The Market for Quacks*

Ran Spiegler†

First version: March 2003
This version: December 2004

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 chose rationally, the market would be inactive. I assume, however, that patients choose according to a boundedly rational procedure, which reflects "anecdotal" reasoning. This element of bounded rationality has significant implications. The market for quacks is active and patients suffer a welfare loss which behaves non-monotonically w.r.t n and α. When quacks are allowed to choose their treatment, they differentiate themselves by betting as much as possible against competitors. The welfare loss that quacks inflict on patients is robust to market interventions which would crowd out low-quality firms in a standard I.O. model. I discuss the implications of these findings for a variety of industries, including forecasting and mutual funds.

*I would like to thank Eddie Dekel, Kfir Eliaz, Gur Huberman, Volker Nocke, Michele Piccione, Andrea Prat, Ariel Rubinstein, Asher Wolinsky, the editor and two referees, as well as seminar audiences at Cambridge, Essex, LSE and UCL, for 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

Imagine a hypothetical market, in which n identical "*healers*" are engaged in price competition over a continuum of "*patients*". If a patient 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 the same probability α. We may refer to the healers as "*quacks*", because they have absolutely no skills relative to the default.

If patients understood this model, they would realize that the entire industry provides a worthless service. Indeed, standard market models presuppose that all market agents, firms and consumers alike, have full understanding of the model. Even if market agents differ in the information they have, at the meta-level they share a "common understanding of the model". If the market model contains stochastic elements, all market agents are perfectly able to understand them. In the case of the "market for quacks" - an industry in which all firms provide a worthless service - such common understanding means that the market will be inactive in equilibrium.

The common-understanding-of-the-model presupposition is often problematic. Firms tend to enjoy frequent interactions with the market, whereas consumers often make a once-and-for-all decision following little exposure to the market. Thus, firms have more opportunities to learn the market model. In addition, the firms' decision process is entirely focused on a specific market, whereas consumers divide their attention among several markets. For these reasons, we should not expect firms and consumers to display the same quality of reasoning about the market.

In this paper I study the market for quacks under the following pair of assumptions: *(i)* healers are standard profit maximizers with respect to a correct probabilistic understanding of the market model; *(ii)* patients follow a boundedly rational decision rule. Instead of reasoning probabilistically with respect to a correct market model, patients reason *anecdotally*: they rely on random, casual stories regarding the quality of treatments, and regard them as being fully informative of the treatments' actual quality. Anecdotal reasoning exposes consumers to exploitation by firms, because they may attribute the firms' random successes to skill rather than luck. The question is whether market competition could mitigate this exploitative effect. I examine the extreme case of the "market for quacks" - namely, a market for a completely worthless good or service - in order to bring this question into sharper focus.

To capture the patients' anecdotal reasoning, I assume that they choose according to a simple procedure, called $S(1)$, which I borrow from Osborne and Rubinstein (1998). A patient samples (once) each of the $n+1$ alternatives (the healers and the default). The 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). A sample point x_i is interpreted as a random anecdote that the patient has gathered about alternative i, either from his own past experience or from a fellow patient. 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 from healer i's treatment. The quacks take into account the patients' choice procedure when determining their strategies.

This behavioral model is relevant to markets for goods or services which generate a random outcome, when it is difficult for consumers to gain hard, persistent evidence regarding quality. I have in mind "soft expertise" industries such as psychotherapy, management consulting, political forecasting (and, of course, unconventional medicine). The effects of skill and luck are hard to disentangle in these industries. Moreover, consumers typically enter them when they face an unexpected crisis or need an ad-hoc solution to some problem, and therefore, their consumption is rarely preceded by a long learning phase. In such circumstances, consumers are more likely to rely on anecdotes such as "a friend of mine started to take vitamin pill X, and he feels much better now", or "we should trust this political analyst: he anticipated the collapse of the USSR".

This type of anecdotal thinking may cause consumers to reward firms for luck rather than skill. Indeed, such a possibility has been a major concern for financial economists who have studied the mutual funds industry (see Jensen (1968), Gruber (1996), Carhart (1997), Wermers (2000), Berk and Green (2004)). Shefrin (2002) and Rabin (2002) argue that to the extent that fund managers are rewarded for sheer luck, the explanation may be investors' tendency to extrapolate naively from small samples ("we should pick mutual fund X, because X alone beat the market last year"). Shefrin and Rabin do not, however, construct an explicit market model that incorporates this insight.

The "imperfect rationality" inherent in the $S(1)$ procedure should not be confused with ordinary imperfect information. Indeed, in Section 5 I argue that a "twin model" with imperfectly informed, rational patients yields very different results. To the extent that the procedure reflects ignorance on the patients' part, this kind of ignorance is more characteristic of early stages of a learning process, in which the model itself (rather than merely the values of its fundamentals) is still poorly known.[1]

There is a unique Nash equilibrium in the price-competition game played among the quacks. The equilibrium is symmetric and mixed. For every α, the "market for quacks" is active. Quacks act as "*charlatans*": they charge positive prices for their worthless treatments. There is a negative relation between α and expected price. In other words, the harder-to-cure the patients' disease, the worse the quacks' charlatanry. The intuition for this result is simple: as α decreases, a patient's sample is less likely to contain multiple successes, and this weakens competitive pressures.

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 w.r.t n. 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 assumption that treatments are exogenous and statistically independent is quite restrictive, especially if we want to extend the scope of the modeling approach to industries such as forecasting or mutual funds. In Section 3 I endogenize treatments: healers cannot alter their success rate, but they can control their correlation with other treatments. In equilibrium, healers differentiate themselves by betting as much as possible against competitors. As the success rate decreases, there is "more room in the space of bets" for them to differentiate themselves, and this weakens competitive pressures. Therefore, the dependency of prices and consumer welfare on the model's parameters is qualitatively the same as in the basic model. If we allow for free entry, there is a negative relation between the success rate and the equilibrium number of active quacks. In other words, *the harder-to-cure the disease, the greater the number of proposed remedies.*

A major conclusion from both the basic model and its extension is that

[1] In a similar vein, the sampling component of the $S(1)$ procedure should not be confused with ordinary models of consumer search. The patient's true payoff from choosing firm i is the expected value $\alpha - p_i$, not the sample realization $x_i - p_i$ (as would be the case in a search model).

merely raising the number of competitors does not eliminate the quacks' adverse effect on consumer welfare. In Section 4, I examine other market interventions, which would push the quacks out of the market in a more standard model: *(i)* raising the success rate of one healer, turning him from a "quack" into a "real expert"; *(ii)* allowing healers of diverse quality to disclose their success rates. In the first case, the quacks' equilibrium behavior and adverse welfare effects remain unchanged. In the second case, all healers (regardless of their quality) choose not to disclose their success rates. The lesson is simple: without lifting consumers' quality of reasoning beyond the anecdotal level, ordinary competition policies may be ineffective.

Related Literature

Osborne and Rubinstein (1998) introduced the $S(1)$ procedure and analyzed games in which all players choose according to this procedure. Their main concern was to construct a game-theoretic equilibrium concept that is appropriate for such situations. In contrast, in the present paper the $S(1)$ is employed by non-strategic agents (the patients), whereas the strategic agents (the healers) are rational. Therefore, the patients' choice procedure does not call for a non-standard equilibrium concept. Osborne and Rubinstein (2003) study a variant on "$S(1)$-equilibrium" in the context of a strategic voting model: each voter chooses his vote as a best-reply to a small sample, regarding his sample as a prediction of the distribution of votes.

The $S(1)$ procedure is related to other departures from standard probabilistic reasoning. Tversky and Kahneman (1971) demonstrated experimentally that people tend to draw exaggerated inferences from small samples, and called this tendency "the law of small numbers". Tversky and Kahneman explained this phenomenon as a consequence of the "*representativeness*" heuristic: people expect a small sample to have the same shape as the underlying probability distribution from which it is drawn. Rabin (2002) proposed a formal model of inference by "believers in the law of small numbers". I discuss Rabin's model in greater detail in Section 5.

The $S(1)$ procedure reflects an extreme version of "the law of small numbers": 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. However, it would be inaccurate to claim that the $S(1)$ procedure is exclusively a model of "the law of small numbers". Rather, it is a model of anecdotal reasoning which may have other sources, such as lack of market experience.

The $S(1)$ procedure is also linked to the model of "case-based reasoning" due to Gilboa and Schmeidler (2001), in which decision makers evaluate an action by recalling its performance in past "cases". Their emphasis, however, is on the question of similarity between past and current cases. To my knowledge, neither Rabin's nor Gilboa and Schmeidler's models have been explicitly incorporated into I.O. models.

This paper belongs to a small literature that has begun to study market interactions between rational firms and agents with boundedly rational perceptions of the market environment. Thadden (1992) studies a repeated buyer-seller interaction, when the buyer uses a non-strategic learning rule to update his beliefs regarding the quality of the seller's good. Given this learning rule, the buyer is not exploited by the seller in the long run. Rubinstein (1993) analyzes monopolistic behavior when consumers differ in their ability to understand complex pricing schedules. Piccione and Rubinstein (2003) studies intertemporal pricing when consumers have heterogeneous ability to perceive intertemporal patterns. Fishman and Hagerty (2003) study voluntary disclosure by firms, when some consumers are unable to understand the content of the disclosure (yet they can draw correct Bayesian inferences from the disclosure decision itself). Chen, Iyer and Pazgal (2002) analyze a model of price competition when consumers have memory imperfections that constrain their ability to conduct market search.

From a different perspective, the paper is related to recent attempts to introduce elements of non-Bayesian reasoning into economic modeling. Eyster and Rabin (2002) analyze players in a Bayesian game, who optimize against the opponents' statistical distribution of actions, rather than their actual Bayesian-game strategies. This model captures a phenomenon called "the attribution fallacy", namely the tendency to assume that other people's behavior is dictated by their personality, rather than the objective situation. Eyster-Rabin use their concept to illuminate the "winner's curse" fallacy. 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". Jehiel applies this behavioral model to familiar finite-horizon games. 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. This belief-selection criterion leads to an "anti folk theorem". These papers are close in spirit to the present one, in that they replace "proper" Bayesian reasoning with more naive inferences from observations.

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 the same probability α. 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. Because healers' success rate is the same as the default rate, I will refer to them as "*quacks*".

Patients choose according to the following procedure, called $S(1)$. 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.[2] 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 $x_i - p_i$.

The quacks take into account the patients' choice procedure when calculating their profits. For example, if $p_1 > p_j$ for every $j > 1$, then quack 1's profits are equal to $p_1 \cdot \alpha \cdot (1 - \alpha)^n$, because the quack'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 quack 1's profits are equal to $p_1 \cdot \alpha \cdot (1 - \alpha)$, because the quack's clientele consists of all the patients who heard a good anecdote about him and a bad anecdote about the default.

Quacks are allowed to use mixed strategies. However, once a price p_i has been realized, quack i is committed to it as far as the patients are concerned.

[2] I employ this particular tie-breaking rule merely to simplify the writing of proofs. It is not necessary for the results.

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 quack 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. For instance, consider the assumption that patients sample *every* quack. It would be more realistic to assume that patients get to hear anecdotes about a subset of quacks. One way to deal with this artificiality is to assume that the number of quacks is infinite, yet the patients can sample only n quacks and he is unaware of quacks whom he did not sample. None of the results in this paper would change under this variant.[3] Another artificial feature is the assumption that the number of observations about quack i is independent of the size of the quack's clientele. Alternatively, we could assume that patients hear more anecdotes about quacks with a larger clientele. This variant would be more difficult to analyze. The endogenization of the number of anecdotes means that this number is a fixed point. The patient's exact sample induces some probabilistic choice, which in turns induces the number of anecdotes per alternative. It is not clear that such a fixed point necessarily exists.

The game induced by the basic model is *formally equivalent* to the following, more conventional model, in which firms compete in prices over consumers with independently drawn, private values. In such a model, let v_i denote the patient's valuation of alternative i, for every $i = 0, 1, ..., n$. The v_i's are independently drawn, and take the value 1 (0) with probability α $(1 - \alpha)$. The v_i's are the patient's private information. A model of this sort, albeit with continuous distributions and without an outside option, was studied by Perloff and Salop (1985).[4] Thus, $S(1)$-patients in the market for quack behave as if they were rational agents with private, independent valuations of the $n + 1$ alternatives.

What are the merits of the market-for-quacks interpretation, in light of this formal equivalence? First, the insights we shall obtain into the question of whether the market rewards experts for sheer luck are entirely due to the market-for-quacks interpretation. Second, this interpretation will lead us to examine *extensions* of the basic model, which would be hard to con-

[3] Under this interpretation, an increase in n cannot be interpreted as market entry, but as an increase in the patients' awareness of available treatments.

[4] Gabaix and Laibson (2004) study asymptotic properties of the Perloff-Salop model, and suggest that consumers' noisy valuations may be a consequence of bounded rationality.

ceive, and even meaningless, under the private-values interpretation. Finally, the market-for-quacks interpretation will be radically different in its *welfare* implications from those of the private-values interpretation.

Let us now turn to equilibrium analysis.

Proposition 1 *There is a unique Nash equilibrium in the price-competition game played among the quacks. Every quack plays the mixed strategy given by the c.d.f:*

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

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

To see the origin of expression (1), suppose that we restricted attention to symmetric equilibria. Let G denote the equilibrium *c.d.f.* For the sake of the argument, assume that G is atomless. Then, for every price p in the support of G, the quacks' payoffs is given by the expression:

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

because for every quack i, $\alpha(1-\alpha)$ 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. 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 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 derivation brings to mind similar characterizations in the literature on equilibrium price dispersion (e.g., Butters (1977), Varian (1980), Burdett and Judd (1983) and Rob (1985)). In this literature, firms are involved in price competition in an imperfectly competitive market, where the imperfection is due to search costs. These works were not strictly game-theoretic, and their formalism essentially assumed that the equilibrium is symmetric. By contrast, Proposition 1 proves that the symmetric equilibrium is the *unique* Nash equilibrium. In this technical sense, the result goes beyond these earlier works.

Corollary 1 *Quacks' expected equilibrium price is strictly decreasing with α. In particular,$E(p) \to 0$ as $\alpha \to 1$ and $E(p) \to 1$ as $\alpha \to 0$.*

Quacks behave as *charlatans* in equilibrium: 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, as the patients' condition becomes more "hopeless", the quacks' charlatanry becomes more pronounced. The intuition behind the comparative statics is simple. As α decreases, the a patient's sample is less likely to contain multiple successes. This weakens competitive pressures and causes prices to go up.

At the same time, a decrease in α causes aggregate demand for quacks (given by the expression $1 - \alpha - (1 - \alpha)^{n+1}$, which is the probability that a patient heard a bad anecdote about the default and at least one good anecdote about a quack) to shrink. As α approaches zero, market equilibrium tends to a state of *monopolistic competition*: every quack faces a demand which is virtually insensitive to competitors' prices, and his profits approach zero. As α approaches one, market equilibrium tends to the rational-patients benchmark. For every α, quacks have a positive clientele. Thus, the "market for quacks" is always active in equilibrium.

A word of caution concerning the validity of comparative statics w.r.t α is due. This exercise is based on the presumption that the patients' choice procedure remains fixed as α changes. If the $S(1)$ procedure is viewed as a solution to some optimization problem that takes α as an input, this presumption is untenable. Therefore, the comparative statics exercise is valid only in so far as the patients' choice procedure itself is independent of the parameters.

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, industry profits are $n\alpha(1 - \alpha)^n$. Because quacks do not contribute any added value, this expression also represents the *welfare loss that quacks inflict on patients in equilibrium.* This expression is not monotonic in α: it attains an maximum at $\alpha^* = \frac{1}{n+1}$. it also behaves non-monotonically w.r.t n. For every α, the number of quacks that maximizes the patients' welfare loss is $n^* = -\frac{1}{\ln(1-\alpha)}$. For every $\alpha \lessapprox 0.4$, $n^* \geq 2$. 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.[5]

The intuition for the comparative statics w.r.t n is simple. On one hand, a greater number of quacks increases the incentive to cut 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 patient has a larger set of available 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. Note, however, that if we allow free entry, the welfare loss vanishes.

For any fixed number of quacks, 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 quacks' equilibrium payoffs. Thus, competition among quacks 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.*

3 Endogenous choice of treatments

In the basic model, the quacks' treatments are exogenously given and statistically independent. This assumption is too restrictive for industries such as forecasting or mutual funds, in which firms have some control over the structure of their bets. In this section I enrich the model by enabling quacks to choose their treatments - in particular, the correlation between their treatment and other alternatives.

[5] The patients' welfare loss can be substantial: for every $\alpha < \frac{1}{2}$, there exists $n \geq 2$, such that the patients' loss exceeds $\frac{1}{4}$. As $\alpha \to 0$, the maximal welfare loss converges to $\frac{1}{e}$.

Let $C = [0,1]$ be a set of *cases* and let $A = \{0, 1, ..., m\}$ be a set of *actions*. A *treatment* is a function $t : C \to A$. A strategy for quack i is a pair (p_i, t_i), where $p_i \in [0,1]$ is the price he charges and t_i is the treatment he offers. The default option (p_0, t_0) satisfies $p_0 = 0$ and $t_0(c) = 0$ for every $c \in C$. A *state* is a pair (c, a), where $c \in C$ and $a \in A$. If a patient acquires treatment t_i, he recovers in state (c, a) if and only if $t_i(c) = a$. Thus, a state consists of a case and the action that leads to recovery in this case. Assume that the state is drawn uniformly. It follows that the probability of recovery is $\frac{1}{m+1}$, independently of the treatment the patient selects.

Patients choose according to the $S(1)$ procedure: they sample every alternative once, and choose the one that maximizes $x_i - p_i$ in their sample, where x_i is the outcome of treatment i in their sample. Let $x_i = 1$ (0) if the outcome is a success (failure). Apply the same tie-breaking rule as in Section 2. This description contains an ambiguity. That is, we need to choose whether: *(i)* the state is drawn independently for each alternative; *(ii)* the state is fixed across the patient's sample. The first assumption reduces the formalism to the basic model of Section 2, with $\alpha = \frac{1}{m+1}$. We shall adopt the second assumption: the patients' collection of anecdotes fixes a case at random and documents the quacks' performance in that particular case.

The forecasting industry provides a natural illustration of this model. Forecasters try to predict the outcome $a \in A$ in every possible situation $c \in C$. By assumption, they are no better in predicting the future than laymen. Consumers ultimately care about the forecasters' true success rate. However, they are bounded in their ability to estimate it, and therefore reason anecdotally. They recall some situation at random, and choose the cheapest forecaster among those who happened to anticipate the outcome of that particular situation.

The quacks take into account the patients' choice procedure when calculating their profits. For example, consider a strategy profile $(p_i, t_i)_{i=1,...,n}$, in which all quacks charge distinct prices. Then, the payoff of quack i is $p_i \cdot A$, where A is the measure of states (c, a) for which $t_i(c) = a$ and $p_j > p_i$ for all other alternatives $j \neq i$ for which $t_j(c) = a$. Note that if all quacks choose the same treatment, then the quack who charges the highest price in the market has no clientele. This quack has an incentive to deviate to another treatment, which assigns different actions to some cases. In this way he will be able to attract some clientele.

Let us turn to characterizing Nash equilibrium in the game played among the quacks. *I allow for mixtures only in the price component,* due to the complexity of the treatment component. Given an equilibrium, define B_i as the set of cases c for which $t_i(c) \neq t_j(c)$ for every $j \neq i$. Let μ_i denote the measure of B_i.

Proposition 2 *If $n \leq m$, then in Nash equilibrium each firm chooses $p = 1$ with probability one, and $\mu_i = 1$ for every quack i.*

Proposition 3 *If $m < n < 2m$, then in Nash equilibrium every quack plays a mixed pricing strategy given by the c.d.f:*

$$G(p) = \frac{p - \mu}{p - p\mu} \tag{4}$$

defined over the support $[\mu, 1]$, where:

$$\mu = 2 - \frac{n}{m}$$

Moreover, for every quack i, $t_i(c) = 0$ for measure zero of cases c; and for every triple of distinct quacks i, j, k, $t_i(c) = t_j(c) = t_k(c)$ for measure zero of cases c.

Proposition 4 *If $n \geq 2m$, then in Nash equilibrium all firms play $p = 0$ with probability one, and $\mu = 0$ for every quack i.*

Corollary 2 *The quacks' equilibrium expected price increases with m and decreases with n.*

The quacks' equilibrium pricing behavior is qualitatively the same as in the basic model. In particular, there is a negative relation between expected equilibrium price and the success rate that characterizes the industry. The intuition is that as m increases, the state space contains "more room" for quacks to differentiate themselves by betting against competitors, and therefore competitive forces are weaker. Notice that unlike the basic model, expected equilibrium price does not vary smoothly with n. First, for low values of n, quacks choose the monopolistic price. Second, there is a finite critical number of firms n^*, such that for every $n > n^*$, prices are equal to zero.

Propositions 2-4 captures the intuition that *when quacks are able to choose their treatments - and totally unable to affect their success rate - they prefer to differentiate themselves by betting as much as possible against competitors*. Recall that according to the $S(1)$ procedure, the state (c, a) is randomly drawn and held fixed throughout a patient's sample. Therefore, given the state (c, a), the patient considers only the treatments t_i for which $t_i(c) = a$, and then chooses the cheapest among them. Since the default is the cheapest alternative, the quack prefers to bet against the default - i.e., $t_i(c) \neq 0$ for every $c \in C$.

For similar reasons, quacks prefer to compete against as few competitors as possible. Indeed, a key step in the proof of Proposition 3 is that as long as the industry is profitable (i.e., in the range $n < 2m$), no more than two quacks ever bet on the same action in a given case. This property means that the quacks' equilibrium treatments are as different from each other as possible. Let us sketch the reasoning behind this result. Consider a case c in which quacks' bets are the most concentrated - i.e., there exists no other case which belongs to more B_i's than c. If there are at all cases in which more than two quacks bet on a particular action, then c surely belongs to this class of cases.

Suppose that quack i belongs to a group of at least three quacks who bet on a particular action in c. Suppose that $c \in B_h$. Then, I argue that $B_i \backslash B_h$ must have measure zero. Assume the contrary - i.e., that there is a positive measure of cases c' which belong to B_i but not to B_h. Then, there is a quack j who bets on the same action as h in c', whereas i alone bets on some other action in c'. If quacks who bet like i in c do not switch from $t_i(c)$ to $t_h(c)$, this is only because quack h is more expensive (probabilistically) than quack i. This is a consequence of the assumption that there are at least two other quacks who bet like i in c. But if quack h is more expensive (probabilistically) than quack i, then the quacks who bet like h in c' necessarily want to switch from $t_h(c')$ to $t_i(c')$. Therefore, $B_i \backslash B_h$ has measure zero. But this in turn implies that there is a positive measure of cases c'' which belong not only to the same B_i's that c belongs to, but also to B_h. This contradicts the definition of c as a case that belongs to as many B_i's as any other case.

Let us turn to welfare analysis. As in the basic model, the patients' welfare loss is equal to industry profits.

Corollary 3 *The patients' welfare loss in Nash equilibrium is:*
(i) $\frac{n}{m+1}$ *for* $n \leq m$.
(ii) $\frac{n}{m+1} \cdot (2 - \frac{n}{m})$ *for* $m < n < 2m$.
(iii) 0 *for* $n \geq 2m$.

As in the basic model, the patients' welfare loss behaves non-monotonically in n and m. The maximal welfare loss is $\frac{m}{m+1}$, attained when $m = n$.

Suppose that we allowed for free entry. As usual in I.O. modeling, a zero-profit condition would characterize equilibrium. The model would then predict that the number of active quacks is $2m$. Thus, *the harder-to-cure the disease, the larger the number of supplied remedies*. In the context of the forecasting industry, the harder it is to make an accurate prediction in a field, the larger the number of "experts" who will try their hand at it.[6]

Possible implications for the mutual-funds literature

A recent theoretical study of mutual funds (Berk and green (2004)) begins with these words:

> One of the central mysteries facing financial economics is why financial intermediaries appear to be so highly rewarded in our economy, despite the apparent fierce competition between them and the uncertainty about whether they add value through their activities. Research into mutual fund performance has provided evidence that deepens this puzzle. Since Jensen (1968), studies have shown little evidence that mutual fund managers outperform passive benchmarks. Recent work has produced several additional findings. The relative performance of mutual fund managers appears to be largely unpredictable from past relative performance. Nevertheless, mutual fund investors chase performance. Flows into and out of mutual funds are strongly related to lagged measures of excess returns (see Chevalier and Ellison (1997); Sirri and Tufano (1998)).

[6] This comparative statics exercise should be taken with an extra grain of salt. There probably exist additional equilibria, which involve mixtures in the treatment dimension, a possibility which I have ruled out.

Berk and Green attempt to account for these anomalies with a rational-choice model. In contrast, Shefrin (2002, Chapter 12)) argues informally that the anomalies may be due to investors' tendency to extrapolate naively from small samples. In a similar vein, Rabin (2002) applies his model of inference by "believers in the law of small numbers" to individual choice of mutual funds by investors. The model presented in this section is in the spirit of Shefrin-Rabin argument. While the model is far too stylized to be considered as a descriptive theory of the mutual funds industry, it may contain the seeds of a future theory, based on the assumption that investors use anecdotal reasoning to evaluate funds.

To substantiate this claim, let us interpret C as a set of "states of the economy", and A as a set of possible sectors. For every $c \in C$, there is exactly one sector which makes for a successful investment. A mutual fund i chooses its fee p_i and its investment policy t_i - i.e., a selection of one sector for every state of the economy. The default option may be viewed as a passive (index) fund. The assumption that (z, a) is drawn according to the uniform distribution implies that actively managed funds have no advantage over the passive fund. The $S(1)$ procedure has a natural interpretation in this environment. From time to time, individual investors get to adjust their choice of fund. When they do so, they evaluate funds according to their most recent performance, which is associated with a particular state of the economy.

The anomalies noted by Berk and Green are directly built into the model and taken for granted. The assumption that investors reason anecdotally immediately explains why actively managed funds may be rewarded for luck rather than skill. The model goes further and produces additional predictions. First, investment policies are as diffuse as possible, with different funds "specializing" in different sectors. Second, for a fixed number of funds, there is a negative relation between the chances of a successful investment and the market fees that characterize the market. Third, in the long run there is a negative relation between the number of active funds and the chances of successful investment that characterizes the market.[7]

I wish to re-emphasize that the model omits a number of crucial elements

[7] The model's assumptions rule out diversification. A fully diversified portfolio would yield a payoff of $\frac{1}{m+1}$ with certainty. I conjecture that if we endow funds with the option of full diversification, the funds will choose to abstain from this option. Also, if we assume that the outside option is fully diversified, I conjecture that the results will change only slightly.

to count as a descriptive theory of the mutual funds industry. Moreover, the comparative statics exercise is severely handicapped by the fact that in reality, the success rate that characterizes investments in the market is endogenous, whereas in the model it is exogenous. Nevertheless, I hope that the model might offer a starting point for a successful theory.

4 Two market interventions

We saw in previous sections that increasing the number of competitors does not necessarily curb the quacks' adverse effects on patients' welfare. As long as n is not too large, the patients' welfare loss increases with n. This section examines two additional market interventions, which would normally be considered as effective competition policies. In contrast, *given the patient's behavioral model, these interventions turn out to be totally ineffective.*

4.1 Replacing a quack with an expert

In this sub-section, I perturb the basic model of Section 2 by replacing one of the quacks with a high-quality healer. The question is whether this intervention will crowd out the quacks, or at least alleviate the welfare loss that they inflict on patients. 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 other healer $i \neq e$ has a success rate α. In other words, healer e is an "expert" while his opponents are "quacks".

If patients knew the market model, then clearly the expert would crowd out the quacks in equilibrium. When patients choose according to the $S(1)$ procedure, we get a very different result:

Proposition 5 *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.*

Turning a quack into an expert does not affect his competitors' equilibrium behavior and performance. The expert ends up luring patients away from the default, not from the quacks. As a result, the patients' welfare loss caused by the $n-1$ quacks remains unaffected.

To get the intuition for this result, suppose that in equilibrium, all G_i's share the same support $[p_L, 1]$. Intuitively, the presence of an expert instead of a quack could not have led the other quacks to raise their prices. Therefore, we do not expect their pricing strategy to place an atom on $p = 1$. The expert's payoff from the monopolistic price $p = 1$ is thus $\alpha_e \cdot (1-\alpha)^n$. But his payoff from p_L is $p_L \cdot \alpha_e \cdot (1-\alpha)$, hence $p_L = (1-\alpha)^{n-1}$. But because p_L also belongs to the support of the quacks' strategies, this means that quacks earn a payoff of $\alpha(1-\alpha)^n$, just as in the basic model. By definition, this is the welfare loss that an individual quack inflicts on patients.

The identity of the supports of the healers' strategies is a consequence of mixed-strategy equilibrium reasoning. The condition for the expert's indifference among all prices in the support of G_e is independent of α_e: it is only expressed in terms of the opponents' success rates and pricing strategies. Therefore, it is the same condition as in the basic model, and it yields the same pricing strategy for the quacks as in the basic model. But this implies that the expert will be indifferent among the same set of prices as in the basic model.

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' anecdotal inferences, although to a lesser degree than the quacks.

4.2 Disclosure of success rates

In the basic model of Section 2, healers have no control over the patients' knowledge. In this sub-section, I perturb the model by assuming 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 is appropriate to allow more general market primitives. Denote the success rate 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 evaluation of the quality of treatment 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.

In the previous models we have considered, the meaning of the $S(1)$ procedure was that patients reason anecdotally about the quality of alternatives. In the context of the model of this sub-section, it also means that patients infer nothing from the healer's disclosure policy itself. In particular, they do not realize that a healer's failure to reveal his quality signals that his quality is relatively low. Thus, the exact sense in which the agents' procedure departs from standard rationality may vary with the model in which embed such agents. This seems to be a general property of models with procedurally rational agents.

Standard adverse selection model with rational, imperfectly informed patients typically assume that the patients know the distribution of success rates, but do not know ex-ante the success rates of individual healers. In sequential equilibrium of such models, every healer would disclose his success rate (except possibly the lowest types).[8] Given our model of the patients' behavior, the result is the complete opposite:

Proposition 6 *For every p, the strategy (p, Y) for healer i is weakly dominated by some other strategy (p', N).*

Thus, given the patients' choice procedure, healers have an incentive not to reveal their success rate, even when they are of the highest quality. 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 $r_i = Y$.)

[8] See, for example, Milgrom and Roberts (1986, Section 2). 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 the full disclosure result breaks down in this case.

Proposition 6 establishes that type revelation is a weakly dominated strategy. In fact, it can be verified that *type disclosure can never be part of Nash equilibrium.* I omit the proof of this result for the sake of brevity. The lesson from these results is that enabling healers to reveal their type is ineffectual when patients choose according to the $S(1)$ procedure.

It is interesting to compare the results of this sub-section to Milgrom and Roberts (1986, Section 3). In their model, consumers are strategically unsophisticated, in the sense of being unable to draw inferences from what firms choose not to reveal. However, they are probabilistically sophisticated: they draw Bayesian inferences from the content of the firms' disclosure. Milgrom and Roberts show that in equilibrium, the full information outcome is attained, thanks to competitive forces. This comparison demonstrates that the no-disclosure result of this sub-section is not merely a consequence of their lack of strategic sophistication.

5 Concluding remarks

Imperfect rationality or imperfect information? Although the models presented in this paper are simple, the modeling procedure they embody is unusual. Our starting point is a standard price-competition model with complete information. Economists typically depart from such a simple benchmark by perturbing its informational structure, without abandoning the meta-level assumption that "the model itself is common knowledge". Instead, in this paper we relaxed the *rationality* of consumers' choice with respect to the complete-information model. The source of the relaxation may be that consumers are inherent bounded rationality, or that they lack opportunities to learn the model.

The question naturally arises, whether the basic model and its various extensions could be "rationalized", in the sense that the same results could be obtained from a price-competition model with imperfectly informed consumers. In other words, *can we replace imperfectly rational patients with imperfectly informed patients, and get the same results*?

Let us begin with the simplest attempt to rationalize the basic model of Section 2. Suppose that the healers' success rate is drawn from some prior distribution over $\{\alpha_L, \alpha_H\}$. The distribution of types is commonly known. Healers know their own success rates, whereas patients observe partially informative signals.

Such a model cannot yield exactly the same results as our model. Recall that in the model of Section 2, patients behave as if they are 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 a "jump" to these extreme posteriors (the low posterior is above 0 or the high posterior is below 1). Thus, although equilibrium strategies will be mixed in the manner of Proposition 1, it will be impossible to reproduce expression (1). More importantly, the two models have different *comparative statics* with respect to the industry's average success rate. In the model of Section 2, equilibrium prices decrease with α. In contrast, in the alternative model proposed here, if we raise α_L and α_H by the same factor, equilibrium prices will increase by this factor as well.

When we turn to some of the extensions of the basic model, the disparity between the two modeling approaches widens. In the model of Section 3, treatments are endogenous, and the expected outcome of selecting a quack is only a function of the treatment that the healer chose. If we tried to rewrite the model as a rational-choice model, then in equilibrium patients would know the healers' strategies. It is hard to see how one could reconcile such equilibrium knowledge with the behavior we attempt to rationalize, without modifying the underlying market situation beyond recognition. In the model of Section 4.2, healers are allowed to reveal their type. We have already observed the contrast between the no-revelation result we obtain and the full-revelation result obtained in a standard model with imperfectly informed patients (e.g., Milgrom and Roberts (1986, Section 2)). Similar differences between the two approaches will emerge in any model in which healers can signal their type.

Extension of the decision procedure. The $S(1)$ procedure captures an extreme case of anecdotal reasoning: 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 evaluates the success rate of treatment i using its average performance in the sample. Thus, patients form an unbiased "point estimate" of the success rate associated with each alternative, but they behave as if there is no sampling error. As K gets larger, the sampling error decreases, and in the limit, the patient's procedure converges to the rational-choice benchmark.

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

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, and that they may even know the success rate associated with the default. Therefore, it makes sense to consider a variant on the model, in which $x_0 = \alpha$ with probability one. The patients form quality assessments of healers as in the basic model.

The essential features of our equilibrium characterization - uniqueness, symmetry, price dispersion and the comparative statics - remain unchanged under this modification. Only fine details have to be modified: the "monopoly price" becomes $1-\alpha$ 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 basic model. The reason patients experience a loss at all is that they compare an alternative they are highly familiar with (the default) with alternatives they know only through anecdotal evidence, without appreciating the difference in informational content between knowing the success rate and hearing an anecdote.

Relaxing quackery. Our analysis is easily extendible to the case in which the default success rate is $\alpha_0 < \alpha$. With standard rational patients, the model is reduced to standard Bertrand competition, such that equilibrium prices and profits are zero, and the patients' utility is $\alpha - \alpha_0$. By comparison, with $S(1)$-patients, Proposition 1 continues to hold. The reason is simple: the default success rate enters the healers' payoff function through the multiplicative term $1-\alpha_0$, and it cancels out as we derive the expression for G. Therefore, the healers' behavior is independent of α_0. The welfare analysis is modified. For instance, when $\alpha_0 = 0$, the patients' net payoff is $\alpha - n\alpha(1-\alpha)^{n-1}$. It follows from this expression that there is a net welfare loss if α is sufficiently low. Similar conclusions hold for the extended models.

References

[1] Berk J. and R. Green (2004), "Mutual Fund Flows and Performance in Rational Markets", *Journal of Political Economy* **112**, 1269-1295.

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

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

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

[5] Carhart M. (1997), "On Persistence in Mutual Funds Performance", *Journal of Finance* **52**, 57-82.

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

[7] Chevalier, J. and G. Ellison (1997), "Risk Taking by Mutual Funds as a Response to Incentives", *Journal of Political Economy* **105**, 1167–1200.

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

[9] 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.

[10] Gabaix X. and D. Laibson (2004), "Gaussian Noise, Consumer Confusion and Asymptotic Markups", Mimeo.

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

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

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

[14] Jensen M. (1968), "The Performance of Mutual Funds in the Period 1945–1964", *Journal of Finance* **23**, 389–416.

[15] Milgrom P. and J. Roberts (1986), "Relying on Information of Interested Parties", *Rand Journal of Economics* **17**, 18-32.

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

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

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

[19] 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.

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

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

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

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

[24] Shefrin H. (2002), *Beyond Greed and Fear: Understanding Behavioral Finance and the Psychology of Investing*, Oxford University Press.

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

[26] Sirri E. and P. Tufano (1998), "Costly Search and Mutual Fund Flows", *Journal of Finance* **53**, 1589-1622.

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

[28] Thadden E. (1992), "Optimal Pricing against a Simple Learning Rule", *Games and Economic Behavior* **4**, 627-649.

[29] Tversky A. and D. Kahneman (1971), "Belief in the Law of Small Numbers", *Psychological Bulletin* **2**, 105-110.

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

[31] 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.

Appendix: Proofs

Proof of Proposition 1

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

Step 1. For every quack i, G_i is continuous over $[0, 1)$.

Proof. 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 quacks i assigns a positive measure to a price that yields zero profits. As we noted in Section 2, the quacks' max-min payoff is $\alpha(1-\alpha)^n > 0$. Therefore, the quack can profitably deviate by shifting this measure to some $p > 0$. Now suppose that $p \in (0, 1)$. If every other quack assigns no weight to the interval $(p, p+\varepsilon)$, then quack i can profitably deviate by shifting the atom from p to $p + \varepsilon/2$. And if some quack $j \neq i$ assigns weight to the interval $(p, p+\varepsilon)$ for arbitrarily small ε, then there exists $\delta > 0$ such that quack j can profitably deviate by shifting this weight to $p - \delta$.

In the remainder of the proof, we shall rely on two additional observations. First, if G_i has an atom on $p = 1$, then no other G_j has an atom on $p = 1$. Otherwise, either of these quacks would be able to deviate profitably by shifting this atom slightly downward. Second, if s_i assigns a positive weight to an interval $(p, p + \varepsilon)$ or $(p, p - \varepsilon)$ for some $p \in (0, 1)$ and $\varepsilon > 0$, then p maximizes quack i's expected payoff against s_{-i}. This is a standard result which follows from Step 1.

Let p_i^L and p_i^H denote the infimum and supremum of the support of G_i. Define $p^L = \min\{p_1^L, ..., p_n^L\}$ and $p^H = \max\{p_1^H, ..., p_n^H\}$.

Step 2. $p^H = 1$.

Proof. Assume that $p^H < 1$. Then by Step 1, none of the G_i's contain an atom on p^H. It follows that the payoff of the quack who charges p^H is $p^H \cdot \alpha(1 - \alpha)^n$, which is below the max-min payoff, a contradiction.

Step 3. All quacks earn the same payoff in equilibrium.

Proof. Assume the contrary, and suppose (w.l.o.g) that quack 1 earns a higher payoff than quack 2. Suppose that quack 2 deviates by playing p_1^L with probability one. Quack 1's payoff before the deviation is:

$$p_1^L \cdot \alpha \cdot (1 - \alpha) \cdot \Pi_{j>1}[1 - \alpha G_j(p_1^L)] \tag{5}$$

whereas quack 2's payoff after the deviation is:

$$p_1^L \cdot \alpha \cdot (1 - \alpha) \cdot \Pi_{j>2}[1 - \alpha G_j(p_1^L)] \tag{6}$$

and since the second expression is at least as high as the first expression, quack 2's deviation is profitable.

Step 4. The quacks' equilibrium payoff is $\alpha(1 - \alpha)^n$.

Proof. By Step 2, $p^H = 1$. We have observed that there exists a quack i whose competitors do no place an atom on $p = 1$. This quack's payoff is $\alpha(1 - \alpha)^n$. By Step 3, all quacks earn this payoff.

Step 5. No quack places an atom on $p = 1$.

Proof. Suppose that quack i places an atom on $p = 1$. We have observed that no other quack places an atom on $p = 1$. Suppose that the second-highest p_j^H is strictly lower than one. Then, p_j^H does not maximize quacks j's payoff. It follows that $p_j^H = 1$. But since quack i places an atom on $p = 1$, quacks i and j earn different payoffs, in contradiction to Step 3.

Step 6. $p_i^L = p^L$ for all quacks i.

Proof. Assume the contrary, and suppose (w.l.o.g) that $p_2^L = p^L$ and $p_1^L > p^L$. Suppose that quack 2 deviates by playing p_1^L with probability one. Expressions (5) and (6) represent quack 1's payoff before the deviation and quack 2's payoff after the deviation, respectively. Because $G_2(p_1^L) > 0$, expression 6 is higher than expression 5. By Step 3, quacks 1 and 2 earn the same payoff prior to the deviation. Therefore, the deviation is profitable.

Step 7. For every quack i, G_i is strictly increasing in $[p_i^L, p_i^H]$.

Proof. Assume the contrary, and suppose (w.l.o.g) that G_1 is flat over some interval $(p, p') \subset [p_1^L, p_1^H]$. By Step 6, $p_1^L = p_2^L$. Then, there must be some other quack (denoted 2, w.l.o.g) who assigns positive weight to the neighborhood of p - otherwise, p would not maximize quack 1's payoff. The two quacks' payoff from the prices p and p' are:

$$\begin{aligned}
\pi_1(p) &= p \cdot \alpha \cdot (1-\alpha) \cdot [1 - \alpha G_2(p)] \cdot \Pi_{j>2}[1 - \alpha G_j(p)] \\
\pi_2(p) &= p \cdot \alpha \cdot (1-\alpha) \cdot [1 - \alpha G_1(p)] \cdot \Pi_{j>2}[1 - \alpha G_j(p)] \\
\pi_1(p') &= p' \cdot \alpha \cdot (1-\alpha) \cdot [1 - \alpha G_2(p')] \cdot \Pi_{j>2}[1 - \alpha G_j(p')] \\
\pi_2(p') &= p' \cdot \alpha \cdot (1-\alpha) \cdot [1 - \alpha G_1(p')] \cdot \Pi_{j>2}[1 - \alpha G_j(p')]
\end{aligned}$$

By Step 3, $\pi_1(p) = \pi_2(p)$. Therefore, $G_2(p) = G_1(p)$. By assumption, $G_1(p) = G_1(p')$, whereas $G_2(p') > G_2(p)$. Therefore, quack 2 can profitably deviate by playing p' with probability one.

Step 8. The equilibrium is symmetric, and the equilibrium strategy is given by expression (1).

Proof. By Step 6, $p_i^L = p^L$ for every quack i. Denote $p^* = \min\{p_1^H, ..., p_n^H\}$. By Step 7, all the G_i's are strictly increasing in $[p^L, p^*]$. By Step 4, all quacks earn a payoff of $\alpha(1-\alpha)^n$. Therefore, for every quack i and every price $p \in [p^L, p^*]$:

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

We have a system of n equations in n variables $G_j(p)$. The equations are symmetric and the R.H.S. is strictly decreasing in the $G_j(p)$'s. Therefore, the system has a unique solution, which is symmetric. In particular, it follows that $p_i^H = p^*$ for every quack i. It is now straightforward to derive expression (1).

Proof of Proposition 2

When $n \leq m$, quack i can set $t_i(c) \neq t_j(c)$ for every $j \neq i$. In this way, he does not affect the chances of success but he reduces competition. Thus, in equilibrium, $\mu_i = 1$. Given this, he can set $p = 1$.

Proof of Proposition 3

The proof proceeds stepwise. borrow the definitions of p_i^L, p_i^H, p^L, p^H from the proof of Proposition 1. Also, as in the proof of Proposition 1, I shall take two facts for granted. First, the quacks' pricing strategies contain no atoms on any $p < 1$. Second, for this reason any price in the support of G_i maximizes quack i's payoff against G_{-i}.

Step 1. All quacks earn the same equilibrium payoff, which is strictly positive.

Proof. Let us first show that at least one quack earns a positive payoff. Because $n < 2m$, for every case $c \in C$ there is at least one quack i, for whom $t_i(c) \neq t_j(c)$ for every $j \neq i$. Therefore, there is at least one quack i for whom $\mu_i > 0$. This quack can secure a strictly positive payoff by setting $p_i > 0$. Now suppose that quack i earns in equilibrium a higher payoff than quack j. Suppose that quack j deviates to a strategy that plays p_i^L with probability one and mimics quack i's treatment. Then, quack j's payoff after the deviation is at least as high as quack i's payoff prior to the deviation, a contradiction.

Step 2. If there is a positive measure of cases c for which $t_i(c) = t_j(c)$, then $p_i^L = p_j^L$ and $p_i^H = p_j^H$. Moreover, neither G_i nor G_j have an atom on $p = 1$.

Proof. Suppose that there is a positive measure of cases c for which $t_i(c) = t_j(c)$. Assume that $p_i^L > p_j^L$. If quack j deviates by playing $t'_j = t_i$ and $p'_j = p_i^L$ with probability one, his payoff will be higher than quack i's payoff prior to the deviation. The reason is that he mimics a strategy that was optimal for quack i prior to the deviation, except that now $p'_j < p_i$ with probability one, whereas prior to the deviation, $p'_j > p_j$ with positive probability. By Step 1, this is a profitable deviation for quack j.

The proof that $p_i^H = p_j^H$ is similar. Assume that $p_i^H > p_j^H$. If quack j deviates by playing $t'_j = t_i$ and $p'_j \in (p_j^H, p_i^H)$ with probability one, his payoff will be higher than quack i's payoff prior to the deviation. The reason is that he mimics a strategy that was optimal for quack i prior to the deviation, except that now $p'_j < p_i$ with positive probability, whereas prior to the deviation, $p'_j > p_j$ with probability one. By Step 1, this is a profitable deviation for quack j.

Finally, suppose that G_i places an atom on $p = 1$. By the preceding argument, $p^H_j = p^H_i = 1$. However, quack j does not place an atom on p^H_j - otherwise, it would be profitable for him to deviate by shifting this atom slightly below $p = 1$. If quack j deviates by playing $t'_j = t_i$ and $p'_j = 1$ with probability one, he will increase his payoff, for the same reason we applied in the preceding argument.

Step 3. $p^H = 1$. Moreover, the quack i with $p^H_i = 1$ satisfies $\mu_i > 0$.

Proof. Consider the quack i for whom $p^H_i = p^H$. If $p^H < 1$, then the payoff attained by this price is $p^H \cdot \mu_i$. But this price fails to maximize the quack's payoff, because deviating to higher price p' would attain a payoff of $p' \cdot \mu_i$. Therefore, $p^H_i = 1$. By Step 2, if there is another quack j such that $t_i(c) = t_j(c)$ for a positive measure of cases c, neither G_i nor G_j place an atom on $p = 1$. Therefore, quack i's equilibrium payoff is equal to μ_i. By Step 1, $\mu_i > 0$.

Step 4. For every quack i, there is a measure zero of cases c for which $t_i(c) = 0$.

Proof. First, let us show that for every quack i, there is a positive probability that $p_i < p_j$ for some other quack j. Assume the contrary - i.e., that for some quack i, $p^L_i \geq p^H_j$ for every $j \neq i$. By Step 2, $t_i(c) = t_j(c)$ for zero measure of cases c, hence $\mu_i = 1$. But this means that quack i's payoff is $\frac{1}{m+1}$. By Step 1, this is the payoff that all quacks earn, which is impossible because $n > m$.

Now suppose that there is a positive measure of cases c for which $t_i(c) = 0$ for some quack i. Clearly, $p_i > p_0 = 0$ with probability one. Suppose that $p_i < p_j$ with positive probability. Then, quack i can deviate by setting $t'_i(c) = t_j(c)$ in these cases c. This would be a profitable deviation.

Step 5. For every quack i, $\mu_i > 0$.

Proof. Suppose that $\mu_i = 0$. Then, for every case c (except for events of measure zero), there is another quack j, such that $t_i(c) = t_j(c)$. By Step 2, $p^H_i = p^H_j$ and neither G_i nor G_j places an atom on $p = 1$. Therefore, the price p^H_i earns a zero payoff for quack i, contradicting Step 1.

Step 6. There is a measure zero of cases c for which $t_i(c) = t_j(c) = t_k(c)$ for three distinct quacks i, j, k.

Proof. For every c, we can count the number of B_i's to which c belongs. Consider a positive-measure set of cases c^* which belong to the maximal number of B_i's, given $(t_i)_{i=1,...,n}$. Denote this number by a. Because $n < 2m$, $a \geq 1$. Moreover, if there is a positive measure of cases in which more than two quacks choose the same action, then c^* necessarily belongs to this class of cases.

Suppose that $t_i(c^*) = t_j(c^*) = t_k(c^*)$, whereas by assumption, $c^* \in B_h$. By Step 2, $p_i^L = p_j^L = p_k^L$. Therefore, for every price p in the support of G_i, $G_j(p), G_k(p) > 0$. If $G_h(p) \leq G_i(p)$, then quack j can profitably deviate by switching from $t_j(c^*)$ to $t_h(c^*)$. It follows that $G_h(p) > G_i(p)$. Let us now show that this implies that the set $B_i \backslash B_h$ has measure zero. Assume the contrary. Then there is a positive measure of cases $c \in B_i \backslash B_h$, such that $c \in B_i$ and $c \notin B_h$. But this means that there is a quack l for whom $t_l(c) = t_h(c)$. By Step 2, $p_l^L = p_h^L$, and therefore there is a price p in the support of G_l, G_h and G_i, such that $G_h(p) > G_i(p)$.Therefore, quack l can deviate profitably by switching from $t_l(c) = t_h(c)$ to $t_i(c)$.

By Step 5, the set B_i has positive measure. Therefore, there must be a positive measure of cases which belong both to B_h and to B_i. But this contradicts the definition of c.

Step 7. The number of B_i's to which every case belongs is $2m - n$ throughout C (except for measure zero events).

Proof. By Step 6, no more than two quacks choose a given action in every case (except for measure zero events). On the other hand, at least one quack chooses a given action in every case (except for measure zero events). Otherwise, because $n > m$, in every case there is at least one action which is assigned two quacks, and one of these quacks can deviate profitably to the action that is not assigned any quack. But when the number of quacks who choose a given action in a given case is either one or two, then the number of B_i's to which the case belongs is fixed. A simple calculation shows that this number must be $2m - n$.

Step 8. All quacks have the same $G(p)$ for every $p > p^L$.

Proof. Consider p in the support of one quack i. If $\mu_i = 1$, then quack i necessarily plays $p = 1$ with probability one. But according to Step 1, this would mean that all quacks earn a payoff of at least $\frac{1}{m+1}$, which is impossible because $n > m$. Therefore, $\mu_i < 1$ - i.e., there is a positive measure of cases c such that $t_i(c) = t_j(c)$ for some $j \neq i$. Because $n < 2m$, there is a quack $k \neq i, j$ such that $c \in B_k$. If $G_k(p) < G_j(p)$ for some price p in the support of G_i, then quack i can deviate profitably by switching from $t_j(c) = t_i(c)$ to $t_k(c)$. Therefore, $G_k(p) \geq G_j(p)$. If the inequality is strict, then by the same argument as in Step 6, the set $B_j \backslash B_k$ has measure zero. But since B_j has a positive measure - by Step 5 - this means that there are cases c' which belong to more B_i's than c, in contradiction to Step 7. Therefore, $G_k(p) = G_j(p)$.

It follows that if $G_k(p) \neq G_j(p)$ for some pair of quacks k, j, then p does not belong to the support of G_i, for any quack i for whom $t_i(c) = t_j(c)$ in a positive measure of cases c. But if there is an interval of such p's, this means that quack j can raise his price.

Step 9. For every quack i, $\mu_i = 2 - \frac{n}{m}$.

Proof. We have established that all quacks have $\mu_i > 0$ and have the same $G_i(p)$ for every $p > p^L$. This means that they also have the same p^H. But this means that they all must satisfy $p^H = 1$ and that they earn a payoff of $\frac{1}{m+1} \cdot \mu_i$. By Step 1, all quacks earn the same equilibrium payoff. Therefore, all quacks have the same μ_i, which is equal to $2 - \frac{n}{m}$, by Step 7.

Step 10. The quacks' equilibrium pricing strategy is given by expression (4).

Proof. By the preceding steps, the quacks' he payoff from every $p > p^L$ is given by the following equation:

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

which immediately implies the result.

Proof of Proposition 4

Suppose that all quacks charge $p = 0$ and that all the B_i's have measure zero. Then it is easy to verify that we are in equilibrium. Now suppose that quacks charge positive prices with positive probability. Then, by reproducing the steps in the proof of Proposition 3, we reach the conclusion that for every case (except for measure zero events), there are no less than one quack and no more than two quacks who play a given action. But since $n \geq 2m$, this means that $\mu_i = 0$ for all quacks, a contradiction.

Proof of Proposition 5

Let us borrow the definitions of p_i^L, p_i^H, p^L, p^H from the proof of Proposition 1. Steps 1 and 2 can also be borrowed. In addition, no more than one healer places an atom on $p = 1$. Consider the case of $n > 2$**.** By the same symmetry arguments as in the proof of Proposition 1, all quacks ($i \neq e$) play the same strategy G. In particular, they all have the same p_i^L, and G does not 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.

Step 1. Healer e's equilibrium payoff is equal to $\alpha_e(1-\alpha)^n$.

Proof. Suppose that $p_e^H < 1$. Then, $p_i^H = 1$ for all $i \neq e$. Suppose that quack $i \neq e$ deviates by playing p_e^H with probability one. The quack's payoff prior to the deviation is $\alpha(1-\alpha)^{n-1}(1-\alpha_e)$. This follows from the fact that $p_i^H = 1$ maximizes the quack's payoff. Healer e's payoff after i's deviation is at least $\alpha_e(1-\alpha)^n$, because this is his max-min payoff. But this means that healer i's payoff after the deviation is at least $\alpha(1-\alpha)^n$, a profitable deviation. It follows that $p_e^H = 1$. Then, $p = 1$ maximizes healer e's payoff, which is therefore $\alpha_e(1-\alpha)^n$.

Step 2. For every price $p \in [p_e^L, 1]$, $G(p)$ is given by expression (1).

Proof. First, let us see that $p_e^L \geq p_i^L$ for every $i \neq e$. By the fact that all quacks play the same strategy, they all have the same p_i^L. If $p_e^L < p_i^L$, then clearly healer p_e^i fails to maximize healer e's payoff (because some other price between p_e^L and p_i^L would be more profitable). Therefore, $p_e^L \geq p_i^L$. Denote $p_i^L = p^L$.

Let p belong to the support of G_e. Because p maximizes healer e's payoff, 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, $G(p)$ is given by expression (1).

Let us now show that the support of G_e is indeed $[p_e^L, 1]$ - i.e., that G_e is strictly increasing in this interval. Assume the contrary - i.e., that G_e is flat in some interval $(p, p') \subset (p_e^L, 1)$. By the symmetry in the quacks' behavior, G is strictly increasing in this interval. A quack's payoff from the prices p and p' is given by:

$$\begin{aligned} \pi(p) &= p \cdot \alpha \cdot (1-\alpha) \cdot [1-\alpha G(p)]^{n-2} \cdot [1-\alpha_e G_e(p)] \\ \pi(p') &= p' \cdot \alpha \cdot (1-\alpha) \cdot [1-\alpha G(p')]^{n-2} \cdot [1-\alpha_e G_e(p')] \end{aligned}$$

By assumption, $G_e(p) = G_e(p')$. Therefore:

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

But according to the expert's equilibrium condition (expression (7)):

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

and since $G(p') > G(p)$, we obtain a contradiction.

Step 3. $p_e^L = p^L$.

Proof. Assume the contrary - i.e., $p_e^L > p^L$. Because p_e^L maximizes healer e's payoff:

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

The limit of the quacks' payoff as $p \to 1$ 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 both p_e^L maximizes the quack's payoff:

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

Because p^L maximizes the quack's payoff:

$$\alpha \cdot (1 - \alpha)^{n-1} \cdot (1 - \alpha_e + \alpha_e A) = p^L \cdot \alpha \cdot (1 - \alpha)$$

such that

$$p^L = (1 - \alpha)^{n-2} \cdot (1 - \alpha_e + \alpha_e A) \tag{10}$$

Suppose that healer e deviates by playing p^L with probability one. Then, his payoff would be $p^L \cdot \alpha_e \cdot (1 - \alpha)$. In order for this to be an unprofitable deviation, it must be the case that:

$$(1 - \alpha)^{n-2} \cdot (1 - \alpha_e + \alpha_e A) \cdot \alpha_e \cdot (1 - \alpha) \leq \alpha_e \cdot (1 - \alpha)^n$$

such that $1 - \alpha_e + \alpha_e A \leq 1 - \alpha$. Applying this inequality to equation (9), we obtain:

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

By assumption, $G(p_e^L) > 0$. Therefore:

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

in contradiction to equation (8). We conclude that $p_e^L = p^L$.

Step 4. The quacks' payoff is $\alpha \cdot (1-\alpha)^n$.

Proof. Consider equation (8). by Step 3, $G(p_e^L) = 0$. Therefore, $p^L = (1-\alpha)^{n-1}$. By equation (10), $1 - \alpha_e + \alpha_e A = 1 - \alpha$, such that the quacks' payoff is $\alpha \cdot (1-\alpha)^n$.

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 $n > 2$.

Proposition 6

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, Y) is dominated by any (p', N).

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:

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

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 $1 - p' > \alpha - p$ as long as $p < \alpha$. Therefore:

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

This 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).