

The Market for Quacks

RAN SPIEGLER
Tel Aviv University

First version received March 2004; final version accepted January 2006 (Eds.)

A group of n “quacks” plays a price-competition game, facing a continuum of “patients” who recover with probability α , whether 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 . In an extended model that endogenizes the quacks’ choice of “treatments”, the quacks minimize the force of price competition by offering maximally differentiated treatments. The patients’ welfare loss is robust to market interventions, which would crowd out low-quality firms in standard models. Thus, as long as the patients’ quality of reasoning is not lifted above the anecdotal level, ordinary competition policies may be ineffective.

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”, as they possess no skills relative to the default.

If patients understood this model, they would realize that the entire industry provides a worthless service, and the “market for quacks” would be inactive. Indeed, standard market models presuppose that all market agents, firms and consumers alike, have “common knowledge of the model”. This assumption is often more plausible for firms than for consumers. Firms interact with the market on a regular basis, while consumers often make a once-and-for-all decision following a brief period of market exposure. Moreover, firms’ decision process is focused on a specific market, whereas consumers divide their attention among many markets. For these reasons, we should not expect firms and consumers to display the same quality of “market reasoning”.

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 react to these stories as if they are fully informative of the actual quality of treatments. As a result, patients are exposed to exploitation by healers, because they attribute their occasional success to skill rather than luck. The question is whether market competition could mitigate this exploitative effect. I examine the extreme case of a “market for quacks”, that is, 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 $n + 1$ treatments (the quacks and the default). The patient’s sample assigns an 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 treatment i , either from his own past experience or from a fellow patient. The patient chooses the treatment i that maximizes $x_i - p_i$ in his sample. The outcome of his decision is a *new, independent draw* from treatment i . The quacks take into account the patients' choice procedure when determining their pricing strategy.

The patients' behavioural 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 of their quality. I have in mind "soft expertise" industries such as psychotherapy, management consulting, forecasting, and alternative medicine. The effects of skill and luck are hard to disentangle in these industries. Moreover, consumers often enter them when they face an unexpected problem, hence their consumption is not preceded by a long learning phase. In such circumstances, consumers are more likely to rely on anecdotes such as "a friend of mine has taken this pill and he feels better now", or "we should trust this political analyst because he foresaw the collapse of the USSR".

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 would yield different results. To the extent that the procedure reflects ignorance on the patients' part, their ignorance is more characteristic of early stages of a learning process in which the model itself (rather than the value of its fundamentals) is still poorly known. Similarly, the sampling component of the $S(1)$ procedure should not be confused with ordinary models of consumer search. The patient's true expected pay-off from choosing i is $\alpha - p_i$, whereas in a search model it would be the sample realization $x_i - p_i$.

The price-competition game played among the quacks has a unique Nash equilibrium, which 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, a more incurable disease generates a greater amount of charlatany. 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 on average worse off than those who end up choosing the default. The welfare loss does not behave monotonically with respect to n . The reason is that the patients' choice procedure induces an aggregate demand function, which is increasing in n . As long as n is not too large, this force outweighs the competitive force generated by a larger number of competitors.

Thus, increasing the number of competitors does not necessarily reduce the quacks' adverse effect on consumer welfare. In Section 3, I examine alternative market interventions, which may appear like effective competition policies at first glance: (i) raising the success rate of one healer, turning him from a "quack" into a "true expert"; (ii) allowing healers of diverse quality to disclose their success rates. In the first case, the quacks' equilibrium behaviour 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 above the anecdotal level, ordinary competition policies may be ineffective.*

The assumption that treatments are exogenous and statistically independent is quite restrictive. Even if quacks cannot alter their success rate, they may be able to control the correlation with other quacks through their choice of treatments. In Section 4 I analyse an extended model that incorporates this consideration. Specifically, I study a stylized model of a *forecasting* industry in which forecasters without any special forecasting ability choose a "forecasting fee" and a rule for predicting the outcomes of "horse races". Consumers use anecdotal evidence to evaluate the forecasters' quality: they recall a past race at random, and pick the cheapest forecaster among those who predicted the winner in that race.

Such anecdotal reasoning implies that a forecaster can try to avoid competition by differentiating his predictions from his competitors'. As it turns out, in Nash equilibrium, forecasters attain the maximal degree of differentiation: their predictions are as diffuse as possible. This result may explain the proliferation of therapeutic methods that we see in alternative medicine and psychotherapy. At the end of Section 4, I argue that it may also be relevant for certain aspects of the mutual funds industry.

1.1. *Related literature*

Osborne and Rubinstein (1998) analysed games in which all players choose according to $S(1)$. Their main concern was to construct a solution concept for such situations. In the present paper the $S(1)$ is employed by non-strategic agents only. Therefore, it does not call for a non-standard equilibrium concept. Osborne and Rubinstein (2003) study a variant on “ $S(1)$ -equilibrium” in a strategic voting model.

The $S(1)$ procedure is related to other departures from standard probabilistic reasoning. Tversky and Kahneman (1971) demonstrated experimentally that people over-infer from small samples. They explained this phenomenon (dubbed “the law of small numbers”) as a consequence of the “*representativeness*” heuristic: people expect a small sample to mirror the probability distribution from which it is drawn. Rabin (2002) proposed a formal model of inference by “believers in the law of small numbers”. The $S(1)$ procedure reflects an extreme version of the “law”: patients maximize utility against the empirical distribution of recoveries given by their sample, as if this were the true distribution. 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 origins, 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, Rabin’s and Gilboa–Schmeidler’s models have not been incorporated into IO models.

More broadly, this paper belongs to a literature that studies market interaction 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) analyses monopolistic behaviour when consumers differ in their ability to understand complex pricing schedules. Piccione and Rubinstein (2003) study inter-temporal pricing when consumers have diverse ability to perceive inter-temporal patterns. Fishman and Hagerty (2003) study voluntary disclosure by firms, when some consumers are unable to understand the content of the disclosure. Chen, Iyer and Pazgal (2002) analyse a model of price competition when consumers have memory imperfections. For other attempts to introduce non-Bayesian reasoning into game-theoretic modelling, see Eyster and Rabin (2005), Jehiel (2005), and Spiegel (2005).

2. A BASIC MODEL

A market consists of a continuum of identical consumers (“*patients*”) and n identical firms (“*healers*”). When a patient acquires the treatment of a healer $i \in \{1, \dots, n\}$, he “recovers” with probability $\alpha \in (0, 1)$. 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 can ignore the patients’ risk attitudes). 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. Because healers' success rate is the same as the default rate, I refer to them as *quacks*.

Patients choose according to the following procedure, called $S(1)$. Each patient independently samples every alternative (including the default) once. For every $i = 0, 1, \dots, n$, let x_i denote the outcome of the patient's sampling 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. I refer to a sample point as an *anecdote*. Given a sample, the patient chooses an alternative $i \in \arg \max_{i=0,1,\dots,n} x_i - p_i$. In case of ties, he chooses the alternative with the highest p_i . If a tie remains, apply the usual symmetric tie-breaking rule.¹ 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 $\alpha - p_i$.

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 $p_1 \alpha (1 - \alpha)^n$, 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 $p_1 \alpha (1 - \alpha)$, 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. Patients know the exact prices; the only source of variation in their sample is the imperfect recovery rate α , which is exogenously given. A mixed strategy inflicts uncertainty on the quack's opponents, not on the 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. The following re-interpretation of the model sidesteps this difficulty: assume that there are infinitely many quacks, yet patients become aware of only n quacks. Under this interpretation, an increase in n cannot be interpreted as market entry, but as an increase in the patients' awareness of available treatments.

Another artificial feature is the assumption that the number of observations per quack is independent of the size of its clientele. Alternatively, we could assume (as in the word-of-mouth learning due to Ellison and Fudenberg, 1995) that patients sample n fellow patients, rather than n firms, such that they get to hear more anecdotes about quacks with a larger clientele. This variant is more difficult to analyse, because the number of anecdotes per alternative is a fixed point: the patient's sample induces a probabilistic choice, which in turn induces the number of anecdotes per alternative. Such a fixed point need not exist.²

A different sort of criticism is that the model is *formally equivalent* to a conventional model of price competition over consumers with private values. In such a model, let v_i denote the patient's valuation of alternative i , for every $i = 0, 1, \dots, n$. The v_i 's are independently drawn, taking the value 1 (0) with probability α ($1 - \alpha$), and they 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, 1986). Thus, $S(1)$ -patients in the market for quacks behave as if they were rational consumers in a market for a "differentiated product".³

1. I employ this tie-breaking rule merely to simplify the writing of proofs.

2. This variant seems to introduce an anti-competitive force. As I illustrate in Section 3.2, quacks prefer that their patients have an inaccurate perception of their quality. If a greater clientele implies more anecdotes, then quacks may have an incentive to raise prices in order to reduce their clientele.

3. Perloff and Salop suggested that the source of differentiation in their model could be mistakes in consumers' perception of brands. Gabaix, Laibson and Li (2005) adopt this interpretation and study asymptotic properties of the Perloff-Salop model.

What are the merits of the market-for-quacks interpretation, in light of this formal equivalence? As it turns out, the observation that $S(1)$ -patients behave as if they have private values, while simple, has a number of useful implications. First, in a market for medical treatments, it is hard to imagine that patients have intrinsic private values for different treatments. The $S(1)$ procedure provides a concrete process that generates the spurious “private values”, thus grounding it in market fundamentals (the recovery rate α). Second, the market-for-quacks interpretation is the source of the insights we shall obtain into the question of whether the market rewards experts for sheer luck. Third, it will lead to extended models, which would be hard to make sense of under the private-values interpretation. Finally, the market-for-quacks interpretation has radically different welfare implications.

Let us 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.:*

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

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

To see the origin of expression (1), restrict attention to symmetric equilibrium and ignore the question of whether asymmetric equilibria exist. The equilibrium strategy is an atomless c.d.f. G . For every price p in the support of G , the quacks’ pay-off is given by the expression

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

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 $x_j - p_j > x_i - p_i$ for any rival quack j . It follows that

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

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).⁴

Corollary 1. *The quacks’ expected equilibrium price is strictly decreasing with α . In particular, $E(p) \rightarrow 0$ as $\alpha \rightarrow 1$ and $E(p) \rightarrow 1$ as $\alpha \rightarrow 0$.*

Quacks behave as *charlatans* in equilibrium: they charge a positive price for a worthless treatment. The false pretence implicit in their over-pricing gets worse as α decreases, that is, when the patients’ condition becomes more incurable. The intuition is rudimentary. As α decreases, a patient’s sample is less likely to contain multiple successes. This weakens competitive pressures and causes prices to go up. Another way of stating this intuition is that as α decreases, the degree of “product differentiation” erroneously perceived by the patients increases.

Because the monopoly price $p = 1$ belongs to the support of G , the quacks’ equilibrium pay-off 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, we may define this expression as the *welfare*

4. This derivation brings to mind similar characterizations in the literature on equilibrium price dispersion (e.g. Butters, 1977; Varian, 1980; Burdett and Judd, 1983; Rob, 1985).

loss that quacks inflict on patients in equilibrium. This expression is not monotonic in α : it attains a 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 $\alpha \gtrsim 0.4$, $n^* \geq 2$. That is, *more competition may increase the patients' welfare loss*. As $\alpha \rightarrow 0$, n^* tends to infinity, such that the perverse effect of greater competition holds for a wider range. However, fixing α , the welfare loss vanishes as $n \rightarrow \infty$.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 set of available treatments is larger, 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.

The "exploitative" and "competitive" effects can be separated in a simple manner. The max-min pay-off in the game is equal to $\alpha(1-\alpha)^n$, which is precisely what quacks earn in equilibrium. Thus, the "exploitative effect" determines the max-min pay-off, and the "competitive effect" does not allow quacks to earn more than their max-min pay-offs.

3. TWO MARKET INTERVENTIONS

We saw in the previous section 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 market interventions, which would normally be considered as effective competition policies. In contrast, *given the patients' behavioural model, these interventions turn out to be totally ineffective*.

3.1. Replacing a quack with an expert

In this subsection, 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 2. *There is a unique Nash equilibrium in the game. Every healer $i \neq e$ plays the mixed strategy given by equation (1), has the same clientele size, and earns the same profits as in the Nash equilibrium of the basic model.*

Turning a quack into an expert does not affect his competitors' equilibrium behaviour 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

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 $\alpha \rightarrow 0$, the welfare loss at n^* converges to $\frac{1}{e}$.

atom on $p = 1$. The expert's pay-off from the monopolistic price $p = 1$ is thus $\alpha_e(1 - \alpha)^n$. But his pay-off from p_L is $p_L\alpha_e(1 - \alpha)$, 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 pay-off 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 up 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 on account of their anecdotal inferences, although to a lesser degree than the quacks.

3.2. Disclosure of success rates

In the basic model of Section 2, healers have no control over the patients' knowledge. In this subsection, 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 $i = 0, 1, \dots, n$.

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 1. 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.

Until this subsection, the $S(1)$ procedure has meant one thing: patients reason anecdotally about the quality of alternatives. Given the model of this subsection, the procedure acquires another meaning: 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 patients' procedure departs from standard rationality may vary with the model in which it is embedded.

Standard adverse-selection models 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).⁶ Given our model of the patients' behaviour, the result is the complete opposite:

Proposition 3. *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 the 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

6. See, for example, Milgrom and Roberts (1986, section 2).

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 $r_i = Y$).

The lesson from Proposition 3 is that enabling healers to reveal their type is ineffectual when patients choose according to the $S(1)$ procedure.⁷ It is interesting to compare this result 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. In contrast, patients in the present subsection are unsophisticated both strategically and probabilistically.

4. FORECASTERS AS QUACKS

In the basic model, quacks' treatments are exogenous and statistically independent. This assumption is too restrictive for a number of potential applications. For instance, industries such as alternative medicine, psychotherapy, or self-help are characterized by a proliferation of apparently distinct "therapeutic methods". In order to illuminate this phenomenon, one would have to study a model in which a quack's treatment, and therefore, the correlation between his own performance and his competitors', is endogenous. In this section I analyse such an extended model.

To better illustrate the extended model's scope of applications, I abandon the alternative-medicine imagery in favour of a stylized model of a *forecasting* industry, in which firms are assumed to provide worthless forecasting services. This assumption is particularly natural in the case of financial forecasting. If one believes in the "efficient market hypothesis", that is, that market prices are fully revealing, then one must accept that financial analysts have no advantage over a rational layman when trying to predict future prices.

Formally, let $H = \{1, \dots, m\}$ be a set of *horses*. Let R be a finite set of *races*. A *forecasting rule* is a function $t : R \rightarrow H$ that predicts a winning horse in each race. Firms move simultaneously. A pure strategy for firm $i \in \{1, \dots, n\}$ is a pair (p_i, t_i) , where $p_i \in [0, 1]$ is the price that the firm charges (its "forecasting fee") and t_i is the forecasting rule it adopts. A *state* is a pair (r, h) consisting of a race and the identity of the horse who wins in the race. The state is drawn according to the uniform distribution over $R \times H$. Firm i 's prediction is accurate in state (r, h) if and only if $t_i(r) = h$.

The value of an accurate prediction for consumers is 1. If a consumer acquires the services of firm i , he adopts its forecasting rule. If a consumer does not acquire the services of any firm, his forecast is totally unpredictable: in each race, he predicts each horse with probability $\frac{1}{m}$. Note that the probability that a consumer makes a correct forecast is $\frac{1}{m}$, regardless of his decision. Thus, firms are "quacks": they have no advantage over laymen.

The assumption that firms can make race-specific forecasts plays an important role in the model. This is what allows them to control the correlation between their own prediction and their competitors'. This correlation device is natural, in the sense that real-life forecasters are expected to condition their forecasts on the specific conditions of the race. Although we assumed that these conditions are immaterial for the outcome of the race, a consumer who enters the market without realizing that this is the case would not be surprised to see firms making race-specific forecasts.

7. It can also be shown that type disclosure can never be part of Nash equilibrium. I omit the proof of this result for the sake of brevity.

Each consumer chooses according to the following adaptation of the $S(1)$ procedure. He randomly samples a state (r, h) ; he then chooses the alternative $i \in \{0, 1, \dots, n\}$ that maximizes $x_i(r, h) - p_i$ in his sample, where $x_i(r, h) = 1$ if $t_i(r) = h$ and $x_i(r, h) = 0$ if $t_i(r) \neq h$. In this description, alternative $i = 0$ represents the default, “lay” forecast, where $p_0 = 0$ and $t_0(r) = h$ with probability $\frac{1}{m}$ for every h . The tie-breaking rule is the same as in Section 2.

The following story fits the procedure. In order to evaluate forecasters, the consumer recalls some past race at random. He checks which of the firms correctly predicted the winner in that race and chooses the cheapest among them. If the lay prediction was accurate, or if none of the firms made an accurate prediction, the consumer does not consult a professional forecaster. This choice procedure seems to capture the way we often judge financial and political analysts. We recall a past episode—the success of eBay, the collapse of the U.S.S.R., etc.—and we give credit to the analysts who anticipated it while discrediting those who did not.⁸

Firms take into account the consumers’ choice procedure when calculating their profits. Consider a strategy profile $(p_i, t_i)_{i=1, \dots, n}$, where $p_1 > \dots > p_n$. Let A_i be the proportion of races r for which $t_j(r) \neq t_i(r)$ for any $j = i + 1, \dots, n$. The pay-off of firm i is $p_i \left(\frac{1}{m}\right) \left(\frac{m-1}{m} A_i\right)$. The second and third terms represent the probabilities that i ’s prediction is accurate and the lay prediction is inaccurate, respectively.

If all firms choose the same forecasting rule, then firm 1 earns zero profits. The firm can profitably deviate by changing its forecasting rule so as to make exclusive predictions in some races. This deviation guarantees a positive clientele, consisting of consumers who sampled states in which firm 1 alone predicted the winner. Such a deviation may be viewed as an attempt to “differentiate the firm’s product”. Indeed, it is possible to regard our forecasting market as an unconventional model of spatial competition, in which firms choose a price as well as a “location” in the space of functions from R to H . In spatial-competition models, firms fight opponents by reducing prices or by differentiating their product (see, for instance, d’Aspremont, Gabszewicz and Thisse, 1979).

My primary objective is to characterize the diffusiveness of the firms’ forecasts (in other words, the degree of “product differentiation”) that emerges in Nash equilibrium. I restrict attention to equilibria in which the forecasting rule component of each firm’s strategy is pure, while allowing for mixing in the price component. I refer to such equilibria as *Nash equilibria in semi-pure strategies*.

The following notation will be useful. Fix a profile of forecasting rules $(t_i)_{i=1, \dots, n}$. Let B_i be the set of races in which firm i makes an *exclusive prediction*. That is, $B_i = \{r \in R \mid t_i(r) \neq t_j(r) \text{ for every } j \neq i\}$. Denote $\mu_i = |B_i|/|R|$. For every (r, h) , let $e(r, h)$ be the number of firms i for which $t_i(r) = h$, that is, the number of firms which predict that h wins in r .

Definition 1. A profile of forecasting rules $(t_i)_{i=1, \dots, n}$ is *maximally differentiated* if $\max_h e(r, h) - \min_h e(r, h) \leq 1$ for every r , or $e(r, h) \geq 2$ for every (r, h) .

In a maximally differentiated profile of forecasting rules, the firms’ predictions in each race are as unconcentrated as possible (as long as there is a “grain of differentiation”—namely, a horse that is predicted by at most one firm in some race).

Proposition 4. *In any Nash equilibrium in semi-pure strategies, the profile of forecasting rules is maximally differentiated.*

8. It is crucial that (r, h) is held fixed across the consumer’s sample. An alternative adaptation of $S(1)$ would assume that (r, h) is drawn independently for each alternative. This would reduce the formalism to the model of Section 2, with $\alpha = \frac{1}{m}$. I find this version less natural in the context of a forecasting industry.

Thus, a necessary condition for Nash equilibrium in semi-pure strategies is that firms make maximally differentiated predictions. When $n \leq m$, the implication is that they make exclusive predictions in each race: $e(r, h) \leq 1$ for every (r, h) . When $m < n < 2m$, every race r is characterized by $n - m$ horses that are predicted by two firms and $2m - n$ horses that are predicted by one firm. When $n \geq 2m$, firms never make exclusive predictions: $e(r, h) \geq 2$ for every (r, h) .

Note that if we knew that all firms play the same pricing strategy, the result would be trivial. In each race, firms would have a clear incentive to pick a horse that is predicted by as few competitors as possible. The importance of Proposition 4 is that it does not rely on any restriction on the firms' equilibrium pricing behaviour.

To see the main argument behind the result, consider the case of $m > n$ and suppose that the profile of forecasting rules is not maximally differentiated. It is easy to show that $e(r, h) = 1$ for some (r, h) . Let r^* be a race with the maximal number of horses for which $e(r, h) = 1$. Then, there must be at least three firms—say, 1, 2, 3—which predict the same horse h in r^* . Consider any firm i that makes an exclusive prediction in r^* . The only reason for firm 1 not to change its prediction in r^* from h to $t_i(r^*)$ is that firm i is (probabilistically) cheaper than firms 2 or 3.

But now consider some other race r^{**} which belongs to B_3 . A central lemma in the proof of Proposition 4 establishes that such a race must exist. If some other firm predicts the same horse h' as firm i , then that firm must strictly prefer changing its prediction in r^{**} from h' to $t_3(r^{**})$, because firm i is cheaper than firm 3. It follows that $r^{**} \in B_i$. But this means that if firms 1, 2, or 3 make an exclusive prediction in r^{**} , then so does any firm i which makes an exclusive prediction in r^* , in contradiction to the definition of r^* .

Let us turn to the firms' equilibrium pricing behaviour.

Proposition 5. *Let $n \leq m$. There exist Nash equilibria in semi-pure strategies. In equilibrium, each firm plays $p = 1$.*

This is an immediate consequence of the “maximal differentiation” result. Because firms make exclusive predictions in each race, they can afford to price monopolistically. The following result deals with the other extreme. When firms never make exclusive predictions, they have no market power and therefore the market price is competitive.

Proposition 6. *Let $n \geq 2m$. There exist Nash equilibria in semi-pure strategies. In equilibrium, the price paid by consumers is 0.*

The case of $m < n < 2m$ is more complex. For simplicity, I restrict attention to equilibria that are symmetric in the price component. In addition, I assume that $|R|$ is a multiple of $\binom{m}{n-m}$, in order to guarantee equilibrium existence.

Proposition 7. *Let $m < n < 2m$. There exists Nash equilibria in semi-pure strategies, which are symmetric in the price component. Each firm plays the pricing strategy given by the c.d.f.*

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

defined over the support $[\mu, 1]$, where $\mu = 2 - \frac{n}{m}$.

The lesson from the last trio of results is that equilibrium pricing behaviour is qualitatively the same as in the basic model. Specifically, expected equilibrium price rises as the number of horses increases or, equivalently, as the probability of making a correct prediction decreases. The

intuition is that as m increases, firms have “more room” to differentiate themselves in the space of functions from R to H . As a result, competitive forces are weaker and prices go up.

Completely mixed strategies. The restriction to strategies with a pure forecasting rule component carries a loss of generality. Once we allow firms to randomize in both components, there exists an equilibrium in which each firm assigns probability $\frac{1}{m}$ to each horse, independently of the race, and plays the pricing strategy given by Proposition 1 with $\alpha = \frac{1}{m}$. Note that when $n \geq 2m$, this equilibrium implies positive industry profits, while equilibrium in semi-pure strategies results in zero industry profits. Conversely, when $n \leq m$, industry profits are higher in semi-pure equilibria. Thus, when the number of firms is small, they are better off playing an equilibrium in which their predictions are as differentiated as possible. But when there are many firms, they are better off playing an equilibrium in which their predictions are statistically independent.

Possible implications for the mutual funds industry. While the analysis in this section has relied on the “forecasting” metaphor, I believe that the results are potentially relevant for a host of markets. Proposition 4 sheds some light on the phenomenon described at the beginning of this section, namely, the proliferation of therapeutic methods in alternative medicine.⁹

The mutual funds industry is another case in point. Since Jensen (1968), financial economists have gathered evidence that fund managers do not systematically outperform passive benchmarks. Moreover, their performance is largely unpredictable from past relative performance. Nevertheless, flows into and out of mutual funds are highly sensitive to their recent relative performance (see Chevalier and Ellison, 1997; Sirri and Tufano, 1998).

In light of these stylized facts, it seems sensible to search for a “market for quacks” account of the mutual funds industry. An index fund may be viewed as a default option, whereas actively managed funds may be viewed as quacks. The sensitivity of flows to past relative performance is an immediate consequence of investors’ anecdotal reasoning. Proposition 4 suggests that fund managers may choose highly specialized investment strategies as a means of avoiding competition. Propositions 5–7 suggest a positive relation between the underlying market risk (captured by the parameter m) and the fees that mutual funds charge. Careful modelling of the industry as a “market for quacks” is left for future research.

5. CONCLUDING REMARKS

This paper presented a model with rational firms and consumers who are boundedly rational, in the sense that they reason anecdotally about firms’ quality. Anecdotal reasoning implies that consumers react to a common-value environment as if they have independent private values. As a result, a market for a worthless service becomes active and displays anomalous features: standard competition policies may be ineffective; a decrease in the quality that characterizes the industry results in higher prices; and firms can avoid competition by making divergent recommendations about the desired action. These results may illuminate phenomena associated with “soft expertise” industries such as alternative medicine, forecasting, and money management.

Imperfect rationality or imperfect information? Although the models presented in this paper are simple, the modelling 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

9. Under an alternative-medicine interpretation of the model, R would be a set of medical conditions, H would be a set of treatments, and t would be a “therapeutic method”.

relaxed the *rationality* of consumers' choice with respect to the complete-information 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. Note 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. 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 modelling approaches widens. In the model of Section 4, the statistical structure of a firm's product is endogenous. If we tried to rewrite the model with imperfectly informed, rational patients, then in equilibrium they would have to know the firms' forecasting rules. It is hard to see how one could reconcile such equilibrium knowledge with the behaviour we attempt to rationalize. In the model of Section 3.2, 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) and called $S(K)$, is to assume that patients sample every alternative K times and maximize their expected pay-off against the empirical distribution generated by their 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 1. The patients form quality assessments of healers as in the basic model.

The essential features of our equilibrium characterization—uniqueness, symmetry, price dispersion, as well as 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, as if a single anecdote has the same informational content as full knowledge of a probability distribution.

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 price is 0, and the patients’ expected 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’ pay-off function through the multiplicative term $1 - \alpha_0$, and it cancels out when we derive the expression for G . Therefore, the healers’ behaviour is independent of α_0 . The welfare analysis is modified. For instance, when $\alpha_0 = 0$, the patients’ expected utility in equilibrium 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.

A dynamic justification for the $S(1)$ procedure. The models presented in this paper are static. However, the interpretation of $S(1)$ as a best reply to a random sample suggests a dynamic learning context. The following is an outline of an explicit dynamic model which justifies our basic model. Quacks commit to their pricing strategy at period $t = 0$ (if a quack plays a mixed strategy, it commits to its realization). At any period $t > 0$, a constant measure of patients enter the market, make a one-shot decision and then exit the market. The outcome of a quack’s treatment is drawn independently each time it is chosen. Patients choose according to the following rule: (i) with probability $1 - \varepsilon$, they imitate the patients who earned the highest pay-off at $t - 1$; (ii) with probability ε , they choose each of the $n + 1$ alternatives with equal probability. This is simply a best-reply dynamics combined with blind experimentation. It can be shown that as ε tends to 0, patients’ long-run average behaviour converges to the average behaviour implied by the $S(1)$ procedure.

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 quack i assigns a positive measure to a price that yields zero profits. As we noted in Section 2, the quacks’ max–min pay-off 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 pay-off 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 $p^L = \min\{p_1^L, \dots, p_n^L\}$ and $p^H = \max\{p_1^H, \dots, p_n^H\}$. \parallel

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 pay-off of the quack who charges p^H is $p^H \alpha (1 - \alpha)^n$, which is below the max-min pay-off, a contradiction. \parallel

Step 3. All quacks earn the same pay-off in equilibrium.

Proof. Assume the contrary, and suppose (without loss of generality, wlog) that quack 1 earns a higher pay-off than quack 2. Suppose that quack 2 deviates by playing p_1^L with probability 1. Quack 1's pay-off before the deviation is:

$$p_1^L \alpha (1 - \alpha) \prod_{j>1} [1 - \alpha G_j(p_1^L)], \quad (5)$$

whereas quack 2's pay-off after the deviation is:

$$p_1^L \alpha (1 - \alpha) \prod_{j>2} [1 - \alpha G_j(p_1^L)], \quad (6)$$

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

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

Proof. By Step 2, $p^H = 1$. We have observed that there exists a quack i whose competitors do not place an atom on $p = 1$. This quack's pay-off is $\alpha(1 - \alpha)^n$. By Step 3, all quacks earn this pay-off. \parallel

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 1. Then, p_j^H does not maximize j 's pay-off, because if j charged a slightly higher price, the probability that he is chosen by a patient would not change. It follows that $p_j^H = 1$. But since quack i places an atom on $p = 1$, quacks i and j earn different pay-offs, in contradiction to Step 3. \parallel

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

Proof. Assume the contrary, and suppose (wlog) that $p_2^L = p^L$ and $p_1^L > p^L$. Suppose that quack 2 deviates by playing p_1^L with probability 1. Expressions (5) and (6) represent quack 1's pay-off before the deviation and quack 2's pay-off 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 pay-off prior to the deviation. Therefore, the deviation is profitable. \parallel

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

Proof. Assume the contrary, and suppose (wlog) 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, wlog) who assigns positive weight to the neighbourhood of p , otherwise, p would not maximize quack 1's pay-off. The two quacks' pay-off from the prices p and p' are

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

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 1. \parallel

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

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

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). By construction, every element in the support of G_i is a best reply to G_{-i} , hence we have a Nash equilibrium. $\parallel \parallel$

Proof of Corollary 1.

Given the formula for $G(p)$ given by Proposition 1, it is easy to calculate the expected equilibrium price:

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

It is straightforward to show that both expressions decrease with α , and that their limits are $\lim_{\alpha \rightarrow 1} E(p) = 0$ and $\lim_{\alpha \rightarrow 0} E(p) = 1$. \parallel

Proof of Proposition 2. Let us borrow the definitions of p_i^L, p_i^H, p^L, p^H from the proof of Proposition 1. 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 pay-off 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 1. The quack's pay-off 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 pay-off. Healer e 's pay-off after i 's deviation is at least $\alpha_e(1-\alpha)^n$, because this is his max-min pay-off. But this means that healer i 's pay-off 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 pay-off, which is therefore $\alpha_e(1-\alpha)^n$. \parallel

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^L fails to maximize healer e 's pay-off (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 pay-off, the following equation holds:

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

Therefore, $G(p)$ is given by expression (1).

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

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

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

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

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

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

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

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

Proof. Assume the contrary—that is, $p_e^L > p^L$. Because p_e^L maximizes healer e 's pay-off

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

The limit of the quacks' pay-off as $p \rightarrow 1$ is $\alpha(1 - \alpha)^{n-1}(1 - \alpha_e + \alpha_e A)$, where A is the size of the atom that G_e places on $p = 1$. Because both p_e^L maximizes the quack's pay-off

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

Because p^L maximizes the quack's pay-off:

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

such that

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

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

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

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

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

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

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

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

Step 4. *The quacks' pay-off is $\alpha(1 - \alpha)^n$.*

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' pay-off is $\alpha(1 - \alpha)^n$.

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 pay-off is just the same as in the case of $n > 2$. $\parallel \parallel$

Proposition 3. *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 pay-off from (p, Y) is 0. 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 0, and the probability that $x_i - p_i = x_0$ is $\alpha(1 - \alpha_0)$. If healer i deviates to $(1 - \varepsilon, N)$, the probability that $x_i - p_i > x_0$ is $\alpha(1 - \alpha_0)$, and the probability that $x_i - p_i = x_0$ is 0. 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 pay-off from the strategy (p, Y) is bounded from above by

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

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

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

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

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

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) . \parallel

Proof of Proposition 4.

Consider a Nash equilibrium in semi-pure strategies, in which the profile of forecasting rules is $(t_i)_{i=1, \dots, n}$. Borrow the definitions of p_i^L, p_i^H, p^L, p^H from the proof of Proposition 1.

Step 1. All firms earn the same pay-off in equilibrium.

Proof. Assume that firm i earns a higher equilibrium pay-off than firm j . Then, $p_i^L > 0$. Let firm j deviate to the pure strategy $(p_j^L - \varepsilon, t_j)$, where ε is arbitrarily small. The probability that j is chosen after the deviation, denoted α , is at least as high as the probability that i was chosen prior to j 's deviation. Therefore, j 's pay-off after the deviation is $\alpha(p_j^L - \varepsilon)$. Firm i 's pay-off prior to j 's deviation was at most αp_i^L . Therefore, j 's deviation is profitable. \parallel

Step 2. $p_i^L = p^L$ for all firms i .

Proof. Suppose that there exist firms i and j such that $p_i^L > p_j^L$. Assume first that $t_i(r) = t_j(r)$ for some race r . Suppose that firm j deviates to the pure strategy $(p_j^L - \varepsilon, t_j)$, where ε is arbitrarily small. The probability that j is chosen is now strictly higher than the probability that i was chosen prior to j 's deviation. Therefore, j 's pay-off after the deviation is higher than i 's pay-off prior the deviation. By Step 1, this is a profitable deviation for j , a contradiction. It follows that for every pair of firms i and j for which $p_i^L > p_j^L$, $t_i(r) \neq t_j(r)$ for every r . Then, firm i 's equilibrium pay-off is $p_i^L (\frac{1}{m}) (\frac{m-1}{m})$ while firm j 's equilibrium pay-off is $p_j^L (\frac{1}{m}) (\frac{m-1}{m})$, in contradiction to Step 1. \parallel

Step 3. Either $B \neq \emptyset$ for all firms, or $B = \emptyset$ for all firms.

Proof. Assume the contrary, that is, that there are firms k and l , such that $B_k \neq \emptyset$ and $B_l = \emptyset$. Note that firm k necessarily earns a positive pay-off. By Step 1, all firms earn a positive pay-off.

Define the following binary relation: $i \succsim j$ if firm i 's pricing strategy assigns positive probability to prices $p \geq p_j^H$. It is easy to verify that \succsim is complete and transitive. Note that if i is \succsim -maximal, then $B_i \neq \emptyset$, otherwise, this firm would earn zero profits, a contradiction. Consider a \succsim -maximal firm j among those with $B = \emptyset$. Then, for every race r , there exists a firm $i \succsim j$ such that $t_i(r) = t_j(r)$, otherwise, j would earn zero profits. Moreover, if $i \sim j$, then both firms place an atom on p_j^H . But in this case, firm j can profitably deviate by shifting this atom slightly below p_j^H . It follows that $i \succ j$. Let i^* be the \succsim -maximal firm among these firms i .

By definition, $B_{i^*} \neq \emptyset$. That is, there is a race r' and a horse h such that i^* is the only firm that predicts h in r' . Note that there is a firm k , such that $i^* \succ k \succ j$ and $t_k(r') = t_j(r') \neq h$. Firm k can deviate by switching to a forecasting rule t'_k that differs from t_k only in that $t'_k(r') = h$. This deviation increases the probability that k is chosen, hence it is profitable. \parallel

Step 4. If $e(r, h) < 2$ for some (r, h) , then $\max_h e(r, h) - \min_h e(r, h) \leq 1$ for every r .

Proof. Assume first that $e(r, h) = 0$ for some (r, h) . If $e(r, h') > 1$ for some other horse h' , then at least one of the firms that predict h' in r can profitably deviate by predicting h in r . Thus, $e(r, h) \leq 1$ for every h , which is possible only when $n \leq m$. It follows that for every (r, h) , either $e(r, h) = 0$ or $e(r, h) = 1$, hence $(t_i)_{i=1, \dots, n}$ is maximally differentiated.

Now suppose that $e(r, h) > 0$ for every (r, h) , and assume that $(t_i)_{i=1, \dots, n}$ is not maximally differentiated. For every r , let $b(r)$ denote the number of horses for which $e(r, h) = 1$. Define $b^* \equiv \max_{r \in R} b(r)$. By assumption, $b^* > 0$. Let r^* be a race satisfying $b(r^*) = b^*$. Then, there must be a horse h such that $e(r^*, h) > 2$.

Let i, j, k be three distinct firms such that $t_i(r^*) = t_j(r^*) = t_k(r^*) = h$. Let l be a firm that makes an exclusive prediction in r^* . Because $b^* > 0$, such a firm must exist. By Step 2, there is a price $p > p^L$, which belongs to the supports of all firms' pricing strategies, such that in particular, $G_i(p), G_j(p), G_k(p) > 0$. If $G_l(p) \leq G_i(p)$, then j can profitably deviate to a pure strategy (p'_j, t'_j) satisfying: (i) $p'_j = p$; (ii) $t'_j(r^*) = t_i(r^*)$ and $t'_j(r) = t_j(r)$ for every $r \neq r^*$. It follows that $G_l(p) > G_i(p)$.

By Step 3, $b^* > 0$ implies that $B \neq \emptyset$ for all firms. Consider a race $r^{**} \in B_l$. Suppose that $r^{**} \notin B_l$, that is, there exists a firm g such that $t_g(r^{**}) = t_l(r^{**})$. We have shown that $G_l(p) > G_i(p)$. But this means that g can profitably deviate to a pure strategy (p'_g, t'_g) satisfying (i) $p'_g = p$, (ii) $t'_g(r^{**}) = t_i(r^{**})$ and $t'_g(r) = t_g(r)$ for every $r \neq r^{**}$. Therefore, it must be that $r^{**} \in B_l$. We have established that $B_i \subset B_l$. But this holds for any firm l among the b^* firms who make an exclusive prediction in r^* . Therefore, $b(r^{**}) > b^*$, contradicting the definition of b^* . $\parallel \parallel$

Proof of Proposition 5.

Existence of Nash equilibrium with semi-pure strategies is easy to verify: for every firm i , let $p_i = 1$ and $t_i(r) = i$. Each firm earns a pay-off of $(\frac{1}{m})(\frac{m-1}{m})$, which is the maximal pay-off that is possible in the model. Suppose that there exist equilibria with $p_i < 1$ or $\mu_i < 1$ for some firm i . Then, this firm attains a pay-off below $(\frac{1}{m})(\frac{m-1}{m})$. The firm can deviate to a pure strategy (p'_i, t'_i) satisfying $p'_i = 1$ and $t'_i(r) \neq t'_i(r)$ for every race r and every $j \neq i$, and attain this pay-off. \parallel

Proof of Proposition 6.

Existence of Nash equilibrium with semi-pure strategies is easy to verify. Suppose that all firms charge $p = 0$, and construct a maximally differentiated profile of forecasting rules satisfying $e(r, h) \geq 2$ for every (r, h) . Then, no firm has any profitable deviation.

Let us now show that the price paid by consumers must be 0 in any equilibrium in semi-pure strategies. Because $n \geq 2m$, Proposition 4 implies that $e(r, h) \geq 2$ for every (r, h) . Therefore, $B = \emptyset$ for each firm. By Step 1 in the proof of Proposition 4, all firms earn the same equilibrium pay-off. Consider the set of firms i with $p_i^H = p^H$. There must be at least one such firm, for which the probability of being chosen by consumers is 0. Therefore, this firm earns zero profits. But this means that all firms earn zero profits, hence the price paid by consumers is 0. \parallel

Proof of Proposition 7.

Because $n > m$, for every race r there is a horse h such that $e(r, h) > 1$. Therefore, the firms' pricing strategy must be mixed and atomless—otherwise, some firms would be able to profitably deviate by shifting the atom to a slightly lower price. Let G denote the firms' mixed pricing strategy. By a standard argument, the support of G is an interval $[p^L, 1]$. Firm i 's pay-off from $p_i = p^H = 1$ is thus $\mu_i(\frac{1}{m})(\frac{m-1}{m})$. Since all firms earn the same equilibrium pay-off, μ is identical for all firms. By Proposition 4, in each race there are $n - m$ horses h with $e(r, h) = 2$ and $2m - n$ horses with $e(r, h) = 1$. Therefore, $\mu = \frac{2m-n}{m}$. Furthermore, for every price p in the support of G :

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

The formula for $G(p)$ follows from this equation.

Let us establish equilibrium existence. Suppose that all firms play the pricing strategy given by expression (4). It remains to construct a profile of forecasting rules which satisfies the necessary condition of Proposition 4, as well as $\mu = \frac{2m-n}{m}$ for all firms. Consider the class of allocations of firms to horses, such that each horse $h \in \{1, \dots, n-m\}$ is allocated two firms, and each horse $h \in \{n-m+1, \dots, m\}$ is allocated a single firm. There are $\binom{m}{n-m}$ such allocations. Partition R into $\binom{m}{n-m}$ subsets of equal cardinality, and attach a distinct allocation to each cell in the partition, such that all races in a given cell have the same allocation of firms to horses. This constitutes a profile of forecasting rules with the desired properties. \parallel

Acknowledgements. I thank Eddie Dekel, Glenn Ellison, 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, Hebrew University, LSE, Tel Aviv, and UCL, for helpful comments.

REFERENCES

- BURDETT, K. and JUDD, K. (1983), "Equilibrium Price Dispersion", *Econometrica*, **51**, 955–970.
 BUTTERS, G. (1977), "Equilibrium Distributions of Sales and Advertising Prices", *The Review of Economic Studies*, **44**, 465–491.
 CHEN, Y., IYER, G. and PAZGAL, A. (2002), "Limited Memory and Market Competition" (Mimeo, New York University).
 CHEVALIER, J. and ELLISON, G. (1997), "Risk Taking by Mutual Funds as a Response to Incentives", *Journal of Political Economy*, **105**, 1167–1200.