)$. That is, healer e is an "expert" while his opponents are "quacks".

Proposition 2 *There is a unique Nash equilibrium in the game. Every healer $i \neq e$ plays the mixed strategy given by Equation (1), has the same clientele size, and earns the same profits as in the Nash equilibrium of the basic model. (The proof appears in the appendix.)*

Thus, when patients choose according to the $S(1)$ procedure, turning a quack into an expert leaves the other healers' equilibrium behavior and performance unaffected. The modification causes patients to switch from the default to the expert, but the expert does not "steal" clients from the quacks.

The reason for this irrelevance result is as follows. Equilibrium strategies are mixed. Symmetry considerations imply that quacks play identical strategies. The expert's equilibrium equation is completely independent of α_e: it is only expressed in terms of the opponents' success rates and pricing strategies, and it is identical to the equilibrium equation of the basic model. This equation yields the quacks' pricing strategy G, which is therefore the same as in the basic model. Continuity considerations imply that the lowest price in the market continues to be $(1-\alpha)^n$. But the profit that a quack makes when he charges this price is also independent of α_e, hence the quacks' equilibrium payoffs are the same as in the basic model.

As to the expert's equilibrium behavior, it can be shown, as a corollary of Proposition 2, that $G_e = \frac{\alpha}{\alpha_e} \cdot G(p)$ for $p \in ((1-\alpha)^n, 1)$, and that G_e contains an atom of measure $1 - \frac{\alpha}{\alpha_e}$ on $p = 1$. A simple calculation shows that a patient who ends us choosing the expert is better off than a patient who ends up choosing a quack. However, both are worse off than a patient who ends up choosing the default. Thus, the expert exploits the patients' bounded rationality, although to a lesser extent than the quacks.

The lesson of this section is that contrary to the case of rational patients, in a market with $S(1)$-patients, it takes more than a single high-quality healer to hurt the market performance of low-quality healers. In order for this to happen, there must be multiple high-quality healers present in the market. To take an extreme case, suppose that we raise the success rate of exactly two healers from α to $\alpha_e = 1$. Then, in equilibrium all healers charge $p = 0$ and make zero profits. However, even in this extreme case, quacks are not absolutely crowded out because they are chosen by a positive fraction of the patients.

5 Disclosure of Success Rates

In the basic model, patients assess the healers' quality according to the $S(1)$ procedure, and healers have no control over the patients' knowledge. In this section, I assume that a healer is able to reveal his type to patients, by credibly providing them with his success rate. If he does not reveal his type, patients continue to assess his quality according to the $S(1)$ procedure. Patients infer nothing from the healer's revelation decision itself.

In this context, it would be more interesting if we allowed market primitives to be more general than in Section 2. First, denote the rate of recovery associated with alternative i by α_i, and allow the α_i's to vary across alternatives, such that $\alpha_i < 1$ for every $i = 0, 1, ..., n$. Second, suppose that the patients' valuations of sure recovery are distributed according to some *c.d.f* F over the interval $[0, 1]$, where F need not be continuous.

Modify the basic model of Section 2 in the following way. A strategy for healer i is a pair (p_i, r_i), where $r_i = Y$ (N) if the healer reveals (does not reveal). As in the basic model, let x_i denote the patient's experience with healer i. When $r_i = Y$, $x_i = \alpha_i$ with probability one. When $r_i = N$, $x_i = 1$ with probability α_i and $x_i = 0$ with probability $1 - \alpha_i$. As before, the patient chooses the alternative that maximizes $x_i - p_i$ in his sample. (Assume the same tie-breaking rule.) Note that the patients' procedure does not distinguish between the informational content of a single observation and an exact knowledge of the underlying probability. This is an extreme case of a "law of small numbers".

In a standard model with Bayesian rational patients, it would be standard to assume that the patients know the success-rate distribution among healers, but do not know ex-ante the exact assignment of success rates to healers. In this case, a high-quality healer would obviously want to reveal himself as such. In fact, every healer with $\alpha_i > \alpha_0$ (except perhaps the lowest-quality healer in this group) would reveal his success rate in equilibrium. The present model predicts different behavior:

Proposition 3 *Any strategy (p, Y) for healer i is weakly dominated by some other strategy (p', N).*

Proof. Denote $\alpha_i = \alpha$, for notational convenience. When healer i takes the strategy (p, Y), every patient whose valuation is below $\frac{p}{\alpha}$ will not choose him. Therefore, healer i's payoff from the strategy (p, Y) is bounded from above by:

$$p \cdot [1 - F(\frac{p}{\alpha})] \cdot \Pi_{j \neq i} \Pr(x_j - p_j \leq \alpha - p)$$

In contrast, when healer i takes the strategy (p', N), his payoff is bounded from below by:

$$p' \cdot \alpha \cdot [1 - F(p')] \cdot \Pi_{j \neq i} \Pr(x_j - p_j < 1 - p')$$

Now, let us show that (p, Y) is weakly dominated by (p', N), where $p' = \frac{p}{\alpha}$. Clearly, $p \leq \alpha$ - otherwise, the strategy (p, Y) would yield zero payoffs for the healer, whereas he can guarantee a positive payoff. Therefore, $p' \in (p, 1]$. The following equality follows:

$$p \cdot [1 - F(\frac{p}{\alpha})] = p' \cdot \alpha \cdot [1 - F(p')]$$

Since $\alpha - p = \alpha \cdot (1 - p')$, it is clear that $1 - p' > \alpha - p$ as long as $p < \alpha$. Therefore:

$$\Pi_{j \neq i} \Pr(x_j - p_j < 1 - p') \geq \Pi_{j \neq i} \Pr(x_j - p_j \leq \alpha - p) \tag{3}$$

The inequality is strict if $G_j(1 - p') > G_j(\alpha - p)$ for at least one healer $j \neq i$ (where G_j is the *c.d.f* induced by healer j's strategy). It follows that (p', N) weakly dominates (p, Y). ■

Thus, given the patients' choice procedure, healers have an incentive not to reveal their success rate, even when they are the highest-quality healers in the market. The decision whether to reveal one's type entails a trade-off. On one hand, the maximal price that healer i can charge from a patient is reduced from 1 to α_i. On the other hand, a fraction $1 - \alpha_i$ of patients are unwilling to choose healer i when $r_i = N$, and they become willing in principle to acquire his treatment with $r_i = Y$, depending on his price. The former consideration turns out to outweigh the latter.

Proposition 3 establishes that type revelation is a weakly dominated strategy. The following result demonstrates that it can never be part of Nash equilibrium.

Proposition 4 *In Nash equilibrium, every healer chooses $r_i = N$.*

Proof. First, note that every healer is able to maker positive payoffs in the game, by choosing $r_i = N$ and a sufficiently low price. Let us now show that in Nash equilibrium, at most one healer chooses $r_i = Y$. Assume the contrary - i.e., that $r_i = r_j = Y$ for two healers i, j. Then, patients choose between i and j just as in standard Bertrand competition, such that healer i will make zero payoffs if $\alpha_i - p_i < \alpha_j - p_j$. Therefore, one of these healers will make zero payoffs in equilibrium, a contradiction. Now suppose w.l.o.g. that $r_1 = Y$ and $r_j = N$ for every $j > 1$. In the relevant domain, there is at least one healer $j > 1$ who plays a strictly increasing G_j. (Otherwise, there is a healer who can deviate profitably by shifting weight upwards.) Therefore, $\Pr(x_i - p_i \leq z)$ is *strictly* increasing in z in the relevant domain. It follows that if healer 1 deviates from the strategy (p_1, Y) to the strategy (p'_1, N) which is constructed in the proof of Proposition 3, and Inequality (3) will be strict, hence the deviation will be profitable. ■

The lesson from Propositions 3 and 4 is that if patients do not assign greater informational content to a healer's entire statistics - or to his mere decision to reveal his entire statistics - relative to a single random observation, then even high-quality healers are reluctant to disclose their success rate. They will prefer patients to continue forming quality judgments on the basis of their limited, random experience. Endowing healers with the technology to reveal their type credibly does not affect their equilibrium behavior.

6 The Generalized $S(K)$ Procedure

The $S(1)$ procedure captures an extreme case of a "law of small numbers": patients form *deterministic* action-consequence correspondences that are justified by a *single* observation per alternative. A natural generalization of this procedure, suggested by Osborne and Rubinstein (1998), is to assume that patients sample every alternative K times and *maximize their expected payoff against the empirical distribution* generated by their sample.

According to this generalized procedure, called $S(K)$, each patient forms the following point estimate of alternative i's success rate:

$$a_i = \frac{\sum_{k=1}^{K} x_i^k}{K}$$

where $x_i^k = 1$ (0) if the outcome of the patient's k-th draw of healer i is "recovery" ("no recovery"). All the x_i^k's are independently drawn. The patient then chooses an alternative that maximizes $a_i v - p_i$, where the v denotes his willingness to pay for sure recovery. Assume that v is distributed over the interval $[0, 1]$ according to a *c.d.f* F, which satisfies standard properties that guarantee a unique monopoly price p^*. (Recall that p^* is also the price that a monopolistic healer would charge when $K = 1$ - see Section 3.)

The generalized choice procedure retains the idea that patients draw sweeping statistical inferences from a small sample, as if it fully represented the true distribution from which it was drawn. Patients form an unbiased "point estimate" of the rate of recovery associated with each alternative, but they neglect the sampling error and behave as if they can form an arbitrarily narrow confidence interval around their point estimate. As K gets larger, the sampling error decreases, and in the $K \to \infty$ limit, the patient's procedure converges to standard Bayesian rational choice. Thus, one merit of the generalized $S(K)$ procedure is that it parameterizes the extent to which our boundedly rational patient departs from Bayesian rationality, while remaining consistent with the "representativeness" heuristic.

At this stage, I am unable to provide a full characterization of equilibria in the price-competition game under the generalized $S(K)$ procedure. In this section, I will settle for an existence result and a pair of asymptotic characterizations.

Proposition 5 *If F is continuous, then the price-competition game with $S(K)$-patients has a Nash equilibrium.*

Proof. In order to apply an existence theorem due to Simon (1987), it is sufficient to show that healer i's utility function is discontinuous only when $p_i = p_j$ for some healer $j \neq i$. Suppose that $p_i \neq p_j$. We need to show that the set of patients who are indifferent between healers i and j is of measure zero. In order for a patient with valuation v to be indifferent between them, he must satisfy: $a_i v - p_i = a_j v - p_j$, where a_i and a_j are multiples of $\frac{1}{K}$. Because $p_i \neq p_j$, $a_i \neq a_j$ as well. Then:

$$v = \frac{p_2 - p_1}{a_2 - a_1}$$

and since F is continuous, the set of patients with this valuation is of measure zero. ■

Continuity of F is crucial for general existence in the generalized model, contrary to the case of $K = 1$. To illustrate the problem, suppose that all patients have $v = 1$, and consider the case of $K = 2$. A fraction $2\alpha(1-\alpha)$ of the patients have $a_i = \frac{1}{2}$. These patients are in principle willing to choose healer i when $p_i < \frac{1}{2}$, because in that case there is a positive probability that $\frac{1}{2} - p_i > a_j - p_j$ for every $j \neq i$. Healer i loses these patients automatically when he charges a price higher than $\frac{1}{2}$. Therefore, his utility function is discontinuous at $p_i = \frac{1}{2}$, irrespective of the opponents' behavior. This kind of discontinuity may be ruinous for existence. However, for arbitrary valuation distributions, Nash equilibrium exists if α is sufficiently small.

Let us turn to asymptotic result. The first result provides a closed characterization of equilibrium behavior in the low-α region.

Proposition 6 *As $\alpha \rightarrow 0$, expected equilibrium prices converge to $\frac{p^*}{K}$.*

Proof. Consider healer 1's decision. If he had no competitors, what would be the optimal price in the face of $S(K)$-patients? A patient whose valuation of sure recovery is v is willing to pay $a_i v$ to healer i. When α is close to zero, $\Pr(a_i = \frac{1}{K}) \gg \Pr(a_i > \frac{1}{K})$. Therefore, a monopolistic healer would target the patients for whom $a_i = \frac{1}{K}$. The optimal price for these patients is $\frac{p^*}{K}$. By continuity, as α tends to zero, the monopolistic healer's optimal price converges to $\frac{p^*}{K}$. Now introduce competition. When α tends to zero, the probability that $a_j > 0$ for some $j \neq i$ is negligible. Therefore, the set of prices that maximize healer i's expected payoffs given s_{-i} must be concentrated in an arbitrarily small neighborhood of $\frac{p^*}{K}$. ■

Proposition 6 shows that in the low-α limit, there is an inversely proportional relation between equilibrium prices and the procedural parameter K. When α is small, the vast majority of patients have a totally unsuccessful experience with healers. As to the rest of the patients, a vast majority of them have only a single success in their sample. Thus, the clientele targeted by healer i consists of patients with $a_i = \frac{1}{K}$ and $a_j = 0$ for every $j \neq i$. Thus, market equilibrium in the low-α limit is characterized by monopolistic competition with respect to the service "recovering with probability $\frac{1}{K}$" - just as in the basic model, market equilibrium in the low-α limit is characterized by monopolistic competition with respect to the service "recovering with probability one". Note that in the low-α limit, not only the market price, but also the patients' welfare loss are inversely proportional to K.

The next result establishes that for every α, as K tends to infinity, the patients' equilibrium payoffs converge to their Bertrand payoffs. Thus, the $S(K)$ procedure itself does not cause a discontinuity of the Nash correspondence, relative to the rational benchmark.

Proposition 7 *For every $\alpha \in (0, 1)$, the patients' expected equilibrium payoffs converge to α, as $K \rightarrow \infty$.*

Proof. Assume the contrary. Then, there must exist a price $p^* > 0$, such that for any arbitrarily high K, the fraction of patients who choose a healer that charges a price above p^* is bounded away from zero. But according to the law of large numbers, if $p_i > 0$, the probability that a patient chooses healer i over the default converges to zero as $K \to \infty$, a contradiction. ■

Throughout this paper, we assumed that the healers offer flat-price contracts, such that payments are independent of the outcome of the healer's treatment. In the $K = 1$ case, this distinction is irrelevant, but it would make a difference when $K > 1$. I leave the analysis of more complex price schedules to future research.

7 Conclusion

In this paper, we considered markets for goods and services that are designed to fix a consumer's "problem", when the probability that the problem will be fixed is strictly between zero and one, regardless of the consumer's decision. In the context of such markets, we examined what happens to the familiar price competition model when rational consumers are replaced with boundedly rational consumers, who obey the $S(K)$ procedure due to Osborne and Rubinstein (1998). This procedure captured a salient feature of the Tversky-Kahneman "representativeness" heuristic, namely people's tendency to interpret a small sample as if it had the exact shape of the distribution from which it was drawn. This aspect of bounded rationality is particularly relevant for industries such as unconventional medicine or self-help, in which consumers lack learning opportunities that might correct their imperfect judgments.

We focused on an extreme case, in which firms are "quacks", in the sense that they have no advantage over the default. In the case of $K = 1$, Nash equilibrium analysis yielded a novel I.O. characterization. As the rate of recovery decreases - i.e., as the situation becomes more "hopeless" - aggregate demand decreases and expected price increases, such that in the limit, the market converges to a state of monopolistic competition. The interpretation of the comparative statics revealed a connection between the phenomena of charlatanry and guruism, and how they respond to the underlying rate of recovery. Welfare analysis established that a greater number of healers in the market can actually increase the welfare loss inflicted on patients. The general $K > 1$ case confirmed the intuition that greater rationality on the consumers' part mitigates the adverse effect of quacks' behavior.

I wish to highlight the simplicity and parsimony of the model. A major obstacle to the construction of viable economic models of bounded rationality is that they often get quite complex even in the simplest scenarios. The present model yields a novel I.O. characterization using elementary tools. This is an encouraging sign, as we move on to more sophisticated models of rational competition over boundedly rational consumers.

Charlatanry in the face of rational patients. In real-life markets, charlatans may earn stable economic rents for two possible reasons. First, consumers may be imperfectly informed about their quality. Second, consumers may draw imperfect statistical inferences. This paper focused on the latter reason. We modeled charlatans as agents who seek economic rents by systematically exploiting consumers' imperfect statistical inferences about the quality of their treatments. Specifically, we traced charlatanry to the behavioral assumption that consumers draw sweeping inferences from limited experience, ignoring the fact that their experience was randomly generated.

To what extent do the results rely on the view of charlatans as exploiters of consumers' imperfect *rationality*, as opposed to the view of charlatans as exploiters of consumers' imperfect *information*? Would it be possible to attain similar results using a standard model with incompletely informed, Bayesian consumers? Clearly, the answer is negative in the basic model, in which all healers are quacks, as long as we maintain the standard assumption that consumers hold correct prior beliefs. If consumers know the quality distribution, they know that all healers are quacks, hence they would choose to stay out of the market. The classical adverse-selection story is relevant only when the industry contains both high-quality and low-quality healers, and the consumers cannot tell which is which. By contrast, the bounded-rationality story is relevant even when healers are homogeneous, because the consumers' inferences may lead to a spurious perception of heterogeneity.

In order for a model with Bayesian patients to replicate the results of this paper, it is therefore necessary to relax the correct-prior assumption (a non-standard modeling decision by itself) or the homogeneity assumption. Even so, I conjecture that such a model would not be able to replicate the main results of this paper, such as unequivocal negative relation between success rates and equilibrium prices, the welfare implications of greater competition, and particularly the impossibility of type revelation. These aspects seem to distinguish the present modeling approach from standard adverse-selection models (e.g., Leland (1979), Wolinsky (1983)), in which consumers rationally infer quality from prices.

What happens when we perturb the basic model, and mix $S(1)$-patients with standard rational patients? Denote the fraction of rational patients by ε. It can be shown that the equilibrium correspondence is continuous in ε. For every $\varepsilon \in (0,1)$ there is a unique equilibrium, which is mixed and symmetric. The patients' equilibrium payoffs increase in ε. That is, rational patients exert positive externalities on boundedly rational patients. As $\varepsilon \to 1$, equilibrium converges to the rational-patients benchmark. Since this result is not surprising and involves the same kind of arguments as the basic model, I omit the details.

Related literature. Slovic, Tversky and Kahneman (1982) contains an excellent collection of psychological studies into the representative heuristic, the law of small numbers and other facets of non-Bayesian inference. In recent years there has been a number of attempts to introduce elements of non-Bayesian statistical inference into economic modeling. Eyster and Rabin (2000) analyze players in a Bayesian game, who optimize against the opponents' statistical distribution of actions, rather than against their Bayesian-game strategies. In Jehiel (2001), players in an extensive game optimize against the statistical behavior of opponents across different histories that belong to the same "analogy class". Rabin (2002) proposes an alternative modeling approach to the "law of small numbers" fallacy. He studies a single decision maker, who make predictions about the evolution of an *i.i.d* stochastic process, as if his observations are drawn from an urn without replacement.

There are few other studies into the $S(K)$ procedure itself. Osborne and Rubinstein (2002) analyze voting games with players who use it. It should be stressed that certain conceptual problems that beset the models of Osborne and Rubinstein (1998,2002) are irrelevant in the present context. In these models, *every* agent in the environment chooses according to the $S(K)$ procedure, which therefore becomes the cornerstone of a novel equilibrium concept. The equilibrating processes that may justify this concept are yet to be explored. (See Sethi (2000) and Miękisz and Ramsa (2003) for studies in this direction.) In the present paper, $S(K)$ is the patients' choice procedure, and there is no strategic interaction among them. The equilibrium concept that is applied to the healers' behavior is Nash equilibrium. In this respect, the bounded-rationality element in the present model does not raise new interpretation problems concerning equilibrium behavior.

For earlier inquiries into the general subject of industrial organization with boundedly rational consumers, see Rubinstein (1993) and Piccione and Rubinstein (2003), who analyze monopolistic behavior when consumers differ in their ability to understand complex pricing schedules. For another attempt to incorporate elements of bounded rationality into the I.O. literature, see Fershtman and Kalai (1993), who study competitive behavior of firms which operate in multiple markets and whose ability to implement complex competition strategies in these markets is bounded.

Finally, the basic model of Section 2 is obliquely related to the search literature, in particular the classic model of Burdett and Judd (1983), which explains equilibrium price dispersion in a market for a homogenous good as a consequence of costly simultaneous search by consumers. In their model, which is not explicitly game-theoretic, the state of a market is characterized by a distribution of prices over a continuum of firms. Consumers have access to a search technology: they can sample a number of firms, for a fixed cost c per sample point. A market equilibrium is a price distribution, such that every price in the support yields the same profit for firms, and every consumer optimally chooses his sample size, given the price distribution. Burdett-Judd show that in equilibrium, some consumers sample one firm, while the others sample two firms. Competitive pressures are curbed because of the former group of consumers, but they are not totally absent thanks to the latter group. The price distribution's support stretches all the way up to the monopoly price, and its exact shape is determined by the value of c.

While the two models bear some resemblance to each other, the differences are substantial. Consumers have limited knowledge of prices in the Burdett-Judd model, whereas in the present paper, they know the prices and have limited knowledge of healers' quality. In the Burdett-Judd model, consumers choose their sample size rationally, given their knowledge of the price distribution. In the present model, patients do not choose their sample size and they respond naïvely to the realization of their sample. The crucial parameter in the Burdett-Judd model is the search cost c. In contrast, the crucial parameter in the basic model of Section 2 is the rate of recovery α, which determines the structure of the the patients' demand for healers, via their $S(1)$ procedure. Most importantly, the comparative statics with respect to α are totally unrelated to the Burdett-Judd model.

8 Appendix

8.1 Proof of Proposition 1

Healer i's strategy s_i induces a *c.d.f* G_i over the interval $[0, 1]$. First, let us show that G_i is continuous over $[0, 1)$. Since G_i is monotonic, it is sufficient to show that s_i contains no atoms on $[0, 1)$. Assume the contrary and suppose that s_i contains an atom on some $p < 1$. If $p = 0$, then healer i assigns a positive measure to a price that yields zero profits. The patients' choice procedure guarantees that at least a fraction $\alpha(1-\alpha)^n$ of the patients will choose healer i. Hence, he can profitably deviate by shifting this measure to $p > 0$. Suppose that $p \in (0, 1)$. There are two cases. First, every other healer may assign measure zero to the interval $(p, p+\varepsilon)$, for some arbitrarily small ε. In this case, healer i can profitably deviate by shifting the atom from p to $p+\frac{\varepsilon}{2}$. Second, for every $\varepsilon > 0$, there may be a healer j who assigns a positive measure to the interval $(p, p+\varepsilon)$. In this case, healer j can profitably deviate by shifting this measure to some $p' < p$ arbitrarily close to p. Thus, we have established that neither healer's strategy assigns an atom to a price $p < 1$. Note that if s_i contains an atom on $p = 1$, then every other healer's strategy does not. Otherwise, one of these healers can profitably deviate by shifting this atom to some smaller p' arbitrarily close to one.

Throughout the rest of the proof, we will use a standard result. If s_i assigns a positive measure to an interval $(p, p+\varepsilon)$ or $(p, p-\varepsilon)$ for every $\varepsilon > 0$, then by a standard continuity argument, p must maximize healer i's expected payoff against s_{-i}.

Define $p_i^L = \sup\{p \in [0, 1]; G_i(p) = 0\}$. Define $p^H = \inf\{p \in [0, 1]; G_i(p) = 1\}$. Let $p^L = \min\{p_1^L, ..., p_n^L\}$ and $p^H = \max\{p_1^H, ..., p_n^H\}$. Our task now is to characterize p^L and p^H. Suppose that $p^H < 1$. The only patients who choose healer i given $p_i = p^H$ are those whose sample has $x_i = 1$ and $x_j = 0$ for every $j \neq i$. Faced with these patients, any price less than one is sub-optimal. Therefore, p^H does not maximize healer i's expected payoffs, a contradiction.

Note that p^L and p^H must yield the same expected payoff for the healers who charge these prices. The reason is as follows. Regardless of the opponents' strategies, the price $p^H = 1$ yields a sure payoff of $\alpha(1-\alpha)$, and the price p^L yields a sure payoff of $p^L \cdot \alpha$. If these payoffs are not equal, then a player who charges the less profitable of these prices can deviate to the more profitable price. It follows that both p^L and $p^H = 1$ yield a payoff of $\alpha(1-\alpha)^n$. Therefore, $p^L = (1-\alpha)^n$.

Let us now show that $p_i^L = p^L$ for every $i = 1, ..., n$. Assume the contrary and suppose that $p_i^L > p^L$ for some healer i. We already established that healer i's payoff must be at least $\alpha(1-\alpha)^n$. Suppose that $p_j^L = p^L$ and that healer i's payoff is strictly larger than $\alpha(1-\alpha)^n$. Recall that player j's payoff is exactly $\alpha(1-\alpha)^n$. Consider the following deviation for healer j, from s_j to the pure strategy $p_j = p_i^L$. this deviation must be profitable, for the following reason. For notational convenience, let $i = 1$ and $j = 2$. Compare healer 1's competitive situation before player 2's deviation and healer 2's competitive situation after the deviation. In the former case, healer 1 faces competition from the *c.d.f*'s $G_2, G_3, ..., G_n$. In the latter case, healer 2 faces competition from the *c.d.f*'s $G_3, ..., G_n$. He does not face competition from G_1 because healer 1's strategy assigns probability one to higher prices than p_1^L. Therefore, healer 2's payoffs following the deviation are at least as high as healer 1's payoffs before the deviation. Therefore, the deviation is profitable.

It follows that all healers make a payoff of $\alpha(1-\alpha)^n$ in equilibrium, because they all have $p_i^L = (1-\alpha)^n$ and because the healers' strategies do not have an atom on this price. Consider a price $p \in (p^L, 1)$, such that every healer assigns a positive measure to the neighborhood of p. Because p maximizes every healer's payoffs, given the other healers' strategy, it follows that for every $i = 1, ..., n$:

$$\alpha(1-\alpha)^n = p \cdot \alpha \cdot \Pi_{j \neq i}[1 - \alpha G_j(p)] \cdot (1-\alpha) \tag{4}$$

We have a set of n equations in n variables $G_j(p)$. The equations are symmetric, and the right-hand side of healer i's equation is strictly decreasing in the $G_j(p)$'s. Therefore, the solution must be unique and symmetric: $G_1(p) = \cdots = G_n(p) \equiv G(p)$ for every $p \in ((1-\alpha)^n, 1)$.

(The system of equations (4) relies on the assumption that every healer's strategy assigns a positive measure to the neighborhood of p. This is indeed the case for *every* $p \in (p^L, 1)$. In other words, none of the healers' strategies contains "holes". Otherwise, the system of equations (4) would consist of m equations, $2 \leq m < n$, and the solution would necessarily mean that some of the G_j's contain atoms at some price below one, a contradiction.)

It is now straightforward to derive Equation (1) from the system of equations given by (4). It can be verified that $G(p) \to 1$ as $p \to 1$ and $G(p) \to 0$ as $p \to (1-\alpha)^n$, such that the healers' equilibrium strategies contain no atoms. This completes the proof.

8.2 Proof of Proposition 2

Let us borrow the definitions of p_i^L, p_i^H, p^L, p^H from the proof of Proposition 1. Several steps in the proof can be borrowed as well. First, equilibrium strategies are mixed, and they contain no atoms below $p = 1$. Moreover, at most one healer's strategy has an atom on $p = 1$. Second, $p^H = 1$. Finally, using the same symmetry arguments as in the proof of Proposition 1, we obtain that all quacks play the same pricing strategy: $G_i = G_j \equiv G$ for every $i, j \neq e$. In particular, they all have the same p_i^L (roughly speaking, the "lowest price in the market") and G cannot contain an atom on $p = 1$. (On the other hand, G_e may contain of some measure A an atom on $p = 1$.

If $p_e^L < p_i^L$, then healer e can profitably deviate by shifting the measure he assigns to the interval (p_e^L, p_i^L) upwards, towards p_i^L. Therefore, $p_e^L \geq p_i^L$. Let us now show that healer e's equilibrium payoff is $\alpha_e \cdot (1-\alpha)^n$. First, suppose that $p_e^H = 1$. By continuity, $p = 1$ must maximize healer e's payoff given the opponents' strategies, such that his equilibrium payoff is $\alpha_e \cdot (1-\alpha)^n$. Second, suppose that $p_e^H < 1$. Consider healer i for whom $p_i^H = 1$. His payoff is $\alpha(1-\alpha)^{n-1}(1-\alpha_e)$. Healer e's payoff cannot be greater than $\alpha_e \cdot (1-\alpha)^n$ - otherwise, healer i would be able to mimic healer e's behavior and obtain a payoff of at least $\alpha(1-\alpha)^n$, a profitable deviation. But $\alpha_e \cdot (1-\alpha)^n$ is equal to healer e's max-min payoff. Therefore, healer e's equilibrium payoff is $\alpha_e \cdot (1-\alpha)^n$.

Select any price $p \in (p_e^L, 1)$, such that s_e assigns a positive measure to the neighborhood of p. Because p maximizes healer e's payoff given the opponents' strategies, the following equation holds:

$$\alpha_e \cdot (1-\alpha)^n = p \cdot \alpha_e \cdot \Pi_{j \neq e}[1 - \alpha G_j(p)] \cdot (1-\alpha) \tag{5}$$

But since $G_j = G$ for every $j \neq e$ and the α_e factor cancels out, we obtain the same equation for G as in the case of Proposition 1, given p.

Just as in the basic model, continuity considerations imply that every healer i's strategy assigns a positive measure to the neighborhood of any price $p \in (p_i^L, 1)$. In particular, Equation (5) holds for every $p \in (p_e^L, 1)$. If $p_e^L > p^L$, then for every $p \in (p^L, p_e^L)$, the quacks' $G_i(p)$ is given by a system of equations which is identical to (4), except that it consists of $n-1$ equations and $G_e(p) = 0$. But this implies that $G_e(p)$ is discontinuous at p_e^L. Therefore, $p_e^L = p^L$.

Since p^L yields a payoff of $\alpha_e \cdot (1-\alpha)^n$ to healer e, $p^L = (1-\alpha)^n$. It follows that the quacks' payoff is $\alpha(1-\alpha)^n$, exactly as in Proposition 1. Also, both p^L and p^H are the same as in Proposition 1, and the quacks' pricing strategy G is defined over the interval $[(1-\alpha)^n, 1)$. Since the quacks play the same strategy and earn the same profits as in Proposition 1, they must be chosen by the same fraction of the patients.

References

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

[2] Eyster E. and M. Rabin (2002), "Cursed Equilibrium", Mimeo, Berkeley University.

[3] Fershtman C. and E. Kalai (1993), "Complexity Considerations and Market behavior", *RAND Journal of Economics* **24**, 224-235.

[4] Jehiel P. (2001), "Analogy-Based Expectations Equilibrium", Mimeo, CERAS.

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

[6] Leland H. (1979), "Quacks, Lemons, and Licensing: A Theory of Minimum Quality Standards", *Journal of Political Economy* **87**, 1328-1346.

[7] Mię̨kisz J. and M. Ramsa (2003), "Dynamical Models of Sampling Procedures", Mimeo, Warsaw School of Economics.

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

[9] Osborne M. and A. Rubinstein (2002), "Sampling Equilibrium with an Application to Strategic Voting", *Games and Economic Behavior*, forthcoming.

[10] Piccione M. and A. Rubinstein (2003), "Modeling the Economic Interaction of Agents with Diverse Abilities to Recognize Equilibrium Patterns", *Journal of European Economic Association*, forthcoming.

[11] Rabin M. (1998), "Psychology and Economics", *Journal of Economic Literature* **36**, 11-46.

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

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

[14] Sethi R. (2000), "Stability of Equilibria in Games with Procedurally Rational Players", *Games and Economic Behavior* **32**, 85-104.

[15] Simon L. (1987), "Games with Discontinuous Payoffs", *Review of Economic Studies* **54**, 569-597.

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

[17] Wolinsky A. (1983), "Prices as Signals of Product Quality", *Review of Economic Studies* **50**, 647-658.

THE FOERDER INSTITUTE FOR ECONOMIC RESEARCH and THE SACKLER INSTITUTE FOR ECONOMIC STUDIES
The Eitan Berglas School of Economics
Tel-Aviv University

Recent List of Working Papers

28-2001	Ran Spiegler	Testing Threats in Repeated Games
29-2001	Bronwyn Hall Adam Jaffe Manuel Trajtenberg	The NBER Patent Citations Data File: Lessons, Insights and Methodological Tools
30-2001	Ran Spiegler	Argumentation and Decision Biases in Multi-Issue Debates
31-2001	Shabtai Donnenfeld Alfred A. Haug	Currency Invoicing of U.S. Imports
32-2001	Itay Goldstein Ady Pauzner	Demand Deposit Contracts and the Probability of Bank Runs
33-2001	Itay Goldstein Ady Pauzner	Contagion of Self-Fulfilling Financial Crises Due to Diversification of Investment Portfolios
34-2001	Yoram Weiss Robert M. Sauer Menachem Gotlibovski	Immigration, Search and Loss of Skill
1-2002	Gene M. Grossman Elhanan Helpman	Outsourcing in a Global Economy
2-2002	Hedva Ber Asher Blass Oved Yosha	Monetary Policy in an Open Economy: The Differential Impact on Exporting and Non-Exporting Firms
3-2002	Ran Spiegler	Justifiability and Delay Tactics in a Concession Game
4-2002	Chaim Fershtman Uri Gneezy Frank Verboven	Discrimination and Nepotism: The Efficiency of the Anonymity Rule
5-2002	Pierre Andre Chiappori Yoram Weiss	Marriage Contracts and Divorce : An Equilibrium Analysis
6-2002	Ariel Rubinstein	Comments on the Risk and Time Preferences in Economics
7-2002	Jean-Marie Viaene Itzhak Zilcha	Human Capital and Cross-Country Comparison of Inequality
8-2002	Chaim Fershtman Aviad Heifetz	Read My Lips, Watch for Leaps: A Theory of Endogenous Instability
9-2002	Sarit Cohen Zvi Eckstein	Labor Mobility of Immigrants: Training, Experience, Language and Opportunities

----

Institute for Economic Research, Tel-Aviv University, Tel-Aviv, 69978 Israel, Tel: 972-3-640-9255; fax: 972-3-640-5815; e-mail: foerder@post.tau.ac.il **Papers from 2000 onwards (and some from 1998 and1999) can be downloaded from our website as follows: http://econ.tau.ac.il/research/search_workingPapers.asp.**

10-2002	Michele Piccione Ariel Rubinstein	Modeling the Economic Interaction of Agents with Diverse Abilities to Recognize Equilibrium Patterns
11-2002	Dmitri Byzalov	A Dynamic Model of OPEC: Collusive Behavior and Demand-Reducing Investment
12-2002	Aviad Heifetz Ella Segev	Escalation and Delay in Long International Conflicts
13-2002	Gene M. Grossman Elhanan Helpman	Outsourcing versus FDI in Industry Equilibrium
14-2002	Gilat Levy	A Model of Political Parties
15-2002	Assaf Razin Efraim Sadka	FDI Flows, Portfolio Flows and Domestic Investment: The Role of Information
16-2002	Itay Goldstein Ady Pauzner	Demand Deposit Contracts and the Probability of Bank Runs
17-2002	Aviad Heifetz Ady Pauzner	Backward Induction with Players who Doubt Others' Faultlessness
18-2002	Allan Drazen	Fiscal Rules from a Political Economy Perspective
19-2002	Allan Drazen	Conditionality and Ownership in IMF Lending: A Political Economy Approach
20-2002	Chaim Fershtman Hans K. Hvide Yoram Weiss	A Behavioral Explanation of the Relative Performance Evaluation Puzzle
21-2002	Neil Gandal	The Effect of Native Language on Internet Usage
22-2002	Martin J. Osborne Ariel Rubinstein	Sampling Equilibrium with an Application to Strategic Voting
23-2002	Kfir Eliaz Ran Spiegler	Anticipatory Feelings and Attitudes to Information
24-2002	Burkhard C. Schipper	Imitators and Optimizers in Cournot Oligopoly
25-2002	Alex Cukierman Itay Goldstein Yossi Spiegel	The Choice of Exchange Rate Regime and Speculative Attacks
26-2002	Alex Cukierman Stefan Gerlach	The Inflation Bias Revisited: Theory and Some International Evidence
27-2002	Aviad Heifetz Chris Shannon Yossi Spiegel	What to Maximize if you Must
28-2002	Maristella Botticini Zvi Eckstein	From Farmers to Merchants: A Human Capital Interpretation of Jewish Economic History
29-2002	Zvi Hercowitz	Concentration of Capital Ownership and Investment Fluctuations
30-2002	Jeffrey Church Neil Gandal	Indirect Network Effects and Adoption Externalities

	David Krause	
1-2003	Zvi Eckstein Yoram Weiss	On the Wage Growth of Immigrants: Israel, 1990-2002
2-2003	Bernhard Eckwert Itzhak Zilcha	The Effect of Better Information on Income Inequality
3-2003	Chaim Fershtman Oren Bar-Gill	Law and Preferences
4-2003	Hans K. Hvide Yoram Weiss	Cultural Diversity, Status Concerns and the Organization of Work
5-2003	Assaf Razin Efraim Sadka	Privatizing Social Security Under Balanced-Budget Constraints: A Political-Economy Approach
6-2003	Assaf Razin Efraim Sadka	A Brazilian-Type Debt Crisis: Simple Analytics
7-2003	Yossi Hadar David Pines	Population Growth and its Distribution between Cities: Positive and Normative Aspects
8-2003	Tomer Blumkin Efraim Sadka	Income Taxation with Intergenerational Mobility: Can Higher Inequality Lead to Less Progression
9-2003	Michele Piccione Ariel Rubinstein	Two Tales of Power and Distribution of Wealth in the Jungle
10-2003	Jacob Glazer Ariel Rubinstein	A Model of Optimal Persuasion Rules
11-2003	Manuel Trajtenberg	Defense R&D in the Anti-terrorist Era
12-2003	Yossi Hadar David Pines	On the Market Failure in a Dixit-Stiglitz Setup with Two Trading Cities
13-2003	Dana Heller	Inter-Generational Risk Sharing through Social Convention
14-2003	Oved Yosha	Financing Innovation: Is Transparency a Two-Edged Sword?
15-2003	Casey B. Mulligan Yona Rubinstein	Household vs. Personal Accounts of the U.S. Labor Market, 1967-97
16-2003	Allan Drazen Nuno Limão	Government Gains from Self-Restraint: A Bargaining Theory of Inefficient Redistribution Policies
17-2003	Aviad Heifetz Martin Meier Burkhard C. Schipper	Interactive Unawareness and Speculative Trade
18-2003	Michele Piccione Ariel Rubinstein	Equilibrium in the Jungle
19-2003	Michele Piccione Ariel Rubinstein	The Curse of Wealth and Power
20-2003	Bernhard Eckwert Itzhak Zilcha	Improvement in Information, Income Inequality and Growth
21-2003	Ran Spiegler	The Market for Quacks