Becoming a Bad Doctor*

Nora Szech

23 January 2009

Abstract

We consider a market model with n rational "doctors" and a continuum of boundedly rational "patients". Following Spiegler (2006), we assume that patients are not familiar with the market and rely on anecdotes.
We analyze the price setting game doctors with given, different healing qualities play. Doctors know their own quality as well as the qualities of their competitors. We find a unique equilibrium in mixed strategies. All doctors, no matter how bad, make positive profits that are typically considerably higher than their maxmin payoffs.
In order to analyze welfare, we introduce a pre-stage where doctors are allowed to choose their qualities themselves simultaneously. Even though a better quality is at no costs, doctors mainly offer mediocre qualities in SPNE. If the highest possible quality is high enough, welfare strictly decreases in n.

JEL Classification: L13, D44, I11

Keywords: bounded rationality, S(1) procedure, quacks, product differentiation, health care, price dispersion

*I would like to thank Mark Armstrong, Olivier Bos, Paul Heidhues, Fabian Herweg, Sebastian Kranz, Benny Moldovanu, Patrick Schmitz, Anja Schöttner, Avner Shaked and seminar participants in Bonn. Nora Szech, Bonn Graduate School of Economics, University of Bonn, Kaiserstraße 1, 53113 Bonn, Germany; email: nszech@uni-bonn.de.

1 Introduction

A consumer who turns to a market he is not familiar with may rely on the advice of other consumers. An example would be a patient with a health problem thinking about a therapy because some friend got cured with it - or the owner of a house not consulting the roofer he heard a bad story on. Whereas judging the quality of a "doctor" is difficult for an unfamiliar "patient" (which makes relying on others' advice a maybe not so bad idea), prices are typically easy to grasp.

In our model, each patient asks one client of each doctor if he got cured or not. The patient assumes that the treatment of a doctor leads to the same result as the treatment of the client he asked for advice. Hence patients only think about attending recommended doctors, i.e. doctors they heard a success story on. Among the recommended doctors, patients pick the cheapest one. Doctors know the market they are in, the qualities, i.e. success probabilities, of their competitors and the way patients look for doctors. They act rationally.

We first consider the pricing game doctors play with possibly different, exogenous qualities. We find that no matter how bad a doctor is, he survives in the market - and his equilibrium payoff is typically much larger than his maxmin payoff. The reason is that good doctors do not feel threatened enough by bad doctors to set low prices. Mixing over rather high prices, they leave room for bad doctors to make considerable payoffs.[1] No doctor, even the best one, has an incentive to reveal his true quality: If he would state his true quality, his position in competition with other recommended doctors would be weakened.

In order to study welfare, i.e. the proportion of patients cured thanks to the doctors, we introduce a pre-stage: Doctors set their qualities themselves before the pricing game takes place. Considering pure quality-setting strategies, we find that in SPNE doctors mainly offer low qualities. A low quality makes a doctor attract fewer patients, but it makes price competition much softer. The latter effect on payoff dominates, thus doctors do not set high qualities, even if setting higher qualities does not incur any direct costs. We find that heavily restricting the number of doctors is helpful as it makes the average quality of treatment a patient receives higher. If the maximal quality doctors could offer is high enough, having a monopolistic

[1] This might explain myriads of obscure nutritional supplements never clinically tested to improve medical conditions, or the survival of medical approaches that in studies do not outperform placebos.

doctor maximizes welfare (and welfare strictly decreases in the number of doctors). It would be even more beneficial for welfare not to limit entry but to replace the pricing stage by exogenously prescribing a fixed price. Then doctors could not take advantage of a weaker price competition and thus had no incentive to choose a lower than maximal quality.

Our model is an extension of Ran Spiegler's "Market for Quacks" (2006a). We adopt his modeling of the pricing game, with the difference that we allow doctors to have generally asymmetric qualities. This generalization to the asymmetric case allows us to let doctors choose their qualities themselves, making the model more suitable for the study of social welfare. The sampling rule our patients apply to evaluate doctors is the S(1) rule which was introduced by Osborne and Rubinstein (1998). Spiegler and Rubinstein have utilized the S(1) rule to model consumer behavior in a variety of settings (see Spiegler (2006a, 2006b) and Rubinstein and Spiegler (2008)). We have used the S(1) procedure as well in a companion paper, Szech (2008).[2]

Besides S(1) there are other related proposals for modeling boundedly rational consumer behavior such as Ellison and Fudenberg's (1995) "word-of-mouth learning" and Rabin's (2002) "law of small numbers". More broadly, our paper contributes to the literature on interactions between rational firms and boundedly rational consumers, see for instance the survey by Ellison (2006). Yet to our knowledge, this paper is the first to extend a price competition game with boundedly rational consumers via introducing a preceding quality setting stage in order to study how bounded rationality affects welfare.

Technically, our analysis has some parallels to papers on price dispersion like Varian (1980) or to papers on complete information all-pay auctions such as Baye, Kovenock and de Vries (1996).[3] Like in these models, the equilibrium of our pricing stage is in mixed strategies. We first specify equilibrium payoffs for all possible equilibria and then identify sequentially the unique equilibrium candidate - a similar approach has been used, for instance, by Siegel (2008) in the context of generalized all-pay auctions.

[2]In there, we consider a variant of Spiegler's (2006a) model where the doctors' qualities are privately known random variables. We identify an equilibrium in monotone pricing strategies of that model and show that welfare goes to zero in the number of doctors. This happens because patients always attend the cheapest among the doctors who are recommended - in monotone strategies this is also the worst recommended doctor.

[3]Indeed, our model is a complete-information first-price procurement auction with stochastic participation.

Reinterpreting the model with rational patients (which would affect the welfare implications, but not the doctors' behavior) one sees the close relation of our model to the advertising model of Butters (1977). The same reinterpretation is of course possible with the Spiegler model, which is then a special (not considered) case of the Perloff and Salop (1985) model of product differentiation. Our extension of the game by introducing a quality-setting stage parallels Shaked and Sutton (1982): They introduce a pre-stage of quality setting to the Gabszewicz and Thisse (1979) pricing model. We do the same with a Perloff-Salop-type pricing model. In Shaked and Sutton (1982), only a limited number of firms can make positive profits. This is not true in our model. The reason lies in the different modeling of consumers' preferences: In Shaked and Sutton (1982), the same ranking of products is valid for all consumers. In our paper, for each product there is a group of consumers which prefer it to all other products.

The paper is structured as follows: Section 2 presents the model and describes the $S(1)$ procedure in detail. In Section 3, we analyze the second stage of decision making in the game (the price setting stage) and identify its unique Nash equilibrium. Equilibria of the quality setting stage and their welfare implications are discussed in Section 4. Section 5 discusses a number of extensions and variations of our model. First we study the robustness of our model to the following modifications: costly quality choice, changes in the timing, allowing doctors to disclose their qualities and more sophisticated patients. We then compare our price setting game to its second price version. Finally we consider the reinterpretation of our model as a model of product differentiation or advertising with rational customers. Section 6 concludes. All proofs are in the Appendix.

2 The Model

We consider a market with n "doctors" that are familiar with the market and can act rationally. This is in contrast to the "patients" in the market: They are not familiar with the market and apply a simple sampling rule described below. Patients form a continuum of mass one. We want to find SPNE of the game the doctors play: First, doctors decide on the qualities α_i they want to offer, then they decide on the prices P_i they want to charge. We assume that doctors know each other very well and hence know the qualities of each other when playing the pricing game. Qualities can be anything between zero and some upper bound $0 < \overline{\alpha} < 1$. A doctor's quality is the probability with which he can cure a patient. With the counter probability, the patient remains ill. Before we specify the patients' behavior, we give the exact **timing** of the model:

1. Doctors simultaneously set their qualities $\alpha_i \in [0, \overline{\alpha}]$.
2. Doctors observe each others' qualities.
3. Doctors simultaneously set their prices P_i.
4. Patients decide if they want to attend a doctor and if so, which one.

Patients are initially ill and have a utility of one from getting cured and a utility of zero from staying ill. They decide according to the behavioral rule **S(1)** as introduced by Osborne and Rubinstein (1998) and utilized in Spiegler (2006a):

- Each patient samples each doctor once.
- With probability α_i, a patient gets a positive signal $S_i = 1$ on doctor i ("**a recommendation**").
- With probability $1 - \alpha_i$, a patient gets a negative signal $S_i = 0$ on doctor i ("**no recommendation**").
- A patient attends the doctor with the highest $S_i - P_i$...
- unless $\max_i S_i - P_i < 0$. Then the patient stays out of the market and expects a utility of 0 at a price of 0.

Note that the last two points implicitly contain a tie-breaking rule: If a patient has to choose between consulting a recommended doctor setting a price of one and staying at home, the patient opts for the doctor. It can be shown that in the pricing stage no equilibrium exists if we depart from this assumption. For all other possible types of ties, no special breaking rule is needed for our results - ties can be broken arbitrarily.

Note that patients rely far too much on the signal they get - they overinfer from their sample. The idea behind the S(1) rule is to capture a simple way of anecdotal reasoning: Each patient independently asks some "former" client of each doctor.[4]

[4] Of course, we are not in a dynamic model here. This is only some motivating story.

A client of doctor i got cured with probability α_i. Thus, with probability α_i, he recommends doctor i to the patient. The patient trusts in this report - he either thinks the doctor can cure him as well for sure or not at all.

Choosing a higher quality comes at no direct costs for the doctors. The motivation for this assumption is that we want to study how the patients' boundedly rational behavior induces doctors to set a low quality. Our model is to be understood as a benchmark case which ignores costs that give doctors another, separate reason for providing a low quality. In Section 5.1 we demonstrate that our results are robust with respect to the introduction of convex cost functions.

3 The pricing stage

We search for SPNE using backwards induction, and thus we start with an analysis of the price setting game for given quality levels α_i.[5] We first show that the equilibrium payoffs of the price setting stage are unique. Then we identify the unique equilibrium of the pricing stage.

Proposition 1 *Take α_1,..., α_n as given. Then in all equilibria of the price setting game, the payoff of doctor i is given by*

$$\pi_i = \alpha_i \prod_{j \neq j^*} (1 - \alpha_j) \tag{1}$$

where $j^* \in argmax_j \alpha_j$.

The intuition for this is as follows: Consider doctor j^* who offers the highest quality of all doctors.[6] This doctor can make a positive profit independent of his competitors' strategies, as there will be patients who only get a positive report on him, but not on the other doctors. (The fraction of these patients is $\alpha_{j^*} \prod_{j \neq j^*} (1 - \alpha_j)$). Our doctor will thus set a positive lowest price in any equilibrium. Consider one equilibrium and assume our doctor makes there a payoff of $\alpha_{j^*} C$ and sets a lowest price of p_L. Via charging an only slightly lower price than this, every other doctor i can make at least a profit arbitrarily close to $\alpha_i C$. Reasoning vice versa we see that

[5]The price setting game is essentially a generalization of the game analyzed by Spiegler who assumes that all (or all but one) doctors have the same α.

[6]If there is more than one, we (arbitrarily) determine as best doctor one of the doctors with highest qualities.

the payoffs indeed have to equal $\alpha_i C$ and cannot be higher. To determine C, we find out that there is one doctor i who has a price of 1 in the support of his equilibrium price setting strategy. This doctor only makes a profit if he is the only recommended doctor, which happens with a probability of $\alpha_i \prod_{j\neq i}(1-\alpha_j)$. Finally, we find that this doctor has to be the best doctor, as otherwise C would be so low that the best doctor would prefer to deviate. Thus we can specify $C = \prod_{j\neq j^*}(1-\alpha_j)$.

From Proposition 1, we see that the quality of the best doctor does not appear in the payoff formulas of the other doctors. Hence it does not matter for the competitors' profits whether the best doctor is equally good as the second best doctor or much better. For the second best doctor, this is not true: His quality always affects the competitors' profits.[7]

One important question left open by Proposition 1 is equilibrium existence to which we turn now. We will see that the equilibrium we find is unique and in mixed strategies. To get an intuition for the equilibrium distribution functions we find, assume all doctors have some price $\widetilde{p}$ in their supports. Then, the equilibrium distribution functions G_i must fulfill:

$$\pi_1(\widetilde{p}) \overset{(1)}{=} \alpha_1 \prod_{j\neq j^*}(1-\alpha_j) = \alpha_1 \widetilde{p} \prod_{j\neq 1}(1-\alpha_j G_j(\widetilde{p})).$$

The middle part of the equation is the equilibrium payoff already specified. The rhs is the payoff doctor 1 makes from playing a price of $\widetilde{p}$ given that the other doctors mix according to G_j: $\widetilde{p}$ is the price he earns if he is consulted. α_1 is the probability that he is recommended. The product is the probability that all competitors are not recommended or pricier than i. The same must hold for doctor 2.

$$\pi_2(\widetilde{p}) \overset{(1)}{=} \alpha_2 \prod_{j\neq j^*}(1-\alpha_j) = \alpha_2 \widetilde{p} \prod_{j\neq 2}(1-\alpha_j G_j(\widetilde{p})).$$

We thus have

$$\frac{1}{\widetilde{p}} \prod_{j\neq j^*}(1-\alpha_j) = \prod_{j\neq 1}(1-\alpha_j G_j(\widetilde{p}))$$

and

$$\frac{1}{\widetilde{p}} \prod_{j\neq j^*}(1-\alpha_j) = \widetilde{p} \prod_{j\neq 2}(1-\alpha_j G_j(\widetilde{p}))$$

[7] Thus, if more than one quack in Spiegler's market would be changed into some expert, the payoffs of the remaining quacks would be diminished, compare Proposition 2 of Spiegler (2006a).

from which we see that

$$1 - \alpha_1 G_1(\widetilde{p}) = 1 - \alpha_2 G_2(\widetilde{p})$$

or, more generally,

$$1 - \alpha_k G_k(\widetilde{p}) = 1 - \alpha_l G_l(\widetilde{p})$$

for all k, l.

We can thus write

$$\prod_{j \neq j^*} (1 - \alpha_j) = \widetilde{p} \prod_{j \neq 1} (1 - \alpha_j G_j(\widetilde{p})) = \widetilde{p}(1 - \alpha_i G_i(\widetilde{p}))^{n-1}$$

and therefore

$$\frac{1}{\alpha_i}\left(1 - \sqrt[n-1]{\frac{\prod_{j \neq j^*}(1 - \alpha_j)}{\widetilde{p}}}\right) = G_i(\widetilde{p})$$

for all i.

This is the basic reasoning behind Proposition 2. Since the doctors may mix over different price intervals, the equilibrium looks a bit more complicated. Before the proposition is stated, the following picture demonstrates what price intervals doctors mix on typically look like.

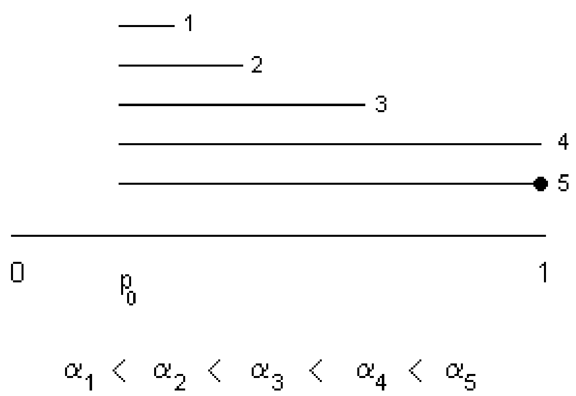

Figure 1: Supports of the doctors' pricing strategies for $n = 5$

We see that the doctors all have the same lowest price in the interval. Furthermore wee see that the better the quality of a doctor, the larger his price interval.

Proposition 2 (***i***) *Consider* $0 < \alpha_1 \leq \ldots \leq \alpha_n \leq \overline{\alpha}$. *Define a sequence of prices* $p_0, \ldots, p_n$ *by*

$$p_i = \frac{(1-\alpha_{i+1}) \cdot \ldots \cdot (1-\alpha_{n-1})}{(1-\alpha_i)^{n-i-1}}$$

for $1 \leq i \leq n-2$,

$$p_0 = \prod_{i=1}^{n-1} (1-\alpha_i)$$

and $p_{n-1} = p_n = 1$. *Then, a mixed strategy Nash equilibrium of the price setting game is given as follows: Each doctor* i *mixes over the interval* $[p_0, p_i]$ *using the distribution function* G_i *defined as*

$$G_i(p) = \frac{1}{\alpha_i}\left(1 - \sqrt[n-j]{\frac{(1-\alpha_j) \cdot \ldots \cdot (1-\alpha_{n-1})}{p}}\right)$$

for $p \in [p_{j-1}, p_j] \subset [p_0, p_i]$ *with* $1 \leq j \leq n-1$. *On* $[0, p_0]$, *define* $G_i = 0$ *and on* $[p_i, 1]$ $G_i = 1$. G_n *places an atom of size* $1 - \frac{\alpha_{n-1}}{\alpha_n}$ *at* 1.

(***ii***) *The equilibrium determined in* (***i***) *is unique.*[8]

Note that the upper boundaries p_i coincide if the corresponding α_i coincide. The distribution functions G_i are continuous except for G_n as doctor n puts an atom on $p = 1$ (if $\alpha_{n-1} \neq \alpha_n$).

We first showed payoff uniqueness (Proposition 1). Knowing the equilibrium payoffs, we sequentially constructed a unique equilibrium candidate (Proposition 2). A very similar approach was recently taken by Siegel (2008) to analyze complete information all-pay auctions with general cost functions. Furthermore, our equilibrium somewhat resembles the mixed-strategy equilibrium of a complete information all-pay auction. Of course, the price setting game we consider is not an all-pay-auction, but a complete-information first price auction with equal valuations and random participation. What these two auction models have in common is that some players can secure a positive maxmin payoff by making a high bid.[9] Such behavior is, how-

[8]Note that Propositions 1 and 2 still hold in the (excluded) case where exactly one doctor has a quality of 1. If there are more than one doctor with a quality of 1, payoffs will be zero due to Bertrand competition. Then, many equilibria are possible as long as two "perfect" doctors set prices of 0.

[9]For the all-price auction, "some" denotes only the player with the (strictly) highest valuation while in our model it means "all". Strictly speaking, for the all-pay auction, we do not mean maxmin payoffs, but maxmin payoffs after elimination of strictly dominated strategies.

ever, not consistent with equilibrium behavior. Thus a mixed strategy equilibrium arises.

Note that in our model in equilibrium weak players typically earn much more than their maxmin payoffs. As an **example**, consider the case $n = 2, \alpha_1 = 0.9, \alpha_2 = 0.3$. Then

$$\pi_1 = \text{maxmin}_1 = 0.9(1 - 0.3) = 0.63,$$

whereas

$$\pi_2 = 0.3(1 - 0.3) = 0.21 > \text{maxmin}_2 = 0.3(1 - 0.9) = 0.03.$$

A good doctor can make a high payoff out of his monopoly situations. Thus he is "unwilling" in face of the weak competitor to play too low prices in equilibrium. Hence there is much room left in the market for the weak doctor who can make a considerable payoff.

4 The quality setting stage

As the equilibrium of the price setting game is unique, the doctors' payoffs in every SPNE of the complete game just depend on the qualities chosen by the doctors in the first stage. From now on, we can thus consider the game as a one-stage quality setting game and search for its Nash equilibria.

For intuition, we first consider the two doctors case $n = 2$ with $\overline{\alpha} > \frac{1}{2}$. By Proposition 1, the payoff of doctor 1, given quality choices α_1 and α_2, is

$$\Pi_1(\alpha_1, \alpha_2) = \alpha_1(1 - \min(\alpha_1, \alpha_2)).$$

It is thus not surprising that the best response curve of doctor 1 contains only $\frac{1}{2}$ and $\overline{\alpha}$:

$$BR_1(\alpha_2) = \begin{cases} \frac{1}{2} & \text{if } \alpha_2 \geq 1 - \frac{1}{4\overline{\alpha}} \\ \overline{\alpha} & \text{if } \alpha_2 \leq 1 - \frac{1}{4\overline{\alpha}}. \end{cases}$$

The two equilibria in pure quality setting strategies are accordingly $(\overline{\alpha}, \frac{1}{2})$ and $(\frac{1}{2}, \overline{\alpha})$.

With more than two doctors it remains true that a doctor's best response to any pure strategies of his opponents is either $\overline{\alpha}$ or $\frac{1}{2}$. Furthermore, if $\overline{\alpha} > \frac{1}{2}$, given that one of his opponents plays $\overline{\alpha}$, doctor i will prefer playing $\frac{1}{2}$ to $\overline{\alpha}$. Technically, this happens because doctor i maximizes an expression $\alpha_i(1 - \min(\alpha_i, \overline{\alpha}))$ times a constant. Intuitively, doctor i chooses a low quality to prevent a fierce price competition with the strong competitor. We can accordingly characterize the SPNE as follows:[10]

Observation 1 (***i***) *If $\overline{\alpha} \in (\frac{1}{2}, 1)$, all SPNE in pure quality setting strategies are of the following form: One doctor i sets $\alpha_i = \overline{\alpha}$, all other doctors j set $\alpha_j = \frac{1}{2}$. Pricing strategies are as calculated in Proposition 2. We will call these equilibria* ***$\frac{1}{2}$-$\overline{\alpha}$-equilibria***.

(***ii***) *If $\overline{\alpha} < \frac{1}{2}$, the unique SPNE is given by all doctors i setting $\alpha_i = \overline{\alpha}$ and playing prices as in Proposition 2.*

The SPNE are in their form essentially robust to the introduction of convex costs as we will see in Proposition 4, only the values of qualities played have to be adjusted. An immediate corollary of Observation 1 is the following:

Corollary 1 *With $\overline{\alpha} > \frac{1}{2}$, there is no SPNE where all doctors offer the highest quality.*

For a doctor, there is a positive effect of raising α above $\frac{1}{2}$: With a higher quality, he gets recommended to more patients. Yet this effect on payoff is dominated by the negative effect of making competition in the market fiercer, if strong competitors are present. This behavior of the doctors already suggests the following conclusion which is made precise in Proposition 3: Having many doctors in the market cannot be good for welfare (i.e. the overall proportion of patients cured). The average quality offered by the doctors would be quite small (close to $\frac{1}{2}$).

We now determine the market size which maximizes social welfare in the $\frac{1}{2}$-$\overline{\alpha}$-equilibria. We find that for $\overline{\alpha} \geq \frac{3}{4}$ monopoly is optimal. For smaller $\overline{\alpha}$ having more but few doctors in the market is best.

[10] We focus on equilibria in pure quality-setting strategies. Yet there are also mixed equilibria. For instance in the two doctors case there is a symmetric mixed strategy equilibrium where both doctors mix over $[\frac{1}{2}, \overline{\alpha}]$ using the distribution function

$$F(\alpha) = \frac{4\alpha^2 - 1}{4\alpha^2}$$

for $\alpha \in [\frac{1}{2}, \overline{\alpha})$, and $F(\overline{\alpha}) = 1$. F has a falling density over $[\frac{1}{2}, \overline{\alpha}]$ and an atom of size $\frac{1}{4\overline{\alpha}^2}$ at $\overline{\alpha}$. There are also equilibria where one doctor sets a quality of $1 - \frac{1}{4\overline{\alpha}}$ while the other plays $\overline{\alpha}$ with probability $\frac{1}{3 - \frac{1}{\overline{\alpha}}}$ and $\frac{1}{2}$ with the counter probability.

Proposition 3 *As n increases, welfare in the $\frac{1}{2}$-$\overline{\alpha}$-equilibria converges to $\frac{1}{2}$ from above. For $\frac{3}{4} < \overline{\alpha} < 1$ welfare strictly decreases in $n \geq 1$. For $\frac{1}{2} < \overline{\alpha} \leq \frac{3}{4}$ welfare increases up to some finite optimal market size n^* and decreases from there on. Furthermore, the optimal market size is bounded from above by*

$$n^* \leq 10.4 + 2.3 \ln \left(1.7 + \frac{0.35}{\overline{\alpha} - \frac{1}{2}} \right).$$

To see why a market with a small number of doctors may generate more welfare than monopoly, consider as an example $\overline{\alpha} = 0.6$. If there is only one doctor, he offers this best possible quality of 0.6 (and sets a price of 1). Thus 60% of the patients get a positive report and attend this doctor. Of this fraction of patients, 60% get cured. Thus, in monopoly, 36% of all patients get cured. With more doctors, more patients get at least one positive report and attend a doctor at all.[11] This is the positive welfare effect of a larger number of doctors. Yet if the number of doctors increases, more and more doctors offer only low quality treatments ($\alpha = \frac{1}{2}$). This is the downside of a high number of doctors, which dominates if n gets larger: The positive effect of increasing the market size vanishes exponentially in n while the market share of the good doctor decreases much slower (roughly like $\frac{1}{n}$). Thus for n sufficiently large the proportion of patients cured is always above $\frac{1}{2}$. While the upper bound on the optimal market size given in the proposition is not very sharp, it makes clear that the optimal market size goes to infinity only very slowly as $\overline{\alpha}$ approaches $\frac{1}{2}$. For instance, for $\overline{\alpha} = 0.50001$ we get $n^* \leq 34$:[12] The presence of a doctor with a slightly better quality drastically reduces the optimal market size as infinitely many doctors would maximize welfare if $\overline{\alpha} = \frac{1}{2}$. The proposition is summed up in Figure 2 which depicts the proportion of patients cured as a function of market size for different values of $\overline{\alpha}$.

[11] In our example, with two doctors, we have $1-(1-0.6)\cdot(1-\frac{1}{2}) = 80\%$ of all patients getting at least one positive report and thus attending a doctor at all. Using formula (10) of the Appendix it can be seen that in this case 44.25% of patients get cured, and that the optimal number of doctors for $\overline{\alpha} = 0.6$ is 7.

[12] The exact value (which can be found numerically) is $n^* = 24$.

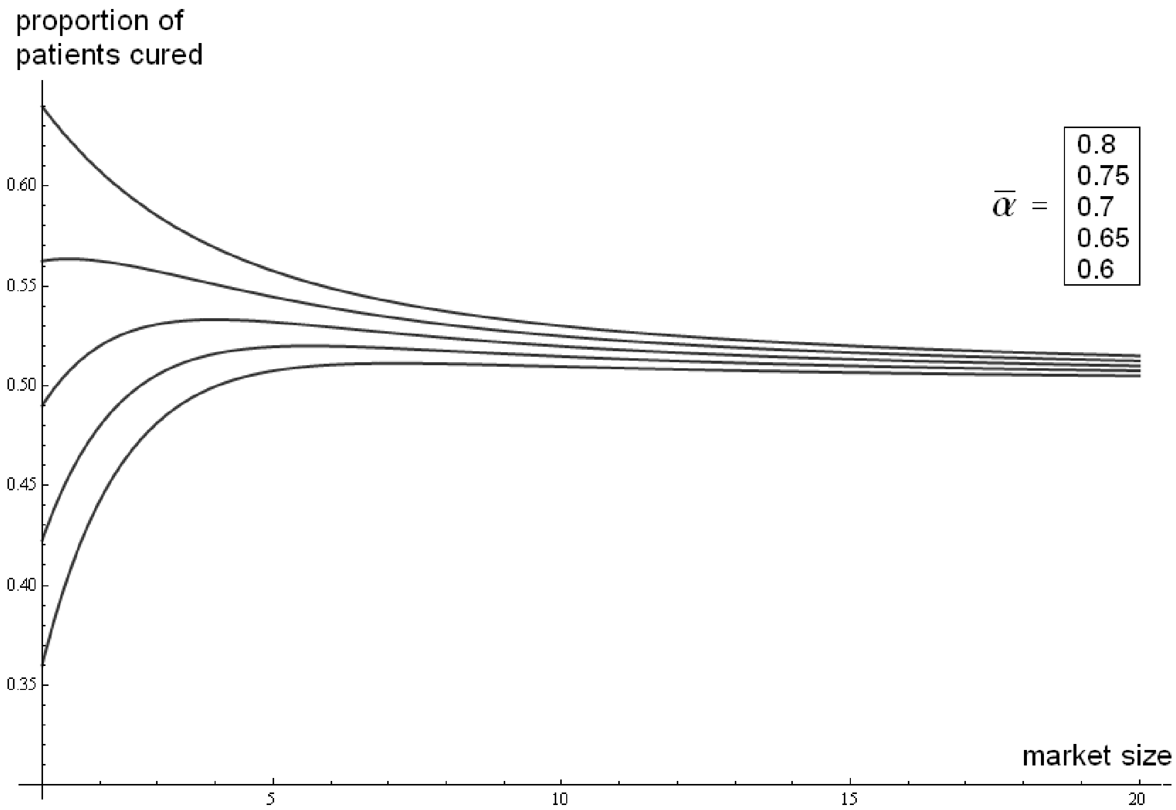

Figure 2: The effect of market size on welfare

Let us now turn to the case where the best possible approved method of healing only leads to a recovery probability of $\overline{\alpha} \leq \frac{1}{2}$. Then, as noted in Observation 1, the unique SPNE is that all doctors offer the best possible treatment with healing probability $\overline{\alpha}$. But what would happen if there was a new technology that could lead to a higher $\overline{\alpha}$? Observation 2 tells us that a rise of $\overline{\alpha}$ in between $[\frac{1}{n}, \frac{1}{2}]$ would not be welcomed by the doctors:

Observation 2 *The equilibrium profit of each doctor*

$$\overline{\alpha}(1-\overline{\alpha})^{n-1}$$

is strictly decreasing in $\overline{\alpha}$ *on* $[\frac{1}{n}, \frac{1}{2}]$.

Hence the doctors would try to block or delay the approval of promising new drugs or treatments.[13]

[13]Note, however, that doctors would, of course, welcome innovations that allowed themselves

For $\overline{\alpha} \leq \frac{1}{2}$, the proportion of patients who are healed increases with the market size. Note, however, that a restriction to a finite market size will not do much harm unless $\overline{\alpha}$ is small since the proportion of patients healed, which is given by

$$\overline{\alpha}\left[1-(1-\overline{\alpha})^{n}\right],$$

converges to $\overline{\alpha}$ exponentially fast.

5 Discussion

We will now consider some extensions of our basic model.

5.1 Costly Quality Choice

We now discuss the robustness of our results with regard to costly quality choice. We find that our results are robust to the introduction of convex costs:

Proposition 4 *Assume that setting α is associated with a cost $c(\alpha)$, where c is a continuously differentiable, increasing and convex function with $0 < c'(0) < 1$ and $c(0) = 0$. Then there exists a "$\frac{1}{2}$-1-type" SPNE, that is, an SPNE where all but one doctors set $0 < \alpha^l < \frac{1}{2}$ and one doctor sets α^h where $\alpha^h > \alpha^l$. α^l solves*

$$c'(\alpha^l) = (1-2\alpha^l)(1-\alpha^l)^{n-2}$$

and α^h solves

$$c'(\alpha^h) = (1-\alpha^l)^{n-1}$$

if $c'(1) \geq (1-\alpha^l)^{n-1}$ and $\alpha^h = 1$ otherwise. Furthermore α^l and α^h are decreasing in n.

Note that all doctors make positive profits in equilibrium no matter how large n is. Of course, if we introduce fixed costs of entry, some doctors may decide to stay out of the market. An implication of the proposition is that the huge importance of the numerical value $\alpha = \frac{1}{2}$ in the basic model which may have seemed a bit awkward was caused by the assumption of no costs.

but not the other doctors to set a higher quality.

5.2 Reversed Timing

What happens if we reverse the timing, such that doctors first choose their prices and then choose their qualities? For a given vector of (positive) prices, it is a unique best response for doctors to set their qualities to $\overline{\alpha}$. Thus the unique SPNE of the game with reversed timing is the following: Doctors mix over prices according to the distribution functions from Proposition 2 for the case of $\alpha_i = \overline{\alpha}$ and set their qualities to $\overline{\alpha}$.

The fact that doctors respond to a fixed vector of prices by setting their qualities to the maximal value also has a policy implication: If a policy maker could prescribe a fixed price for the doctors' services, the problem with the low qualities would vanish. Doctors would set their qualities as high as possible.

If we consider the simultaneous choice of qualities and prices we thus get - in the case $\overline{\alpha} > \frac{1}{2}$ - two types of Nash equilibria with pure quality setting strategies. Among these, the $\frac{1}{2}$-$\overline{\alpha}$-equilibria strictly Pareto-dominate the $\overline{\alpha}$-equilibria.

5.3 Disclosure of Qualities

Having results such as Milgrom and Roberts (1986) in mind one might expect that if doctors had a chance to disclose their qualities they would do so in equilibrium and most of our results would break down. Spiegler (2006a) addresses this by assuming that simultaneously with setting prices doctors can choose between being evaluated according to $S(1)$ or being recognized as offering a utility of α_i. Spiegler shows that any strategy involving disclosure is weakly dominated by some strategy involving no disclosure. Furthermore, in the working paper version Spiegler (2003), Spiegler shows that no Nash equilibrium involves disclosure of his quality by any doctor. The equilibrium we have identified in Theorem 2 persists because if some doctor could profitably deviate to a strategy involving disclosure there would also exist a strategy involving no disclosure he could deviate to. But then our equilibrium would not have been an equilibrium in the original game. Furthermore, because there are no equilibria involving disclosure, the equilibrium from Theorem 2 is still the unique equilibrium. There are two reasons for this difference between our model and standard models with rational patients: The first reason is that a doctor who is assumed to have a quality of 1 by a fraction α_i of patients is in a better position when competing with other doctors than a doctor who is known to have quality

α_i by all patients. The second and more fundamental reason is that unlike rational agents, our boundedly rational patients do not draw any conclusions from a doctor's decision not to disclose his quality.

5.4 More Sophisticated Reasoning and Outside Options

Let us now consider the extension that patients do not only get one report on every doctor, but several reports. Let us assume that each patient gets k_i reports on doctor i. One can imagine different ways of how these reports may be interpreted by a patient. We will consider two different interpretations.

First, we assume that a patient attributes a gross utility of zero to every doctor on which he got at least one negative report. (The negative report serves as a deterrence.) Thus a patient only attributes a positive healing potential (and then a perfect healing potential) to a doctor i if all k_i reports on doctor i's quality are positive. This does not change much of our analysis. We only have to consider that in the price setting stage, the expected equilibrium profit of each doctor i changes from $\pi_i = \alpha_i \prod_{j \neq j^*}(1 - \alpha_j)$ to $\pi_i = \alpha_i^{k_i} \prod_{j \neq j^*}(1 - \alpha_j^{k_j})$. The analoga to the $\frac{1}{2}$-$\overline{\alpha}$-equilibria of our basic model are then equilibria of the type $\sqrt[k_i]{\frac{1}{2}}$ - $\overline{\alpha}$ where all doctors but one set their qualities α_i to $\sqrt[k_i]{\frac{1}{2}} \geq \frac{1}{2}$. Note that if the numbers of reports k_i differ among the doctors, also their qualities chosen will differ: A doctor on which patients get more reports will choose a higher quality, as failures of his treatments will have larger damage on his expected profits. Hence we see that patients getting more reports help to improve the health care services of the doctors.[14]

Second, we assume that patients expect the average of the reports they get to be the doctor's true quality. This is the $S(k_i)$ rule of Osborne and Rubinstein (1998). As demonstrated in Spiegler (2003), some small change in modeling is necessary to ensure equilibrium existence. For instance, the valuations of patients for getting cured can be made a bit noisy, such that sharp discontinuities in the payoff functions of the doctors get smoothed out.[15] The price setting game becomes much more

[14]Similar reasoning shows that the introduction of an independent quality control which closes each doctors' business with probability $1 - \alpha_i$ improves the quality of the healthcare system. We assume that after qualities and prices are chosen each doctor has to go through a quality check. The result is either good or bad and made public. Then $1 - \alpha_i$ is the probability that doctor i performs badly during the check. This model is equivalent to the one discussed in the previous paragraph with $k_i = 2$. Hence we get a $\sqrt[2]{\frac{1}{2}}$ - $\overline{\alpha}$ equilibrium.

[15]For instance, without noise, there is a sharp discontinuity at the price $\frac{1}{2}$ in the S(2) model.

difficult to analyze in this case[16], but the model still shows many properties of the $S(1)$ model: The equilibrium is still in mixed strategies and there are positive maxmin payoffs. Furthermore it is still the case that the lowest prices a good doctor competing against a bad doctor is willing to set are so high that the bad doctor earns considerably more than his maxmin payoff.

Finally, let us consider the introduction of an $S(1)$-evaluated outside option, like staying in bed, that offers a healing potential of $\alpha_{out} < 1$, but at no costs. An outside option like this appears in Spiegler (2006a) where it is called the default option. Such an outside option does not affect the doctors' behavior (though it affects their equilibrium profits, which have to be scaled down by a factor of $(1 - \alpha_{out})$): With probability α_{out}, patients get a good report on the outside option, rely on their self-healing powers and stay at home. With the counter probability, patients get a negative report on the outside option and turn to the market of doctors. Even if the outside option had a healing prognosis $\overline{\alpha} > \alpha_{out} > \frac{1}{2}$, the $\frac{1}{2}$-$\overline{\alpha}$-equilibria would persist. So, there would be doctors offering treatments that made patients more likely to stay ill.[17]

5.5 Second-Price Version

Let us have a brief look at the second-price version of the auction we have considered so far. In that game the cheapest recommended doctor is chosen and the price he gets is the bid of the second-cheapest recommended doctor. If only one doctor is recommended he gets a price of 1. Clearly, in this game doctors only have to consider non-monopoly situations when making their pricing decision. It follows easily that the unique symmetric Nash equilibrium of this game is that all doctors set a price of 0. In that equilibrium each doctor only earns his maxmin payoff. There are further equilibria where one doctor i sets a price $p > 0$ while all other doctors j set a price of 0. In these equilibria all doctors but i earn strictly more than their maxmin payoffs. In the equilibrium where the best doctor n plays 1 while all other doctors j play 0, doctor n earns his maxmin payoff while the others earn

$$\pi_j = \alpha_j \prod_{k \neq j,n} (1 - \alpha_k) > \alpha_j \prod_{k \neq n} (1 - \alpha_k).$$

[16] Note for instance that under $S(2)$ depending on whether $\alpha_i < \frac{2}{3}$ or not, doctor i attains his maxmin payoff either by charging a price of $\frac{1}{2}$ or 1.

[17] Of course, behind this is the assumption that the doctors can make use of "cures" that worsen the health situation of the patient, like the prescription of a very restricted and inappropriate diet.

Note that there are no other equilibria, because given a competitor possibly plays prices above 0 playing a price of 0 is a strictly dominant strategy for all other doctors. Thus it depends on the equilibrium played whether payoffs are higher in the first or in the second price version of our auction. It is also easy to check that if the α_i differ from each other, no equilibrium in the second price auction generates the same payoffs to all doctors as the unique equilibrium of the first price auction. There is no revenue equivalence.

5.6 Rational Agents

Now we discuss two further interpretations of our model which have some interesting implications concerning firm behavior:

Our model can be seen as a model of product differentiation with fully rational consumers. This is because the Spiegler model, i.e. our price setting stage, can be reinterpreted as a variant of Perloff and Salop's (1985) model of product differentiation (see also Gabaix, Laibson and Li (2005)). In this interpretation, consumers (independently) attribute to a firm i's service a valuation of 1 with probability α_i (and a valuation of 0 otherwise).[18] Thus we reinterpret "quality" as "mass appeal" here. A higher mass appeal incurs no costs to the firms. Yet our results show that under competition most firms produce niche products, i.e. products with a lower mass appeal. A quality setting stage (or mass appeal setting stage) is not present in other papers based on the Perloff Salop model. Shaked and Sutton (1982), however, consider a quality setting stage extending the Gabszewicz and Thisse (1979) model of product differentiation.

Our model can also be seen as a model of advertising: We assume that all firms offer services consumers value with one. Consumers need some flyer to become aware of a firm. Each firm can specify the proportion α_i of consumers (but not the concrete identities of consumers) that should get a flyer. Our analysis then says that firms do not necessarily inform as many consumers as possible, even though advertising is at no costs.[19] See Butters (1977) for a seminal paper in the advertising literature

[18] Perloff and Salop consider continuous distributions instead of this Bernoulli distribution. This leads to pure strategy equilibria of the price setting game.

[19] Note, however, that under the reinterpretations, social welfare is identical no matter which firm is selected. Thus the decrease in social welfare described in Proposition 3 is not present in the reinterpretations. (There is only a decrease in welfare between monopoly and duopoly when consumers end up not consulting any firm. This may happen in the case of two firms in the mixed quality-setting equilibria.)

which considers a related model.

6 Conclusion

We have seen that if patients rely on anecdotes, all doctors, no matter how bad, survive in the market. More than that, bad doctors typically earn much more than their maxmin payoffs. If doctors can choose, they will mostly opt for low qualities: A lower quality makes the doctors attract less patients. Yet it also softens price competition, thus allows them to set higher prices in equilibrium. The latter, positive effect on payoffs dominates.

Having many doctors in the market does not help to cure the problem. Indeed, welfare falls for larger numbers of doctors, as the average quality offered decreases. Depending on the maximally possible quality, a monopoly or oligopoly of doctors is best for welfare.

Fixing prices exogenously would destroy the incentives of our doctors to choose low qualities. Then, doctors would all offer the best possible qualities. An increase in qualities offered would also be induced by making the market more transparent to the patients.

A Proofs

Proof of Proposition 1:
Wlog assume $\overline{\alpha} \geq \alpha_n \geq \alpha_{n-1} \geq \ldots \geq \alpha_1 \geq 0$.
The maxmin payoff of doctor n

$$\alpha_n \prod_{j \neq n} (1 - \alpha_j)$$

is strictly positive. So doctor n must have a strictly positive payoff in any equilibrium. Fix such an equilibrium. p_n^l, the lowest value in the support of doctor n's equilibrium strategy, must be strictly positive. Let π_n be doctor n's expected equilibrium payoff. Define a positive constant C, depending on $\alpha_1, \ldots, \alpha_n$, via $\pi_n = \alpha_n C$.

In this equilibrium, it must then be that each doctor i has an expected payoff π_i of $\alpha_i C$:

If some doctor $j < n$ earned more, say $\alpha_j D$ where $D > C$, doctor n could deviate to a price p^* that was smaller but close to p_j^l. Through this, doctor n could make a profit arbitrarily close to $\alpha_n D$ or even higher. This is because, setting the price p^*, doctor n faces as many or even less cheaper competitors as doctor j did in equilibrium at price p_j^l (as $p^* < p_j^l$). So, given he gets a good report (for which the probability is α_n), doctor n earns at least something close to D: The price is slightly lower, the possibility to be chosen by the patients is equal or higher (due to the slightly lower price).

If some doctor $j < n$ earned less than $\alpha_i C$, he could use an analogous type of deviation (thus deviate to a price slightly below p_n^l) and raise his expected profits.

Hence for any equilibrium there is some constant C such that for all i the equilibrium payoff of doctor i is given by

$$\pi_i = \alpha_i C. \tag{2}$$

Next we show that there is some doctor who does not earn more than his maxmin payoff. We start by establishing the following facts:

1) In equilibrium, no doctor places an atom in the interval $[0, 1)$: Placing an atom on $p = 0$ cannot be part of an equilibrium strategy as it leads to zero payoffs. (Recall

that we have shown above that equilibrium payoffs are positive.) Placing an atom on $p \in (0,1)$ cannot be part of an equilibrium strategy either: Assume some doctor i did so. Then consider a small price interval $(p, p+\epsilon)$. If no other doctor sets prices in this interval, doctor i is better off by shifting his atom a bit to the right. If there is some doctor j in this interval, doctor j should better shift the mass he has in the interval to some price slightly left from p. This would give doctor j a slightly lower price in case patients buy his services, but the probability that patients choose him increases substantially. (If ϵ is small enough, the second effect will dominate the first one.)

2) In equilibrium, at most one doctor will place an atom on $p = 1$: If several doctors had an atom on $p = 1$, each of them would have an incentive to deviate and shift his atom to a slightly lower value.[20]

3) In equilibrium, the support of some doctor's pricing strategy must go up to 1: If $p^h < 1$ was the highest price in the supports of the doctors' strategies, some doctor whose support went up to p^h could earn more by shifting probability mass from a neighborhood of p^h to 1. This would give him a substantially higher price while only slightly diminishing his chances of winning. (Recall that no doctor would set an atom on $p^h < 1$.)

We can now show that, in equilibrium, there is some doctor k^* who only earns his maxmin payoff: From 3), we know that there is a doctor whose strategy support goes up to 1. This doctor can only earn more than his maxmin payoff in $p = 1$ if he has a competitor who sets an atom on $p = 1$. But then, this competitor cannot earn more than his maxmin payoff (since there cannot be two doctors setting an atom on 1, compare 2)).

Thus, there is some doctor k^* earning his maxmin payoff $\pi_{k^*} = \alpha_{k^*} \prod_{j \neq k^*}(1 - \alpha_j)$ implying $C = \prod_{j \neq k^*}(1 - \alpha_j)$ for all doctors in equilibrium. Thus, in equilibrium,

$$\pi_i = \alpha_i C = \alpha_i \prod_{j \neq k^*}(1 - \alpha_j)$$

[20]This argument would not be valid in Varian's (1980) model of sales where, so to say, either all or only one doctor (seller) is recommended. There (as long as other doctors set lower prices with certainty) two or more doctors setting an atom in 1 do not stand in competition because each of them is only attended by patients who are exclusively aware of him. In contrast, in our model for any subset of doctors there is a group of patients who are aware of these doctors and no others. This is the reason for the contrast between our unique equilibrium and the multiplicity of equilibria in the Varian model pointed out by Baye, Kovenock and de Vries (1992).

for all i. It remains to be shown that $k^* \in \text{argmax}_j \alpha_j$:

Assume $\alpha_{k^*} < \alpha_n$. Then the payoff of doctor n is

$$\pi_n = \alpha_n \prod_{j \neq k^*} (1 - \alpha_j),$$

which is strictly smaller than doctor n's maxmin payoff $\alpha_n \prod_{j \neq n}(1 - \alpha_j)$ since $1 - \alpha_{k^*} > 1 - \alpha_n$. This would give doctor n an incentive to deviate. Thus we get $k^* = n$. □

Proof of Proposition 2

(***i***) Let $\pi_i(p)$ denote the payoff of doctor i given that he chooses p with certainty and given that the other doctors mix according to the distribution functions G_k described in the proposition. Clearly

$$\pi_i(p) = p\alpha_i \prod_{k \neq i}(1 - \alpha_k G_k(p)). \tag{3}$$

We have to show that, for all i, there exists a set S_i with mass 1 under G_i so that π_i is constant on S_i and π_i is weakly greater on S_i than on S_i^C. We will show this with $S_n = [p_0, p_n]$ and with $S_i = [p_0, p_i)$ for $i < n$.

For $p \in [0, p_0)$ $\pi_i(p) \leq \pi_i(p_0)$ as the other doctors do not put any mass on $[0, p_0)$: Deviating below p_0 gives the same chances of attracting a patient as playing p_0 but at a smaller price.

Now we show that $\pi_i(p)$ is constant on $[p_0, p_i] \setminus \{1\}$. The case $p = 1$ will be treated separately afterwards. By (3) and the definition of the G_k

$$\pi_i(p) = p\alpha_i(1 - \alpha_1) \cdot \ldots \cdot (1 - \alpha_{j-1}) \left[\sqrt[n-j]{\frac{(1 - \alpha_j) \cdot \ldots \cdot (1 - \alpha_{n-1})}{p}} \right]^{n-j}$$

for $p \in [p_{j-1}, p_j] \subset [p_0, p_i]$ where we use the fact that $n - j$ of i's opponents have p in the supports of their pricing strategies, while the $j - 1$ remaining doctors put all their probability mass below p. This immediately implies that

$$\pi_i(p) = \alpha_i(1 - \alpha_1) \cdot \ldots \cdot (1 - \alpha_{n-1}). \tag{4}$$

The case $p = 1$ poses a minor technical difficulty due to tie-breaking: If doctor n does not have an atom on 1 (i.e. $\alpha_{n-1} = \alpha_n$), the calculation for $p < 1$ still goes through. If doctor n has an atom on 1, it is easy to check that doctors $i \neq n$ with $p_i = 1$ earn less than the payoff from (4) when playing 1. But this does not give them an incentive to deviate from playing G_i as the probability of playing 1 is zero under G_i. Doctor n plays 1 with positive probability but he does not face opponents who do: He gets his maxmin payoff which corresponds to his payoff from (4) when playing $p = 1$. Thus doctor n does not have an incentive to deviate either.

To conclude the proof, we have to show that doctor i with $p_i < 1$ does not have an incentive to set prices $p \in (p_i, 1]$: Consider $p \in [p_{j-1}, p_j]$ with $j > i$ (so that $p_i \leq p_{j-1}$). By (3) we have

$$\begin{aligned}\pi_i(p) &= p\alpha_i \left[\sqrt[n-j]{\frac{(1-\alpha_j)\cdot\ldots\cdot(1-\alpha_{n-1})}{p}} \right]^{n-j+1} \prod_{\substack{1 \leq k \leq j-1 \\ k \neq i}} (1-\alpha_k) \\ &= \alpha_i \sqrt[n-j]{\frac{(1-\alpha_j)\cdot\ldots\cdot(1-\alpha_{n-1})}{p}} \prod_{\substack{k \neq n \\ k \neq i}} (1-\alpha_k).\end{aligned}$$

The above expression is strictly decreasing in p, so that it is sufficient to consider deviations to p_{j-1}. Using the definition of p_{j-1} we see that

$$\pi_i(p_{j-1}) = \alpha_i(1-\alpha_{j-1}) \prod_{\substack{k \neq n \\ k \neq i}} (1-\alpha_k).$$

This is weakly smaller than the payoff from (4) as $p_i \leq p_{j-1}$ implies $1-\alpha_i \geq 1-\alpha_{j-1}$. Thus, doctor i does not have an incentive to deviate to prices above p_i.

(***ii***) This proof proceeds in a number of steps. We start with some preliminary observations:

First, recall from the proof of Proposition 1 that in any equilibrium doctors will not place atoms except for one doctor possibly putting an atom on 1. From Proposition 1 we also know that, for all i, doctor i's equilibrium payoff is given by

$$\pi_i = \alpha_i C \text{ where } C := \prod_{j=1}^{n-1} (1-\alpha_i) > 0. \tag{5}$$

Note that the first equality states that all doctors' expected equilibrium payoffs *conditional on being recommended* must be identical. Furthermore, recall from the proof of Proposition 1 that the union of the agents' strategies' supports must go up to 1. Due to the positive equilibrium payoffs this union must be bounded away from zero. Note also that the union of equilibrium payoffs must be an interval $[p_L, 1]$, i.e. there cannot be any gaps in the union of supports: If there was an interval $[\underline{p}, \overline{p}]$, $\underline{p} > p_L$, where no doctor was active, a doctor who would be playing prices right below $\underline{p}$ could deviate by moving probability mass from a small interval below $\underline{p}$ to $\overline{p}$ yielding a substantially better price at a marginally lower probability of winning. Thus the union of supports must be an interval $[p_L, 1]$. Note also that there cannot be a subset $[\underline{p}, \overline{p}] \subset [p_L, 1]$ where only one doctor is active: That doctor could profitably deviate by concentrating all probability mass of the interval in an atom at $\overline{p}$ which yields a higher price at the same probability of winning.

Armed with these insights we turn to the first major step in the proof:

1) In any equilibrium, the support of every doctor goes down to the same $p_L > 0$. Furthermore, in any equilibrium, $p_L = C$.

Proof of 1): Consider two doctors i and j with supports S_i and S_j. Assume $p_L^i < p_L^j$ where $p_L^k = \inf S_k$ for $k = i, j$. Note that with positive probability doctor i plays a price from $[p_L^i, p_L^j]$ and that agent j's payoff from playing p_L^j equals $\alpha_j C > 0$. But this implies that doctor i can earn more than his equilibrium payoff of $\alpha_i C$ by playing p_L^j: Since - unlike doctor j - doctor i does not have himself as a possibly cheaper competitor when playing p_L^j, his expected payoff conditional on being recommended must be higher than that of j. This is, however, a contradiction to (5) and thus the support of every doctor goes down to the same p_L. To see that $p_L = C$, note that, for all j, doctor j's payoff from playing p_L must be $\alpha_j p_L$: The other doctors ask for higher prices with probability 1 and thus doctor j gets all the patients to which he was recommended and they pay him p_L. These payoffs of $\alpha_j p_L$ are, however, only consistent with (5) if $p_L = C$.

The next step further characterizes the functional form of the doctor's equilibrium distribution functions:

2) Let $\mathcal{D} \subset \{1, \ldots, n\}$ denote the set of doctors who are active on some interval $I = (\underline{p}, \overline{p})$ in some arbitrary but fixed equilibrium. Assume all doctors $j \in \mathcal{D}$ are active at any $p \in I$ and let $m = \#\mathcal{D}$. (Note that from our preliminary observations

we have $m \neq 1$.) We will show that for all $j \in \mathcal{D}$ the equilibrium distribution functions $G_j(p)$ must satisfy for all $p \in I$

$$G_j(p) = \frac{1}{\alpha_j}\left(1 - \sqrt[m-1]{\frac{H}{p}}\right) \tag{6}$$

where $H > 0$ depends on the α_i and on the probability mass placed below $\underline{p}$ by the doctors from $\mathcal{D}^C$, namely,

$$H = \frac{C}{\prod_{i \in \mathcal{D}^C}(1 - \alpha_i G_i(\underline{p}))}.$$

Proof of 2): To see this, note that for all $j \in \mathcal{D}$ and all $p \in I$ the expected payoff of doctor j from playing p must equal the equilibrium payoff of $\alpha_j C$. Using (3) and the fact that distribution functions of the inactive doctors are constant over I this condition reads

$$\alpha_j C = p\alpha_j \left[\prod_{i \in \mathcal{D}^C}(1 - \alpha_i G_i(\underline{p}))\right]\left[\prod_{k \in \mathcal{D}\setminus\{j\}}(1 - \alpha_k G_k(p))\right].$$

Rearranging and using the definition of H yields for all $p \in I$ and $j \in \mathcal{D}$

$$\prod_{k \in \mathcal{D}\setminus\{j\}}(1 - \alpha_k G_k(p)) = \frac{H}{p}. \tag{7}$$

Now consider (7) for two different doctors $j_1, j_2 \in \mathcal{D}$. Taking the quotient of (7) for j_1 and (7) for j_2 yields that for all $p \in I$

$$1 = \frac{1 - \alpha_{j_2} G_{j_2}(p)}{1 - \alpha_{j_1} G_{j_1}(p)}$$

which implies that there is a function $h(p)$ such that $h(p) = \alpha_k G_k(p)$ for all $k \in \mathcal{D}$. Substituting $h(p)$ for $\alpha_k G_k(p)$ on the lhs of (7) and then solving for h yields

$$h(p) = 1 - \sqrt[m-1]{\frac{H}{p}}$$

and thus

$$G_j(p) = \frac{1}{\alpha_j}\left(1 - \sqrt[m-1]{\frac{H}{p}}\right)$$

as required.

The last main step before we can conclude the proof shows that all doctors' supports are intervals, i.e. no doctor will be inactive over some range of prices (above p_L) but put positive probability mass on prices above that range:

3) For all j the support of doctor j's strategy is of the form $[p_L, p_H^j]$ for some $p_L < p_H^j \leq 1$.

Proof of 3): The proof is by contradiction. Assume that some price $\overline{p}$ is in the support of the strategy of doctor j but j is inactive on some interval below $\overline{p}$. Choose $\underline{p} < \overline{p}$ such that for all $p \in I = (\underline{p}, \overline{p})$ the set of doctors who are active at p is identical. (This is possible since there are no atoms and thus the G_i are continuous.) Denote the set of doctors active on I by $\mathcal{D}$. Using (3) as in step 2) we can write the payoff of player j from playing some $p \in I \cup \{\overline{p}\}$ as

$$\pi_j(p) = \alpha_j p \left[\prod_{i \in \mathcal{D}^C \setminus \{j\}} (1 - \alpha_i G_i(\underline{p})) \right] \left[\prod_{k \in \mathcal{D}} (1 - \alpha_k G_k(p)) \right]$$

Defining the constant factor from the other inactive doctors as

$$K := \left[\prod_{i \in \mathcal{D}^C \setminus \{j\}} (1 - \alpha_i G_i(\underline{p})) \right].$$

and using (6) from the last step, we can express $\pi_j(p)$ as

$$\pi_j(p) = \alpha_j p K \left(\sqrt[m-1]{\frac{H}{p}} \right)^m = \alpha_j K H^{\frac{m}{m-1}} \sqrt[m-1]{\frac{1}{p}}$$

where the constant H is as defined in Step 2. Note that this implies that $\pi_j(p)$ is strictly decreasing in p over $I \cup \{\overline{p}\}$. By assumption, doctor j is active at $\overline{p}$ and thus earns his equilibrium payoff there:

$$\pi_j(\overline{p}) = \alpha_j C.$$

But since $\pi_j(p)$ is decreasing, this implies that for $p \in I$

$$\pi_j(p) > \alpha_j C.$$

such that doctor j can profitably deviate - which is a contradiction.

To conclude the proof, we still have to show that the vector of strategies considered in (***i***) is actually the only candidate for an equilibrium. We have seen that all supports start at $p_L = C$ and since doctors do not set atoms or leave gaps in their supports, all doctors remain active up to the price p_1 where the first doctor(s) j have $G_j(p_1) = 1$. Note that on any interval $[p_L, \overline{p}]$ where all doctors are active, all distribution functions are uniquely pinned down by Step 2. This also uniquely determines p_1 and the set of agents who have $G_j(p_1) = 1$. Above p_1, all doctors who still have probability mass to spend must remain active. By Step 2, distribution functions above p_1 are again uniquely determined, determining in turn uniquely the price $p_2 > p_1$ where the next supports end. Continuing this procedure sequentially until $p = 1$ or until all or all but one distribution functions equal 1 determines a unique candidate for an equilibrium. It is easy to check that this unique candidate is actually the vector of strategies considered in (***i***). □

Proof of Proposition 3

The quantity of interest in this proof is $w(\overline{\alpha}, n)$, the proportion of patients healed in the $\frac{1}{2}$-$\overline{\alpha}$-equilibria with n doctors where $n \geq 1$ and $\frac{1}{2} < \overline{\alpha} \leq 1$. Clearly, $w(\overline{\alpha}, n)$ can be written as

$$w(\overline{\alpha}, n) = p_g \overline{\alpha} + p_b \frac{1}{2} + p_0 0 \tag{8}$$

where p_g, p_b and p_0 denote the fractions of patients consulting the good doctor (i.e. the doctor offering $\overline{\alpha}$), the fraction consulting the other doctors and the fraction who stays at home, respectively. Note that p_g, p_b and p_0 are unique since the price setting game has a unique equilibrium by Proposition 2 and that they can be calculated from the equilibrium strategies given there: The good doctor mixes over $[2^{-(n-1)}, 1]$ using the distribution function

$$F_g(p) = \frac{1}{\overline{\alpha}} \left(1 - \sqrt[n-1]{\frac{1}{2^{n-1} p}} \right)$$

and puts an atom of size $1 - \frac{1}{2\overline{\alpha}}$ on 1. The remaining doctors mix over the same interval with

$$F_b(p) = 2 \left(1 - \sqrt[n-1]{\frac{1}{2^{n-1} p}} \right).$$

Note that the pricing strategy of the good doctor can be interpreted in the following way: With probability $1-\frac{1}{2\overline{\alpha}}$ he sets a price of 1 and with probability $\frac{1}{2\overline{\alpha}}$ he uses exactly the same pricing strategy as the other doctors. In the first case the good doctor only gets patients if no other doctor is recommended (which happens with probability $2^{-(n-1)}$). In the second case, the good doctor has exactly the same chances to acquire a patient as the other recommended doctors. Let the random variable r_n denote the number of bad doctors who are recommended to a patient. Then the market share of the good doctor can be written as

$$p_g = \overline{\alpha}\left[\frac{1}{2\overline{\alpha}}E\left[\frac{1}{1+r_n}\right] + (1-\frac{1}{2\overline{\alpha}})2^{-(n-1)}\right] \tag{9}$$

where the leading factor $\overline{\alpha}$ results from the fact that the doctor is only competing for the patients to which he is recommended. Note that r_n is distributed binomially with parameters $n-1$ and $\frac{1}{2}$ and thus

$$\begin{aligned} E\left[\frac{1}{1+r_n}\right] &= \sum_{k=0}^{n-1}\binom{n-1}{k}2^{-(n-1)}\frac{1}{1+k} = \frac{1}{n}\sum_{k=0}^{n-1}\binom{n}{k+1}2^{-(n-1)} \\ &= \frac{2}{n}\sum_{k=1}^{n}\binom{n}{k}2^{-n} = \frac{2}{n}\left(-2^{-n}+\sum_{k=0}^{n}\binom{n}{k}2^{-n}\right) \\ &= \frac{2}{n}(1-2^{-n}) \end{aligned}$$

where in the first step we have used that $n\binom{n-1}{k} = (k+1)\binom{n}{k+1}$ and in the final step we used that the sum equals 1 since it simply adds up all probabilities of a Binomial $(n, \frac{1}{2})$ distribution. Putting this into (9) and rearranging terms gives us

$$p_g = \frac{1}{n}(1-2^{-n}) + (\overline{\alpha}-\frac{1}{2})2^{-(n-1)}.$$

Clearly $p_0 = (1-\overline{\alpha})2^{-(n-1)}$. Thus we see that - as claimed in the main text - as n gets large p_g goes to zero like $\frac{1}{n}$ and thus much slower than p_0 which decreases exponentially. Inserting these expressions for p_g and p_0 and $p_b = 1-p_g-p_0$ into (8) and rearranging we obtain

$$w(\overline{\alpha}, n) = \frac{1}{2} + (\overline{\alpha}-\frac{1}{2})\frac{1}{n}(1-2^{-n}) + (2\overline{\alpha}^2-\overline{\alpha}-\frac{1}{2})2^{-n}. \tag{10}$$

The remainder of the proof consists of an analysis of the function $w(\overline{\alpha}, n)$.[21] We will

[21] While this rather technical analysis is of course necessary to complete the proof, the impatient

proceed in the following way: First, we will show that for every α as n gets large w approaches $\frac{1}{2}$ from above. Then we will show that $w(\overline{\alpha}, x)$, where $x \in \mathbb{R}_{\geq 1}$, is monotonically decreasing in x for $\alpha \geq \frac{1}{4}(1+\sqrt{5}) \approx 0.809$. For $\frac{1}{2} < \overline{\alpha} < \frac{1}{4}(1+\sqrt{5})$ we will show that w increases in x up to some value $x^* \geq 1$ and decreases from there on. From this we can conclude that $w(\overline{\alpha}, 1) \geq w(\overline{\alpha}, 2)$ is a sufficient condition for $n^* = 1$ being a maximizer of $w(\overline{\alpha}, n)$. Then we will verify that $w(\overline{\alpha}, 1) \geq w(\overline{\alpha}, 2)$ if and only if $\overline{\alpha} \geq \frac{3}{4}$. From these results we can conclude that for $\overline{\alpha} > \frac{3}{4}$ $w(\overline{\alpha}, n)$ is maximized at $n^* = 1$ while for $\overline{\alpha} \geq \frac{3}{4}$ there is a finite $n^* > 1$ which maximizes $w(\overline{\alpha}, n)$. Furthermore it follows that there are at most two maximizers and if there are two, these must be subsequent integers. In the final part of the proof we will show an upper bound on n^* in terms of $\overline{\alpha}$.

It is immediate that $w(\overline{\alpha}, n)$ converges to $\frac{1}{2}$ as n goes to infinity. To show that this convergence is from above we have to show for every fixed $\overline{\alpha}$ that $w(\overline{\alpha}, n) > \frac{1}{2}$ for sufficiently large n. Here and in the following we will substitute n by the real-valued parameter x and view w as the weighted sum of two functions f and g which do not depend on $\overline{\alpha}$:

$$w(\overline{\alpha}, x) = \frac{1}{2} + (\overline{\alpha} - \frac{1}{2})f(x) + (2\overline{\alpha}^2 - \overline{\alpha} - \frac{1}{2})g(x)$$

where

$$f(x) = \frac{1}{x}(1 - 2^{-x}) \text{ and } g(x) = 2^{-x}.$$

Note that the coefficient of f is always positive while the coefficient of g is zero at $\overline{\alpha} = \frac{1}{4}(1+\sqrt{5})$ and strictly increasing over $[\frac{1}{2}, 1]$. Note furthermore that g and f are strictly decreasing and positive. For g this is clear, for f note that

$$f'(x) = -[2^x - (1 + x\ln(2))]x^{-2}2^{-x}.$$

The factors outside the brackets are clearly positive. The term in squared brackets is positive for all $x \geq 1$ since it is the difference between the convex function 2^x and its first order Taylor approximation in 0. Thus $f'(x) < 0$ for all $x \geq 1$.

Note that in the case $\overline{\alpha} \geq \frac{1}{4}(1+\sqrt{5})$ where the coefficients of f and g are both non-negative it is immediate that w is strictly decreasing in n, always greater than $\frac{1}{2}$ and

reader is invited to skip it. Ultimately the calculations only verify that Figure 2 delivers the complete picture.

maximized by $n^* = 1$. The case where the coefficient of g is negative requires more work however. The key observation driving the argumentation that follows now is that f decreases much slower than g and thus the term with the positive coefficient will eventually dominate - even if for $\overline{\alpha}$ near $\frac{1}{2}$ this coefficient is very small.

Note first that we can rewrite w to

$$w(\overline{\alpha}, x) = \frac{1}{2} + g(x)\left[(\overline{\alpha} - \frac{1}{2})\frac{f(x)}{g(x)} + (2\overline{\alpha}^2 - \overline{\alpha} - \frac{1}{2})\right].$$

In order to show that this is greater than $\frac{1}{2}$ for x sufficiently large (and finite) it is sufficient to show that f/g tends to infinity as x gets large because this guarantees that the first (positive) summand in the squared brackets will eventually dominate the second one making the term in the squared brackets positive. Now f/g is easily calculated to be

$$\frac{f(x)}{g(x)} = \frac{2^x - 1}{x}$$

which obviously tends to infinity as x gets large.

Next we show for fixed $\overline{\alpha}$ that $w(\overline{\alpha}, x)$ has exactly one local maximum in x. This is equivalent to showing that (depending on $\overline{\alpha}$) the x-derivative of $w(\overline{\alpha}, x)$ is either negative for all x or changes signs exactly once on $[1, \infty)$ from positive to negative. Note that we can write this derivative as

$$\frac{\partial w}{\partial x}(\overline{\alpha}, x) = g'(x)\left[(\overline{\alpha} - \frac{1}{2})\frac{f'(x)}{g'(x)} + (2\overline{\alpha}^2 - \overline{\alpha} - \frac{1}{2})\right]. \quad (11)$$

Recalling that $g' < 0$ it is clear that it is sufficient to prove that f'/g' is monotonically increasing and tends to infinity. It is easily calculated that

$$\frac{f'(x)}{g'(x)} = \frac{-1 + 2^x - x\ln(2)}{x^2\ln(2)} \quad (12)$$

which clearly tends to infinity as x gets large. In order to show monotonicity consider the derivative of f'/g' which can be written as

$$\frac{d}{dx}\frac{f'(x)}{g'(x)} = \frac{2^x(x\ln(2) - 2) - (-x\ln(2) - 2)}{x^3\ln(2)}.$$

To determine the sign of this expression we can concentrate on the numerator. Note that the numerator is exactly the difference between the function $2^x(x\ln(2)-2)$ and its first order Taylor approximation around 0. Since $2^x(x\ln(2)-2)$ is strictly convex on $(0,\infty)$ (its second derivative is $2^x x\ln(2)$) this difference is positive for $x > 0$. Thus f'/g' is monotonically increasing for $x \geq 0$ as desired. Thus we have shown for every $\overline{\alpha} > \frac{1}{2}$ that $w(\overline{\alpha}, x)$ has a unique maximizer $1 \leq x^* < \infty$.

Now we will consider $w(\overline{\alpha}, n)$ with an integer parameter n again. From the previous analysis it is clear that $w(\overline{\alpha}, n)$ is globally maximized by $n = 1$ if and only if

$$w(\overline{\alpha}, 1) - w(\overline{\alpha}, 2) \geq 0$$

(and $n = 1$ is the unique maximizer if the inequality holds strictly). Now we have

$$w(\overline{\alpha}, 1) - w(\overline{\alpha}, 2) = \frac{1}{2}\overline{\alpha}^2 + \frac{1}{8}\overline{\alpha} + \frac{3}{16}.$$

Since this is a quadratic polynomial, it is easily seen that it is increasing over $[\frac{1}{2}, 1]$ and zero for $\overline{\alpha} = \frac{3}{4}$. Thus $n^* = 1$ for $\overline{\alpha} \geq \frac{3}{4}$ and $n^* > 1$ for $\overline{\alpha} < \frac{3}{4}$. That for fixed $\overline{\alpha}$ there are at most two integer maximizers of $w(\overline{\alpha}, n)$ and that, if there are two, those must be subsequent integers follows trivially from the fact that $w(\overline{\alpha}, x)$ has a unique real-valued local (and thus global) maximizer.

In the final part of the proof we show an upper bound on n^* in terms of $\overline{\alpha}$. We will again consider the function $w(\overline{\alpha}, x)$ with real valued argument $x \geq 1$. Note that since $w(\overline{\alpha}, x)$ has a unique local maximum for fixed $\overline{\alpha}$, any point x where the x-derivative of $w(\overline{\alpha}, x)$ is negative must lie to the right of the maximizer x^*. We will construct a function $B(\overline{\alpha})$ with the property that

$$x > B(\overline{\alpha}) \Rightarrow \frac{\partial}{\partial x} w(\overline{\alpha}, x) < 0.$$

This implies $x^* < B(\overline{\alpha})$ and thus $n^* < B(\overline{\alpha}) + 1$.

Note that from (11) and (12) $\frac{\partial}{\partial x} w(\overline{\alpha}, x) < 0$ is equivalent to

$$\frac{-1 + 2^x - x\ln(2)}{x^2 \ln(2)} > \frac{\overline{\alpha} + \frac{1}{2} - 2\overline{\alpha}^2}{\overline{\alpha} - \frac{1}{2}}.$$

We will now try to find a sufficient condition for this which is of the desired form.

Note that the numerator of the rhs only fluctuates between $\frac{1}{2}$ and $-\frac{1}{2}$ and thus a sufficient condition is

$$\frac{2^x}{x^2\ln(2)} > \frac{1 + x\ln(2)}{x^2\ln(2)} + \frac{1}{2(\overline{\alpha} - \frac{1}{2})}.$$

Now note that the first term on the rhs is at most $(1 + \ln(2))/\ln(2)$ and thus a sufficient condition is

$$\frac{2^x}{x^2} > 1 + \ln(2) + \frac{\ln(2)}{2(\overline{\alpha} - \frac{1}{2})}.$$

Taking logarithms on both sides yields

$$x\ln(2) - 2\ln(x) > \ln\left(1 + \ln(2) + \frac{\ln(2)}{2(\overline{\alpha} - \frac{1}{2})}\right).$$

Since $\ln(x)$ is concave we can bound it from above by its first order Taylor approximation in 8 which is $\ln(8) + (1/8)x$.[22] Thus a sufficient condition is

$$x > \frac{2\ln(8)}{\ln(2) - \frac{1}{4}} + \frac{1}{\ln(2) - \frac{1}{4}}\ln\left(1 + \ln(2) + \frac{\ln(2)}{2(\overline{\alpha} - \frac{1}{2})}\right).$$

Since this bound is not very sharp anyway we can afford to improve readability by bounding the logarithms by real numbers in a way that the condition remains sufficient and get

$$x > 9.4 + 2.3\ln\left(1.7 + \frac{0.35}{\overline{\alpha} - \frac{1}{2}}\right) =: B(\overline{\alpha}).$$

As argued above this implies

$$n^* \leq 10.4 + 2.3\ln\left(1.7 + \frac{0.35}{\overline{\alpha} - \frac{1}{2}}\right).$$

□

Proof of Observation 2

This observation follows immediately from the fact that

$$\frac{\partial}{\partial\overline{\alpha}}\,\overline{\alpha}(1 - \overline{\alpha})^{n-1} = (1 - n\overline{\alpha})(1 - \overline{\alpha})^{n-2}$$

[22]The choice of 8 as the expansion point is essentially arbitrary, it only matters that the value is large enough so that the lhs of the equation remains increasing in x.

is negative for $\overline{\alpha} \in (\frac{1}{n}, 1)$. $\square$

Proof of Proposition 4

We first search for a symmetric Nash equilibrium of the quality setting game played by the doctors 1,... $n-1$ given that doctor n chooses a higher α (and thus does not affect the payoffs of doctors 1,... $n-1$). Then we show that doctor n will indeed respond to this behavior by setting a higher quality than the other doctors.

Given that doctor n chooses a high α_n, doctors $i < n$ maximize

$$\alpha_i(1-\alpha_i)\left[\prod_{j\neq i,n}(1-\alpha_j)\right] - c(\alpha_i)$$

which yields a first order condition of

$$(1-2\alpha_i)\left[\prod_{j\neq i,n}(1-\alpha_j)\right] - c'(\alpha_i) = 0.$$

Assuming symmetry for doctors 1,... $n-1$ this becomes

$$(1-2\alpha)(1-\alpha)^{n-2} = c'(\alpha).$$

There is a unique $\alpha^l \in (0, \frac{1}{2})$ which solves this equation: the lhs decreases strictly on $(0, \frac{1}{2})$ taking all values between 1 and 0 while the rhs is strictly increasing starting with $0 < c'(0) < 1$ as c is assumed to be convex.

We now determine doctor n's best response to the other doctors playing α_l. Note that doctor n will not play qualities below α_l: On $[0, \alpha^l]$ doctor n has the same optimization problem as the other doctors in the previous step,

$$\max_{\alpha} \alpha(1-\alpha)(1-\alpha^l)^{n-2} - c(\alpha)$$

which is solved by $\alpha = \alpha^l$. On $[\alpha^l, 1]$ doctor n maximizes

$$\alpha(1-\alpha^l)^{n-1} - c(\alpha)$$

which gives a FOC of

$$(1-\alpha^l)^{n-1} = c'(\alpha).$$

Note that the other doctors' FOC implies that

$$(1-\alpha^l)^{n-1} > c'(\alpha^l).$$

Thus we have to distinguish two cases: If

$$(1-\alpha^l)^{n-1} < c'(1),$$

by the monotonicity of c', there is a unique best response $\alpha^h \in (\alpha^l, 1)$ which solves

$$(1-\alpha^l)^{n-1} = c'(\alpha^h).$$

If

$$(1-\alpha^l)^{n-1} \geq c'(1),$$

it is optimal for doctor n to play $\alpha^h = 1$. Finally, note that doctors 1,..., $n-1$ do not want to deviate to qualities above α^h: If this were a profitable deviation for one of them, α^h would not have been optimal for doctor n (who faces weaker competitors than the other doctors).

To conclude the proof we have to show that α^l and α^h are decreasing in n: To make the dependence on n clearer we write $\alpha^l(n)$ and $\alpha^h(n)$. Recall that from the bad doctors' FOC $\alpha^l(n)$ was the value of α where $c'(\alpha)$ and $(1-2\alpha)(1-\alpha)^{n-2}$ intersect. Increasing n to $n+1$ shifts the function $(1-2\alpha)(1-\alpha)^{n-2}$ downwards, implying that it intersects c' at a smaller value of α since c' is increasing. Thus $\alpha^l(n)$ is decreasing in n.

We now show that $\alpha^h(n)$ is weakly decreasing: As a preliminary observation, note that $c'(\alpha^l(n))$ is decreasing in n and thus by the bad doctors' FOC

$$(1-2\alpha^l(n))(1-\alpha^l(n))^{n-2} \tag{13}$$

is decreasing as well. If for some n the good doctor's FOC is not binding (which implies $\alpha^h(n) = 1$) we obviously have $\alpha^h(n) \geq \alpha^h(n+1)$. Now consider some n where the good doctor's FOC is binding so that $\alpha^h(n)$ solves

$$(1-\alpha^l(n))(1-\alpha^l(n))^{n-2} = c'(\alpha^h(n)). \tag{14}$$

Clearly, to show that $\alpha^h(n) \geq \alpha^h(n+1)$ we need to prove that the lhs of (14) is decreasing: Recall that (13) is decreasing and note that the leading factor $(1-2\alpha^l(n))$ is increasing. Thus the second factor $(1 - \alpha^l(n))^{n-2}$ must be decreasing strongly enough to make the product (13) decreasing. The lhs of (14) has the same second factor as (13). The first factor $(1 - \alpha^l(n))$ is however increasing less strongly than the first factor of (13). Thus the lhs of (14) is dominated by its second, decreasing factor. $\square$

References

[1] Baye, Michael, Dan Kovenock and Casper de Vries. It takes two to tango: Equilibria in a model of sales. Games and Economic Behavior, 4: 493-510, 1992.

[2] Baye, Michael, Dan Kovenock and Casper de Vries. The all-pay-auction with complete information. Economic Theory, 8: 91-305, 1996.

[3] Butters, Gerard. Equilibrium Distributions of Sales and Advertising Prices. Review of Economic Studies, 44(3):465-491, 1977.

[4] Ellison, Glenn. Bounded Rationality in Industrial Organization. In: R. Blundell, W. Newey and T. Persson (eds.), Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, Cambridge University Press, 2006.

[5] Ellison Glenn and Drew Fudenberg. Word-of-Mouth Communication and Social Learning. Quarterly Journal of Economics, 110, 93-125, 1995.

[6] Gabaix, Xavier, David Laibson and Hongyi Li. Extreme Value Theory and the Effects of Competition on Profits. Working Paper, 2005.

[7] Gabszewicz, Jean and Jacques-François Thisse. Price Competition, Quality and Income Disparities. Journal of Economic Theory, 20: 340-359, 1979.

[8] Milgrom, Paul and John Roberts. Relying on Information of Interested Parties. Rand Journal of Economics, 17, 18-32, 1986.

[9] Osborne, Martin and Ariel Rubinstein. Games with Procedurally Rational Players. American Economic Review, 88: 834-849, 1998.

[10] Perloff, Jeffrey M. and Steven C. Salop. Equilibrium with Product Differentiation. Review of Economic Studies, 52(1): 107-120, 1985.

[11] Rabin, Matthew. Inference by Believers in the Law of Small Numbers, Quarterly Journal of Economics, 117, 775-816, 2002.

[12] Rubinstein, Ariel and Ran Spiegler. Money Pumps in the Market. Journal of the European Economic Association, 6(1): 237-253, 2008.

[13] Shaked, Avner and John Sutton. Relaxing Price Competition Through Product Differentiation. Review of Economic Studies, 49(1): 3-13, 1982.

[14] Siegel, Ron. All-Pay Contests. Ph.D. Thesis, Stanford Graduate School of Business, 2007.

[15] Spiegler, Ran. The Market for Quacks. Working Paper No.21-2003, The Foerder Institute for Economic Research and The Sackler Institute of Economic Studies, 2003.

[16] Spiegler, Ran. The Market for Quacks. Review of Economic Studies, 73(4): 1113-1131, 2006.

[17] Spiegler, Ran. Competition over Agents with Boundedly Rational Expectations. Theoretical Economics, 1(2): 207-231, 2006.

[18] Szech, Nora. Are too many healers spoiling the cure? Working Paper, Bonn Graduate School of Economics, 2008.

[19] Varian, Hal. A Model of Sales. American Economic Review, 70(4): 651-659, 1980.