THE PINHAS SAPIR CENTER FOR DEVELOPMENT
TEL AVIV UNIVERSITY

The Market for Quacks*

Ran Spiegler[1]

Discussion Paper No. 6-2004

June 2004

∗ I would like to thank Eddie Dekel, Kfir Eliaz, Volker Nocke, Michele Piccione, Andrea Prat, Ariel Rubinstein and Asher Wolinsky, as well as seminar audiences at Cambridge, Essex, LSE and UCL, for helpful comments

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

Abstract

A group of n "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, they would choose to default and the market would be inactive. I assume, however, that patients choose according to a boundedly rational procedure due to Osborne and Rubinstein (1998), which captures people's tendency to make exaggerated inferences about treatments' quality from very limited data. This element of bounded rationality has significant implications. The market for quacks is active. Quacks inflict a welfare loss on patients, which is not monotonic in n and α. Given the patients' model of behavior, certain market interventions which might seem to alleviate the adverse welfare effects (replacing a quack with a "real expert", enabling quacks and real experts to credibly their quality) end up leaving the phenomenon unaffected.

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 that the systematic departures from Bayesian inference documented by psychologists 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 might affect our analysis of 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 "industrial organization" 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 a simple market model. A group of n identical firms (referred to as "*healers*"), faces a continuum of identical consumers (referred to as "*patients*"), and plays a standard price-competition game. Every patient enters the market with some problem. If he acquires the "*treatment*" offered by one of the healers, he recovers with probability $\alpha \in (0,1)$. If the patient chooses to acquire none of the treatments offered in the market, he recovers with probability $\alpha_0 \in (0,1)$. The patient's utility is equal to his recovery rate minus the price that he pays.

I will focus on a special case, in which *the patients' rate of recovery is independent of their decision*, i.e., $\alpha_0 = \alpha$. In this case, I 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. Indeed, as this image and our terminology suggest, the unconventional medicine industry will serve as a running example in this paper, because it seems to provide a natural context for discussing the phenomenon of quackery. Of course, I do not wish to impugn all practitioners of unconventional medicine as quacks. Furthermore, even in the most blatant cases of quack medicine, placebo effects endow practitioners with some advantage over the default. I focus on the ideal case of $\alpha_0 = \alpha$ for expositional purposes. As we shall see, extending the analysis to the case of $\alpha \neq \alpha_0$ is straightforward.

If patients were rational and endowed with a correct understanding of the market model, they would regard the entire industry as providing a worthless treatment, and 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 samples (once) each of the $n + 1$ alternatives. A patient's sample assigns some outcome $x_i \in \{0, 1\}$ to alternative i, where $x_i = 1$ (0) means that the outcome was a success (failure). The patient chooses the alternative i that maximizes $x_i - p_i$ in his sample. The outcome of his decision is a new, independent draw, such that his true expected payoff from his decision is $\alpha - p_i$ (and not $x_i - p_i$).

The patients' choice procedure is borrowed from Osborne and Rubinstein (1998), who called it $S(1)$ and studied normal-form games with players who choose their strategy according to this procedure. In the present context, the $S(1)$ procedure serves as a model of the element of bounded rationality that I wish to focus on, namely consumers' tendency to draw sweeping inferences about the quality of treatments from limited data, ignoring the randomness of their data. The sampling should not be taken literally. *This is not a model of how patients actively gather data, but rather a model of how they draw inferences from casual observations.* Patients in this model do not perceive their stochastic environment by forming probabilistic beliefs. Rather, they perceive it through *anecdotes*. The anecdotes are randomly generated, either from the patient's own experience or from hearsay.

This model of behavior fits situations in which consumers have poor understanding of the market. The most basic assumption in standard economic models is that the model is common knowledge among all agents. Different agents may have different information, but they all share the same understanding of how the economy works. In contrast, the present model assumes that the consumers' understanding of the market falls *below* the level of being able to form a probabilistic prior. Instead, they perceive their environment through casually collected anecdotes, and they commit the fallacy of attaching too much significance to their anecdotes.

(In Section 8, I discuss whether one could replicate the results of the model, by replacing the assumption that patients and healers differ in their understanding of the model with the more standard assumption of common knowledge of the model combined with asymmetric information.)

Tversky and Kahneman (1974) used the term "the law of small numbers" to describe this fallacious tendency to over-infer from small samples, and provided experimental evidence for this "law". They explained the fallacy as a consequence of the "*representativeness*" heuristic. (See Part II in Kahneman, Slovic and Tversky (1982).) 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. The $S(1)$ procedure captures an extreme version of "the law of small numbers", and it is consistent with the representativeness-based explanation. Patients in our model maximize utility against the empirical distribution of recoveries given by their sample, as if this were the true distribution. Patients' quality judgments end up being insensitive to the prior recovery rate α and to their sample size.

I believe that this model - namely, the reliance on anecdotes instead of probabilities when forming quality judgments, as well as the tendency to over-infer from the anecdotes - is fairly descriptive of the way consumers in the unconventional medicine industry form quality judgments. Choosing an unconventional treatment for an illness (when standard treatments are relatively ineffective) is a choice problem 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 their experience candidly, or because personal circumstances are too idiosyncratic. In such circumstances, it is particularly natural for consumers to form quality judgments of unconventional treatments on the basis of anecdotes, and the consumers' fallacious tendency to over-infer from anecdotes becomes particularly relevant.

Let us turn to a description of the results. The price-competition game played by the healers has a unique Nash equilibrium, which is symmetric and mixed. The equilibrium strategy is given by a simple formula. For every α, the "market for quacks" is active. Quacks act as "*charlatans*": they charge positive prices for worthless treatments. Expected equilibrium price is decreasing in α; it tends to the competitive price $p = 0$ as $\alpha \to 1$, and to the monopoly price $p = 1$ as $\alpha \to 0$. 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. The characterization is reminiscent of the literature on equilibrium price dispersion, e.g. Butters

(1977), Varian (1980), Burdett and Judd (1983) and Rob (1985), although technically, it goes beyond these papers in establishing the symmetric equilibrium as the *unique* Nash equilibrium.

Activity in the market for quacks inflicts a *welfare loss* on patients: those who end up acquiring the quacks' treatments are worse off in expectation than those who end up choosing the default. I define the patients' welfare loss as the difference between their equilibrium expected payoff and their expected payoff from the default. The welfare loss is given by the expression $n\alpha(1-\alpha)^n$, which does *not* behave monotonically in α and n. It follows that *greater competition (in the sense of larger n) may 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 force generated by a larger number of competitors. The patients' welfare loss can be substantial: for every $\alpha < \frac{1}{2}$, there exists $n \geq 2$ such that the loss exceeds $\frac{1}{4}$.

The rest of the paper extends the basic model in various directions. In each extension, I introduce an intervention that might seem a priori to curb the quacks' adverse welfare effects. Indeed, in standard models with rational consumers, such interventions typically improve consumer welfare. First, I endow healers with the ability to disclose costlessly and credibly their success rate. (In this extension, I allow success rates to vary across healers.) In a standard adverse selection model, in which Bayesian rational patients are imperfectly informed of the healers' quality, this intervention would result in a complete crowding out of low-quality healers. In contrast, given our model of the patients' behavior, *disclosure of success rates is a dominated strategy*, even for high-quality healers.

In the second extension, I raise the recovery rate of one healer, turning him from a "quack" into a "real expert". The other $n-1$ healers remain quacks. In a standard model with (perfectly informed) rational patients, this intervention may help crowding out the low-quality healers. However, in the present model, the intervention has no impact whatsoever on the quacks' equilibrium strategy and payoff.

The third extension of the model generalizes the patients' choice procedure. Each patient samples each of the $n+1$ alternatives K times. Let $a_i \in \{0, \frac{1}{K}, ..., 1\}$ denote the empirical success rate of alternative i in the patient's sample. The patient chooses the alternative i that maximizes $a_i - p_i$ in his sample. Osborne and Rubinstein (1998) called this extended procedure $S(K)$. The parameter K indicates the extent of the patients' departure from the rational benchmark. As K increases, our $S(K)$-patient and a standard rational patient become more likely to reach the same decisions. I provide a pair of asymptotic results, which illustrate the effect of K on the competitiveness of market equilibrium. As $\alpha \to 0$, expected equilibrium price converges to $\frac{1}{K}$. For every α, expected equilibrium price converges to zero as $K \to \infty$.

The lesson from the extensions of the basic model is that ordinary pro-competition market interventions - which would crowd out low-quality firms and increase consumer welfare in standard market models with Bayesian rational consumers - have very different implications when consumers choose according to the $S(1)$ procedure. As long as we retain the consumers' boundedly rational statistical inferences, such interventions do not make the market more competitive.

The paper proceeds as follows. Section 2 presents the basic model. Section 3 analyzes market equilibrium. Section 4 analyzes disclosure of success rates. Section 5 perturbs the basic model by replacing a quack with a "real expert". Section 6 analyzes the generalized $S(K)$ model. Section 7 discusses the effect of the $S(1)$ procedure in markets for truly differentiated goods. Section 8 discusses the model's interpretation - in particular, the extent to which it can be "rationalized". Section 9 discusses related literature. Some proofs are relegated to the appendix.

2 A Basic Model

A two-sided market consists of a continuum of measure one of identical consumers ("*patients*") on one side and n identical firms ("*healers*") 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 a 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 $\alpha_0 = \alpha$. That is, *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.

As emphasized in the introduction, the sampling procedure should not be taken literally. I do not describe how patients produced their sample, and sampling costs are irrelevant. The $S(1)$ procedure is not a model of the patients' data-gathering process, but a model of how they make inferences from sparse anecdotal evidence. When a patient chooses alternative i, the outcome of treatment i is a new, independent draw, such that the patient's true expected utility from this decision is $\alpha - p_i$, not $x_i - p_i$. This is in sharp contrast to the rational-search literature, which focuses attention on the way rational consumers optimally design samples when there are sampling costs.

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 heard a good anecdote only about him. 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 heard a good anecdote about him and a bad anecdote about the default.

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. The patients know the exact prices; the only source of variance in their samples is the imperfect recovery rate α, which is exogenously given. Thus, when a healer plays a mixed strategy, he introduces uncertainty into his opponents' environment, but not into his patients'.

The simplicity of the $S(1)$ procedure inevitably means that it is artificial in a number of ways. It also naturally raises the question of whether our model can be "rationalized" - i.e., substituted with a standard incomplete-information game, without changing the results. I discuss these issues in Section 8.

3 Equilibrium

This section is devoted to analyzing Nash equilibrium in the price-competition game with $S(1)$-patients. The following proposition is our basic result.

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

Thus, the price-competition game has a unique Nash equilibrium, which is symmetric and mixed. To use the terminology of the search literature, the game results in equilibrium price dispersion. The quacks' 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 $[\frac{1}{2}, 1]$.

To see the origin of expression (1), suppose that we restricted attention to symmetric equilibria. Let G denote the equilibrium *c.d.f* and let π denote the quacks' equilibrium payoff. Then, for every price p in the support of G:

$$\pi = p \cdot \alpha \cdot (1 - \alpha_0) \cdot [1 - \alpha G(p)]^{n-1} \tag{2}$$

because for every quack i, $\alpha(1-\alpha_0)$ is the probability that $x_i - p_i > x_0$ in a patient's sample, and $1 - \alpha G(p)$ is the probability that in the patient's sample, $x_j - p_j > x_i - p_i$ for some other quack

j. (Although we assume that $\alpha_0 = \alpha$, I identify α_0 as such in the expression for expositional purposes.) From expression (2), it follows that:

$$G(p) = \frac{1}{\alpha} \cdot [1 - \frac{c}{\sqrt[n-1]{p}}] \tag{3}$$

where c is some constant. By fairly standard arguments, the monopoly price $p = 1$ belongs to the support of G. Therefore, we can retrieve the value of c by plugging $p = 1$ and $G(1) = 1$ in expression (3).

This is quite similar to derivations of price distributions in the literature on equilibrium price dispersion (e.g., Butters (1977), Varian (1980), Burdett and Judd (1983) and Rob (1985)). The more novel and substantive part of the proof of Proposition 1 is the demonstration that the symmetric equilibrium is the *unique* Nash equilibrium. In this technical sense, the result goes beyond the above-cited papers.[1]

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

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

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

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

$$E(p) \underset{\alpha\to1}{\to} 0$$
$$E(p) \underset{\alpha\to0}{\to} 1$$

[1] In the models analyzed by the above-cited papers, attention is restricted to symmetric mixed-strategy equilibria, either explicitly or, more often, implicity, by pursuing a non-game-theoretic definition of market equilibrium, which is equivalent to symmetric mixed-strategy equilibrium in a game-theoretic formulation.

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 aggregate demand for quacks to shrink. (The fraction of patients who acquire some treatment is $1-\alpha-(1-\alpha)^{n+1}$: this is the probability that a patient heard a bad anecdote about the default and at least one good anecdote about a quack.) 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. For every α, healers have a positive clientele. That is, the "market for quacks" is always active in equilibrium.

The results illuminate the phenomenon of *charlatanry* in the market for quacks. In equilibrium, quacks behave as charlatans: they charge a positive price for a worthless treatment. The false pretense implicit in their over-pricing gets worse as α decreases. In other words, their charlatanry becomes more pronounced as the situation becomes more "hopeless". In addition, as α decreases, every quack attracts a small number of patients who are willing to pay large amounts for his treatment, while being dismissive of alternative treatments. Thus, our model characterizes the phenomenon of charlatanry in the market for quacks, and traces it to the consumers' boundedly rational inferences.

Welfare analysis

Because the monopoly price $p=1$ belongs to the support of G, the quacks' equilibrium payoff is $\alpha(1-\alpha)^n$, according to expression (2). Therefore, the industry's equilibrium profits are given by the following expression:

$$\Pi = n\alpha(1-\alpha)^n \tag{4}$$

Because quacks do not contribute any added value, industry profits are equal to the welfare loss inflicted on patients. Thus, Equation (4) also provides the expression for the patients' equilibrium welfare loss.[2] The R.H.S of Equation (4) is not monotonic in α: it attains an maximum at $\alpha^* = \frac{1}{n+1}$. 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 (in the sense of larger* n*) may increase the patients' welfare loss.* As $\alpha \to 0$, n^* tends to infinity, such that the perverse effect of greater competition holds for a larger domain of n.

The intuition 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, an increase in n leads to higher aggregate demand for quacks. This "exploitative" effect is a consequence of the $S(1)$ procedure: when the market contains a larger variety of treatments, there is a higher chance of hearing a good anecdote about some treatment. 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 payoff. However, the max-min payoff is positive because patients err with positive probability. *The "exploitative effect" determines the max-min payoff, and the "competitive effect" does not allow quacks 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 continue to hold when $\alpha \neq \alpha_0$. The

[2]Equation (3) is also consistent with the case of $n = 1$, because the maximal payoff of a monopolistic healer is $\alpha(1-\alpha)$.

healers' equilibrium behavior is independent of the default success rate, because it enters the healers' payoff function through the multiplicative term $1 - \alpha_0$, and cancels out as we derive the expression for G.

Although relaxing quackery does not affect the healers' equilibrium behavior, it affects the welfare analysis. For example, when $\alpha > 0$ and $\alpha_0 = 0$, the patients' net payoff is $a - na(1-a)^{n-1}$. From this expression, it can be seen that there is a net welfare loss if a is sufficiently low.

The patients' knowledge of the default

The basic model assumes that the patients' choice procedure treats the default and the healers symmetrically: patients sample each of them once. It could be argued that patients are more familiar with the default than with the healers. For example, patients may enter the market for unconventional medicine only after having learned the true probability of recovery under conventional treatments. Therefore, it is interesting to check a variant on the model, in which $x_0 = \alpha_0$ with probability one. That is, patients know the default recovery rate, while forming quality assessments about healers on the basis of anecdotes.

The essential features of our equilibrium characterization - uniqueness, symmetry, price dispersion and the qualitative comparative statics - remain unchanged under this modification. 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. In particular, the patients' welfare loss is lower than in the case in which they do not know the default recovery rate.

4 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 disclose his success rate to patients. If he does not reveal his success rate, patients continue to assess his quality according to the $S(1)$ procedure. In this context, it would be appropriate to allow more general market primitives than in Section 2. Denote the rate of recovery associated with alternative i by α_i, and allow the α_i's to vary across alternatives, where $\alpha_i \in (0, 1)$ for every $i = 0, 1, ..., n$.

Formally, a strategy for healer i is a pair (p_i, r_i), where $r_i = Y$ (N) if the healer reveals (does not reveal) his success rate. As in the basic model, x_i denotes the patient's impression of the quality of 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. Note that the meaning of this procedure is that patients do not infer anything from the healer's disclosure policy itself. That is, this extension of the $S(1)$ procedure to the model with disclosure means that patients do not take into account adverse-selection effects.

In a standard adverse selection model with rational, imperfectly informed patients, it would be standard to assume that the patients know the distribution of success rates across among healers, but do not know ex-ante the exact assignment of success rates to healers. In Bayesian equilibrium, every healer would disclose his success rate (except possibly the lowest types).[3] Given our model of the patients' behavior, the result is the complete opposite:

Proposition 2 *For every p, the strategy (p, Y) for healer i is weakly dominated by some other strategy (p', N).*

Proof. Denote $\alpha_i = \alpha$, for notational convenience. If $p > \alpha$ and $r_i = Y$, then clearly no patient will choose healer i, and therefore, the healer's payoff from (p, Y) is zero. In this case, (p', N) clearly dominates (p, Y).

Let $p = \alpha$. Then, in a patient's sample, the probability that $x_i - p_i > x_0$ is zero, and the probability that $x_i - p_i = x_0$ is $\alpha \cdot (1 - \alpha_0)$. If healer i deviates to $(1 - \varepsilon, N)$, the probability that $x_i - p_i > x_0$ is $\alpha \cdot (1 - \alpha_0)$, and the probability that $x_i - p_i = x_0$ is zero. Therefore, this deviation is profitable, regardless of the other healers' strategies. Therefore, $(1 - \varepsilon, N)$ strictly dominates (p, Y).

Finally, consider the case of $p < \alpha$. In this case, healer i's payoff from the strategy (p, Y) is bounded from above by:

$$p \cdot \Pi_{j \neq i} \Pr(x_j - p_j \leq \alpha - p)$$

[3] Board (2003) constructs a variant on such a model, in which firms choose their disclosure policy prior to the price-setting stage. He shows that in this case, the full disclosure result breaks down.

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

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

Now, let us show that (p, Y) is weakly dominated by (p', N), where $p' = p/\alpha$. Then, $p' \in (p, 1)$. Since $\alpha - p = \alpha \cdot (1 - p')$, it is clear that $1 - p' > a - p$ as long as $p < a$. Therefore:

$$\Pi_{j \neq i} \Pr(x_j - p_j < 1 - p') \geq \Pi_{j \neq i} \Pr(x_j - p_j \leq \alpha - p) \tag{5}$$

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, when a healer deviates from $r_i = Y$ to $r_i = N$, the "monopoly" price jumps from α_i to 1. On the other hand, the fraction of patients who are willing to pay anything to healer i shrinks from 1 to $1 - \alpha_i$. The reason that the former consideration outweighs the latter is simple. Suppose that $p < \alpha$. By deviating from (p, Y) to $(p/\alpha_i, N)$, the healer replicates his monopoly profits. At the same time, he attains an edge over competitors because conditional on $x_i = 1$, the patient's perceived utility from choosing healer i is $1 - p/\alpha_i$, compared to $\alpha_i - p$ (the patient's perceived utility from choosing healer i when $r_i = Y$.)

Proposition 2 established that type revelation is a weakly dominated strategy. The following result verifies that type disclosure can never be part of Nash equilibrium. The proof appears in the appendix.

Proposition 3 *In Nash equilibrium, every healer chooses $r_i = N$.*

The lesson from Propositions 2 and 3 is that given our model of the patients' behavior, even high-quality healers prefer that patients form their quality judgments on the basis of random, anecdotal evidence. This effect is a consequence of the two main aspects of the patient's choice procedure. First, they do not distinguish between the informative content of full knowledge of a treatment's success rate and a single random anecdote. Second, they do not draw any inference from the healer's actual disclosure policy.

5 Replacing a Quack with an Expert

In this section, I perturb the basic model of Section 2 by replacing one of the quacks with a high-quality healer. The question is to what extent this intervention will crowd out the quacks. Formally, modify the basic model 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 $(= \alpha_0)$. In other words, healer e is an "expert" while his opponents are "quacks".

Proposition 4 *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 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. It also follows that the lowest price in the market continues to be $(1-\alpha)^{n-1}$. 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 4, 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 $p = 1$. Finally, turning to the patients' welfare, 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.

6 A Generalized $S(K)$ Procedure

The $S(1)$ procedure captures an extreme case of a "law of small numbers": patients form *deterministic* action-consequence correspondences on the basis of 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 v denotes his willingness to pay for sure recovery. Assume that v is distributed over the interval $[0,1]$ according to a continuous *c.d.f* F. Let $p^* = \arg\max_{p\in[0,1]} p[1-F(p)]$. Assume that F satisfies the usual properties that guarantee that p^* is interior and unique.

The $S(K)$ 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 success rate 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 as $K \to \infty$, 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.

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 *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. ■

Contrary to the case of $K = 1$, continuity of F is needed to ensure existence in the generalized model. To see why, let $K = 2$, and suppose that $v = 1$ for all patients. 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 when he charges a price above $\frac{1}{2}$. Therefore, his utility function can be discontinuous at $p_i = \frac{1}{2}$, even when $p_j \neq p_i$ for every $j \neq i$. This kind of discontinuity may be ruinous for existence.[4]

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

[4]For a discontinuous F, Nash equilibrium exists if α is sufficiently small. In this range, $p = \frac{1}{K}$ dominates any price $p > \frac{1}{K}$. In mixed-strategy equilibrium, only prices in the interval $(0, \frac{1}{K}]$ receive a positive density, and the discontinuity problem becomes irrelevant.

Proposition 6 *As $\alpha \to 0$, 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 never hear a good anecdote about healers. A vast majority of the remaining patients have a single success in their sample. The clientele targeted by healer i consists of patients with $a_i = \frac{1}{K}$ and $a_j = 0$ for every $j \neq i$. Therefore, 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 the rational benchmark.

Proposition 7 *The patients' expected equilibrium payoffs converge to α, as $K \to \infty$.*

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. ■

Propositions 6 and 7 confirm the intuition that greater rationality on the patients' part makes the market more competitive.

7 $S(1)$-Patients in a Differentiated Market

In Section 3, we saw how even in the case of $\alpha > \alpha_0$, the patients' choice procedure causes a market for a genuinely homogeneous good to be organized like a market for a differentiated good. As a result, the market equilibrium outcome is less competitive than in the rational-patients benchmark.

What is the implication of the $S(1)$ procedure in a market that is genuinely differentiated? Consider the following example. Let $n = 2$ and $\alpha_0 = 0$. Suppose that the population of patients is divided into two groups of identical size, A and B. The healers' success rates are group-specific. Let α_i^j denote healer i's success rate with group j. Assume that $\alpha_1^A = \alpha_2^B = 1$ and $\alpha_1^B = \alpha_2^A = \alpha \leq \frac{1}{2}$. That is, healer 1 (2) is an "expert" for group A (B) and an "amateur" for group B (A). In a standard price-competition model with rational patients and complete information, the market is completely differentiated in equilibrium, such that every healer charges the monopoly price $p = 1$ and specializes in his expertise group.

The following result describes market equilibrium when patients choose according to the $S(1)$ procedure. (I omit the proof, because it follows the same lines as the proof of Proposition 1.)

Proposition 8 *There is a unique Nash equilibrium in the price-competition game played between the healers. Every healer plays the same mixed strategy, given by the c.d.f:*

$$G(p) = \frac{1}{1+\alpha} \cdot [2 - \frac{1-\alpha}{p}]$$

over the support $[\frac{1-\alpha}{2}, 1]$. *Expected price is given by:*

$$Ep = -\frac{1-\alpha}{1+\alpha} \ln(\frac{1-\alpha}{2})$$

Thus, when patients choose according to the $S(1)$ procedure, the market outcome is *more competitive* (in the sense of lower prices) than in the rational-patients benchmark. Expected equilibrium price falls below the monopoly price $p = 1$. At the same time, the market outcome is inefficient, because patients from one group choose with positive probability a healer who specializes in the other group. Despite this inefficiency, patients' welfare can be higher than in the rational benchmark. When $\frac{1}{4} < \alpha < \frac{1}{2}$, the patients' true equilibrium expected payoff (measured by $\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \alpha - Ep$) is positive, whereas in the rational benchmark, it is zero.

The intuition for this result is that the $S(1)$ procedure creates a spurious perception among patients that healers are more homogeneous than they really are, just as in the case of a genuinely homogeneous market, the procedure creates a spurious perception among patients that healers are more heterogeneous than they really are.

8 Discussion

This paper explored the implications of consumers' boundedly rational statistical inferences - specifically, their tendency to draw exaggerated inferences from sparse anecdotal evidence - on market equilibrium and consumer welfare. For expositional purposes, I focused on a market for a worthless good or service, and called it "the market for quacks". Throughout the paper, I used the unconventional medicine industry as a typical example for such a market. Although the model is of course extremely stylized and abstract, I believe that its underlying assumptions are quite descriptive of this industry. Consumers of unconventional medicine often have a poor understanding of the underlying stochastic variables; they often rely on anecdotal evidence when choosing a treatment; and therefore, they are quite vulnerable to exploitation because of their tendency to over-infer from anecdotes.

The results of this paper - the healers' charlatanry, the market's tendency to a state of monopolistic competition as the probability of recovery decreases, the increase in aggregate demand caused by the multiplicity of healers, the healers' reluctance to reveal their exact success rates - can be viewed as a theoretical characterization of some broad I.O. features of the unconventional medicine industry, which may serve as a benchmark for systematic empirical analysis of this industry (which, to my knowledge, has not been carried out yet).

To a varying extent, the modeling approach of this paper may be relevant to other markets for expert service. For instance, even in the mutual funds industry - in which investors have much more statistical information at their disposal than in the unconventional medicine industry

- they sometimes choose a mutual fund according to its performance in the very recent past. They ascribe the fund's success to its expertise, rather than to chance. This type of behavior is consistent with the "law of small numbers" aspect of the $S(1)$ procedure. It may provide a basis for an explanation of the puzzling observation that investors are willing to pay substantial fees to actively managed mutual funds, which are not justified by their performance relative to passive (index) funds. (See Gruber (1996) and Wermers (2000).) This and other extensions are left for future research.

Rationalizing the model

Despite its simplicity, the modeling procedure in this paper is non-standard. Our starting point was a standard price-competition model with complete information. Usually, market models depart from such a simple benchmark by perturbing the informational structure. Instead, the model of Section 2 perturbs the consumers' rationality. The question naturally arises, whether the basic model of Section 2 and its various extensions could be "rationalized", in the sense that the same results could be obtained from a price-competition model with asymmetric information.

Let us begin with the basic model of Section 2, starting with what may seem like the most natural rationalization of this model: substituting the patients' imperfect rationality with imperfect information. In such a model, the healers' success rates are drawn from some prior distribution, which is commonly known among all market agents. Healers know their own success rates, whereas patients observe partially informative signals about the quality of healers.

To see why such an adverse-selection model cannot yield the same results as our model, note that in the model of Section 2, patients always end up being absolutely certain of the quality of each alternative, and consequently their willingness to pay "jumps" to 1 or 0. A partially informed rational patient would not display such a "jump" in his willingness to pay, unless he can infer the healers' quality from their equilibrium pricing behavior. However, because in this model healers may only differ in their success rates, there is no mechanism that could prevent low-quality healers to pool with high-quality healers. It follows that this type of adverse-selection model cannot yield the same pricing behavior as our model.

Let us now turn from a "rationalizing model" with common values and imperfect information to a model in which patients have independent, private values. That is, assume that v_i, the patient's valuation of healer i, is the patient's private information. The v_i's are independently drawn, and take the value 1 (0) with probability α $(1-\alpha)$. This is a price-competition model in which consumers have private, heterogeneous valuations, similar in spirit to Perloff and Salop (1985).

As far as the healers' behavior is concerned, this reinterpretation works well: the healers have the same payoff function as in our basic model, and therefore the Nash equilibrium is the same. However, the *welfare analysis* in the two models is entirely different. In the private-values rationalization, patients never make errors. Therefore, they do not suffer a welfare loss, and their welfare is monotonically increasing in n.

In general, any model with rational patients will fail to reproduce the welfare analysis of Section 3. Rational patients may choose sub-optimal alternatives with positive probability if they are imperfectly informed. However, they cannot be fooled systematically. Therefore, their equilibrium payoff cannot possibly be lower than the expected payoff that they obtain from the default.

Does the private-values rationalization survive the extensions of the basic model? First, observe that the rationalization is entirely meaningless for the model of Section 4. When a patient's valuations are his private information, it makes no sense to talk about disclosure by healers. As to the extended models of Sections 5 and 6, each model by itself can be rationalized along similar lines, where the probability distributions from which valuations are drawn are tailored specifically to each model. However, the comparison between these models to the basic model becomes meaningless.

For these reasons, I find that the private-values rationalization has very limited merit. It cannot reproduce the welfare analysis of Section 3, and it renders the extended models (or their relation to the basic model) meaningless. Moreover, the insights that the basic model provides into the phenomenon of market charlatanry depend on the bounded-rationality interpretation of the patients' model of behavior. These insights disappear under the private-values interpretation.

Variations on the patients' behavioral model

As mentioned in Section 2, the simplicity of the $S(1)$ procedure inevitably means that it contains some artificial features. For instance, consider the assumption that patients sample *every* healer. It would be more realistic to assume that patients get to hear anecdotes about a subset of healers. One way to deal with this artificiality is to assume that the patients' sample size is m, such that if $n > m$, the patient randomly samples m out of the n healers, whereas he is totally unaware of the remaining $n - m$ healers. In this case, m should replace n in expressions (1)-(4). An increase in m cannot be interpreted as market entry, but as an increase in the patients' awareness of the set of available treatments.

Another artificial feature is the assumption that the number of observations about healer i is independent of the size of the healer's clientele. Alternatively, we could assume that patients hear more anecdotes about healers with a larger clientele. I expect this variant to be much less tractable than the basic model. However, I do not believe that it would reverse the results of Section 3. When patients have a larger number of observations about a healer, they are less likely to err in their assessment of his success rate. Therefore, quacks have an incentive to reduce the number of observations K_i that patients have about them. The model of Section 4 provides an extreme illustration of this argument. If K_i is increasing in healer i's clientele size, this means that healer i has an incentive to reduce his clientele, and he can only achieve this by raising prices. Therefore, the assumption that K_i increases in clientele size may even weaken competitive pressures in the market for quacks.

What happens when we introduce standard rational patients into the population? 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 $\varepsilon \rightarrow 1$, equilibrium payoffs converges to the rational benchmark.

9 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 aspects of non-Bayesian inference. In recent years there has been a number of attempts to introduce elements of non-Bayesian inferences into economic modeling. Eyster and Rabin (2002) analyze players in a

Bayesian game, who optimize against the opponents' statistical distribution of actions, rather than against their Bayesian-game strategies. In Jehiel (2003), players in an extensive game optimize against the statistical behavior of opponents across different histories that belong to the same "analogy class". Spiegler (2003) analyzes repeated games with players who are willing to believe a threat if and only if they observe it being realized on the play path. 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 $S(1)$ 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 Osborne-Rubinstein models are irrelevant in the present context. In these models, *every* agent in the environment chooses according to the $S(1)$ procedure. The agents are strategic, and the fact that they use a non-standard choice procedure impels Osborne and Rubinstein to invent a new game-theoretic equilibrium concept. Sethi (2000) and Mię̧kisz and Ramsa (2003) study equilibrating processes that may justify this concept. In the present paper, $S(1)$ is the patients' choice procedure. The patients are not strategic. The only strategic agents are the healers, and they are standard, profit-maximizing agents. The solution concept that describes their strategic behavior Nash equilibrium. Thus, the bounded-rationality element in the present model does not raise new problems concerning equilibrium behavior.

The present paper belongs to a growing literature on market competition among rational firms over boundedly rational agents. Rubinstein (1993) analyzes monopolistic behavior when consumers differ in their ability to understand complex pricing schedules. Piccione and Rubinstein (2003) study the behavior of agents with heterogeneous ability to perceive intertemporal patterns, in competitive and monopolistic markets. Fishman and Hagerty (2003) study a model of voluntary disclosure by sellers, when some consumers cannot understand the content of the disclosure. (They can, however, make correct Bayesian inferences from the seller's disclosure decision itself.) Chen et. al. (2002) analyze a model of price competition when consumers have memory imperfections that constrain their ability to conduct market search. Mullainathan and Shleifer (2002) propose a model of competition in the media market, in which consumers of news have bounded memory.

As mentioned previously in this section, the basic model is formally related to models of competition over consumers with private, heterogeneous valuations, in the spirit of Perloff and Salop (1985). In a recent paper, Gabaix and Laibson (2004) extend the Perloff-Salop model in various directions, and give it a bounded-rationality interpretation, according to which the consumers' probabilistic choice is a consequence of the complexity of the firms' products.

Finally, as noted in Section 3, the technique of characterizing mixed-strategy equilibrium in the basic model is related to the literature on equilibrium price dispersion (e.g., Butters (1977), Varian (1980), Burdett and Judd (1983), Rob (1985)). In this literature, firms are involved in price competition in an imperfectly competitive market, where the imperfection is due to search costs. As emphasized in Section 3, the above-cited papers restrict attention to symmetric equilibria, whereas Proposition 1 establishes the uniqueness of symmetric equilibrium.

10 Appendix

10.1 Proof of Proposition 1

Healer i's equilibrium strategy s_i induces a *c.d.f* G_i over the interval $[0,1]$. The main task in this proof is to show that the equilibrium is symmetric.

First, let us show that for every healer i, 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 none of the G_i's contains an atom on any price $p<1$.

Note that if G_i contains an atom on $p=1$, then there exists no other healer j, such that G_j's contains an atom on $p = 1$. Otherwise, i or j would be able to deviate profitably by shifting this atom slightly downward.

In 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 satisfies $x_i = 1$ 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, even if different healers charge p^L and p^H. The reason is as follows. Regardless of the opponents' strategies, the price $p^H = 1$ yields a payoff of $\alpha(1-\alpha)^n$, and the price p^L yields a payoff of $p^L \cdot \alpha \cdot (1-\alpha)$. If these payoffs are not equal, then a player who charges the less profitable of these prices can profitably deviate to a strategy that assigns probability one to the more profitable of the two prices. It follows that both p^L and p^H yield a payoff of $\alpha(1-\alpha)^n$. Therefore, $p^L = (1-\alpha)^{n-1}$. By the same argument, it follows that every healers's equilibrium payoff does not fall below $\alpha(1-\alpha)^n$.

Let us now show that $p_i^L = p^L$ for every $i = 1, ..., n$. Assume the contrary and suppose, without loss of generality, that $p_1^L > p^L$. We already established that healer 1's payoff must be at least $\alpha(1-\alpha)^n$. Suppose, without loss of generality, that $p_2^L = p^L$ and that healer 1's payoff is strictly larger than $\alpha(1-\alpha)^n$. Recall that player 2's payoff is exactly $\alpha(1-\alpha)^n$. Consider the following deviation for healer 2, from s_2 to the pure strategy $p_2 = p_1^L$. This deviation must be profitable, for the following reason. Compare healer 1's situation before player 2's deviation and healer 2's 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 by definition, s_1 assigns probability one to prices above p_1^L. Therefore, healer 2's payoff after the deviation cannot fall below healer 1's payoff before the deviation. Therefore, healer 2's deviation is profitable.

It follows that all healers make a payoff of $\alpha(1-\alpha)^n$ in equilibrium, because they all have

$p_i^L = (1-\alpha)^{n-1}$ and because the healers' strategies do not contain 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 $i = 1, ..., n$:

$$\alpha \cdot (1-\alpha)^{n-1} \cdot (1-\alpha_0) = p \cdot \alpha \cdot (1-\alpha_0) \cdot \prod_{j \neq i} [1 - \alpha G_j(p)] \tag{6}$$

We have a set of n equations in n variables $G_j(p)$.[5] 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 $p \in ((1-\alpha)^{n-1}, 1)$.

The system of equations (6) 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' $c.d.f$'s contains "holes". If there existed a price $p \in (p^L, 1)$, such that only $m < n$ healers assigned a positive measure to the neighborhood of p, then the system (5) would consist of m equations, and the solution would necessarily mean that some of the G_j's are discontinuous at p, a contradiction.

It is now straightforward to derive Equation (1) from the system of equations given by (5). 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.

10.2 Proof of Proposition 3

First, note that every healer is able to earn positive payoffs in the game, by choosing $r_i = N$ and a sufficiently low price. Let us 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 $\alpha_i - p_i < \alpha_j - p_j$. By standard "Bertrand" reasoning, these two healers must earn 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 order for patients choose

[5] Despite the assumption $\alpha_0 = \alpha$, I write α_0 explicitly in expression (5), in order to make it evident that the characterization of equilibrium strategies does not depend on this assumption.

healer 1, it must be the case that whenever $x_j = 1$, $\alpha_1 - p_1 > 1 - p_j$. There must be at least one healer j, for whom G_j is atomless on $p_j = 1 - \alpha_1 + p_1$, and strictly increasing in the neighborhood of $p_j = 1 - \alpha_1 + p_1$. Otherwise: (1) if one G_j contained an atom on p_j, healer 1 would not assign positive density to p_1; (2) if all G_j's were constant in the neighborhood of p_j, healer 1 could profitably deviate by shifting weight from p_1 upwards. The case of $r_1 = Y$ and $p_1 \geq \alpha_1$ is ruled out by the same arguments used in the proof of Proposition 3. Therefore, in equilibrium, if $r_1 = Y$, $p_1 < \alpha_1$.)

Because G_j is strictly increasing in the neighborhood of $p_j = 1 - \alpha_1 + p_1$ for at least one healer $j > 1$, the function $\Pr(x_j - p_j \leq \alpha_1 - p_1)$ is *strictly* decreasing in p_1. 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 2 for the case $p_1 < \alpha_1$, inequality (4) is strict, such that the deviation is profitable.

10.3 Proof of Proposition 4

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, there is at most one healer whose equilibrium strategy places an atom on $p = 1$. Second, $p^H = 1$.

Let us first analyze the case of $n > 2$, such that there are at least two quacks in the market. Using the same symmetry arguments as in the proof of Proposition 1, all quacks play the same strategy: $G_i = G$ for every $i \neq e$. In particular, they all have the same p_i^L, and G cannot place an atom on $p = 1$. In contrast, G_e may contain an atom on $p = 1$. Denote the size of this atom by A.

Let us establish that healer e's equilibrium payoff is equal to $\alpha_e(1-\alpha)^n$. First, First, suppose that $p_e^H = 1$. Then, $p = 1$ maximizes healer e's payoff, given the quacks' strategy G. Recall that G does not place an atom on $p = 1$. Therefore, healer e's payoff is $\alpha_e(1-\alpha)^n$. Second, suppose that $p_e^H < 1$. Then, $p_i^H = 1$ for all $i \neq e$. It follows that the quacks' payoff is $\alpha(1-\alpha)^{n-1}(1-\alpha_e)$. Healer e's payoff cannot be greater than $\alpha_e(1-\alpha)^n$. Assume the contrary. Consider a deviation of a healer $i \neq e$ to the pure strategy $p = p_e^H$. Healer e's payoff after i's deviation is at least $\alpha_e(1-\alpha)^n$, because the probability that $p_e < p_i$ did not decrease as a result. Therefore, healer

i's payoff after his deviation is at least $\alpha(1-\alpha)^n$, a profitable deviation. But healer e's payoff cannot be lower than $\alpha_e(1-\alpha)^n$, because this is his max-min payoff.

If $p_e^L < p_i^L$, then healer e can profitably deviate by shifting the measure he assigns to the neighborhood of p_e^L upwards, towards p_i^L. Therefore, $p_e^L \geq p_i^L$. Denote $p_i^L = p^L$ for $i \neq e$. 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:

$$\alpha_e \cdot (1-\alpha)^{n-1} = p \cdot \alpha_e \cdot (1-\alpha) \cdot [1-\alpha G(p)]^{n-1} \tag{7}$$

Therefore, we obtain the same expression for G as in the case of Proposition 1. Furthermore, it must be the case that G_e assigns a positive density to every $p \in (p_e^L, 1)$. Otherwise, equation (7) would imply that G is discontinuous at some $p < 1$, a contradiction.

It remains to show that $p_e^L = p^L$. Assume the contrary - i.e., $p_e^L > p^L$. For $p > p_e^L$, $G(p)$ is given by equation (7). Note that G is continuous at $p = p_e^L$. Therefore, in the limit, as p tends to p_e^L from above, we obtain:

$$(1-\alpha)^{n-1} = p_e^L \cdot [1-\alpha G(p_e^L)]^{n-1} \tag{8}$$

Note that $p_i^H = 1$ for every $i \neq e$. The reason is as follows. All quacks play the same strategy. Therefore, $p_e^H \leq p_i^H$ - otherwise, healer e can profitably deviate by shifting weight from the neighborhood of p_e^H downwards, toward p_i^H. If $p_i^H < 1$, then any healer $i \neq e$ can profitably deviate by shifting weight from p_i^H upwards.

The quacks' payoff at $p = 1$ is $\alpha \cdot (1-\alpha)^{n-1} \cdot (1-\alpha_e+\alpha_e A)$, where A is the size of the atom that G_e places on $p = 1$. Because $p_i^H = 1$, p maximizes the quacks' payoff. Therefore, $p^L = (1-\alpha)^{n-1} \cdot (1-\alpha_e+\alpha_e A)$. As p tends to p_e^L from below, we obtain:

$$\alpha \cdot (1-\alpha)^{n-1} \cdot (1-\alpha_e+\alpha_e A) = p_e^L \cdot \alpha \cdot (1-\alpha) \cdot [1-\alpha G(p_e^L)]^{n-2} \tag{9}$$

Now, if healer e deviated to the pure strategy $p = p^L$, his payoff would be $p^L \cdot \alpha_e \cdot (1-\alpha)$. In order for this to an unprofitable deviation, it must be the case that $1 - \alpha_e + \alpha_e A \leq 1 - \alpha$. Therefore, by equation (9):

$$p_e^L \cdot [1 - \alpha G(p_e^L)]^{n-2} \leq (1-\alpha)^{n-1}$$

but by assumption, $G(p_e^L) > 0$, which implies that:

$$p_e^L \cdot [1 - \alpha G(p_e^L)]^{n-1} < (1-\alpha)^{n-1}$$

in contradiction to equation (8).

Therefore, $p_e^L = p^L$. From equation (8) we deduce that $p^L = (1-\alpha)^{n-1}$. Therefore, for every $p \in [(1-\alpha)^{n-1}, 1]$, $G(p)$ is given by equation (8), which yields the same expression for G as in Proposition 1. From equation (9) we deduce that $1 - \alpha_e + \alpha_e A = 1 - \alpha$, such that the quacks' payoff is $\alpha \cdot (1-\alpha)^n$. Thus, we established that in equilibrium, the quacks play the same strategy and earn the same payoff as in the equilibrium of the basic model. It follows immediately that the fraction of patients who choose the quacks is the same as in the basic model.

The case of $n = 2$ should be handled separately, because there is one expert and one quack, and so the argument that all quacks play the same strategy is irrelevant. However, in this case it is much more straightforward to show that $p_e^L = p_i^L$ and $p_e^H = p_i^H = 1$. From this point, the derivation of the quack's strategy and payoff is just the same as in the case of $n > 2$.

References

[1] Board O. (2003), "Competition and the Impact of Mandatory Disclosure Laws", Mimeo.

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

[3] Butters G. (1977), "Equilibrium Distributions of Sales and Advertising Prices", *The Review of Economic Studies* **44**, 465-491.

[4] Chen Y., G. Iyer and A. Pazgal (2002), "Limited Memory and Market Competition", mimeo.

[5] Eyster E. and M. Rabin (2002), “Cursed Equilibrium”, Mimeo.

[6] Fishman M. and K. Hagerty (2003), "Mandatory versus Voluntary Disclosure in Markets with Informed and Uninformed Customers", *Journal of Law, Economics and Organization* **19**, 45-63.

[7] Gabaix X. and D. Laibson (2004), “Competition and Consumer Confusion”, Mimeo.

[8] Gruber M. (1996), “Another Puzzle: The Growth in Actively Managed Mutual Funds”, *The Journal of Finance* **51**, 783-810.

[9] Jehiel P. (2003), “Analogy-Based Expectations Equilibrium”, *Journal of Economic Theory*, forthcoming.

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

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

[12] Mię̧kisz J. and M. Ramsa (2003), “Dynamical Models of Sampling Procedures”, Mimeo, Warsaw School of Economics.

[13] Mullainathan S. and A. Shleifer (2002), “Media Bias”, NBER Working Paper No. w9295.

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

[15] Osborne M. and A. Rubinstein (2003), “Sampling Equilibrium with an Application to Strategic Voting”, *Games and Economic Behavior* **45**, 434-441.

[16] Perloff J. and S. Salop (1985), “Equilibrium with Product Differentiation”, *The Review of Economic Studies* **52**, 107-120.

[17] 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.

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

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

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

[21] Rubinstein A. (1998). *Modeling Bounded Rationality.* MIT Press.

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

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

[24] Spiegler R. (2003), "Testing Threats in Repeated Games", Mimeo.

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

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

[27] Wermers R. (2000), "Mutual Fund Performance: An Empirical Decomposition into Stock-Picking Talent, Style, Transactions Costs, and Expenses", *The Journal of Finance* **55**, 1655-1695.

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