Welfare in Markets where Consumers Rely on Word of Mouth*

Nora Szech

April 2011

Abstract

We analyze a market with rational firms knowing the distributions from which their opponents' qualities are drawn. Firms engage in price competition. Following Spiegler (2006a) we assume that consumers only see the firms' prices and rely on word of mouth to judge the firms' different qualities.
We prove equilibrium uniqueness for the special case of complete information on the firms' side. With this result, we characterize all equilibria of the incomplete information model. Different equilibria generate identical payoffs for the firms, but different welfare results. In the monotone pricing equilibrium, welfare converges to zero in the number of firms.

JEL Classification: D03, D44, L13

Keywords: Word of Mouth, Welfare, Anecdotal Reasoning, Bounded Rationality, Health Care

*I would like to thank Konstanze Albrecht, Mark Armstrong, Olivier Bos, Gary Charness, Drew Fudenberg, Nicola Gennaioli, Susanne Goldlücke, Paul Heidhues, Fabian Herweg, Sebastian Kranz, Benny Moldovanu, Patrick Schmitz, Anja Schöttner, Avner Shaked, Dan Silverman, and Rani Spiegler, as well as seminar and conference audiences in Bonn, Dortmund, Glasgow, Maastricht, and Mannheim. Nora Szech, University of Bonn, Lennéstr. 37, 53113 Bonn, Germany; email: nszech@uni-bonn.de.

1 Introduction

There are plenty of markets in which consumers are not fully aware of the different qualities of specialists and rely on word of mouth. Whereas consumers are unfamiliar with the market, specialists know the market well: They know how good they are, they have a good idea about their competitors' qualities, and they know how consumers search for them. Specialists act rationally.

Consumers who come into an unfamiliar market see the prices charged by the firms, but not the different firms' qualities. They rely on word of mouth to get a rough idea about the qualities. If several firms seem to offer a good quality, consumers focus on the prices as the ultimate selection device. Even though high quality firms get recommended more often, they may have to compete in prices against much lower quality firms if those get recommended as well. Anticipating this, firms play a very different pricing game than in a traditional market model with rational consumers. Competition is not necessarily beneficial in this market - as soon as some firms compete against each other, welfare may decrease substantially. This is the situation explored in this paper.

In the following, we mostly stick to markets for health care and health insurance as leading examples. The "healers" in our model can hence be seen as specialized therapists or as providers of health care insurance. Yet our analysis applies to any market with which consumers are not familiar, such that they rely on anecdotal evidence to judge the qualities of different firms, e.g. markets for car repair or markets for financial advice. In these markets, prices are salient and easy to grasp, but quality is not.

A large body of experimental research shows that anecdotes serve as a compelling and convenient tool for transporting information and influencing behavior.[1] In the medical literature, there is broad evidence that patients rely on anecdotal reasoning.[2] Even if statistical information on different forms of therapy is available (which often is not the case, e.g., for surgical treatments)[3], patients tend to prefer to rely on personal stories. Fagerlin et al. (2005) point out identification and emotional feelings as driving factors behind this. Patients may find it much easier to identify with a "natural person" than with the "statistically average person".[4] Additionally, in situations of uncertainty, people are often driven by emotions, and anecdotes transport more emotions than statistical results.[5]

Fagerlin et al. (2005) see another compelling characteristic about anecdotes in that "anecdotal information often provides a clear dichotomy − either an individual was cured or not" (p. 399). This kind of information may be much easier to grasp for a lay person than some statistical percentage of getting cured, and hence be much more easy to relate to. Indeed, most people have difficulties in understanding percentages and basic statistical concepts. For example, the importance of sample size is typically not recognized by untrained subjects. This has been shown in general studies as well as in medical contexts like cancer treatment or cancer screening.[6]

We assume that patients rely on word of mouth regarding quality in an otherwise standard market model: Patients only think about attending a healer if they heard some good story about him, and avoid those healers on which they heard something negative. When patients heard some favorable report about several healers, they opt

[1] See, e.g., Kahneman and Tversky (1973), Borgida and Nisbett (1977).

[2] Compare, e.g., Fagerlin et al. (2005) and the references therein and Enkin and Jadad (1998).

[3] Compare Gattellari et al. (2001) and McCulloch et al. (2002). McCulloch et al. (2002) state that "treatments in general surgery are half as likely to be based on RCT [Randomised Control Trials] evidence as treatments in internal medicine" (p. 1448).

[4] Compare also Jenni and Loewenstein (1997).

[5] Compare Loewenstein et al. (2001) and Finucane et al. (2000).

[6] Compare among others Tversky and Kahneman (1971), Hamill et al. (1980), Garfield and Ahlgren (1988), Yamagishi (1997), Schwartz et al. (1997), Weinstein (1999), Lipkus et al. (2001), Weinstein et al. (2004).

for the recommended healer with the lowest price. The way we model the patients' behavior is known as the $S(1)$ rule (where $S(1)$ stands for "sampling once") and goes back to Osborne and Rubinstein (1998). It has been applied as well by Spiegler (2006a, 2006b), Rubinstein and Spiegler (2008) and Szech (2010). Of these papers, Spiegler (2006a) is closest to the present study as we discuss below. Szech (2010) complements our welfare analysis by shedding light on the case where healers choose their qualities themselves, see the discussion in Section 5.3.

Besides $S(1)$, other related approaches for modeling boundedly rational consumer behavior are Ellison and Fudenberg's (1995) "word-of-mouth learning" and Rabin's (2002) "law of small numbers". Further related models from the literature on bounded rationality are reviewed in Spiegler (2006a). More broadly, our paper contributes to the literature on interactions between rational firms and boundedly rational consumers as surveyed by Ellison (2006).

We generalize the model of Spiegler (2006a). He analyzes a market of quacks who all have the same qualities and do not succeed better than some costless outside option the patients could choose instead. In our model, healers have true healing powers, but may differ strongly in their healing qualities. Additionally, we assume that the healers do not know the qualities of their competitors perfectly: Healers only know the distributions from which the qualities of their competitors are drawn. Spiegler shows different types of market failure, e.g. that patients' surplus may fall in the number of quacks for a low overall number of quacks in the market. Yet this negative effect of competition disappears if the number of quacks gets larger. Harsh competition among many quacks drives the prices down. As the quacks offer identical (low) qualities, patients fare better as competition becomes strong. In contrast, in our model, both, patients' surplus and overall welfare typically increase for low numbers of healers, but start to decrease substantially when too many healers are active. Both even go to zero if healers employ monotone price strategies. This negative effect of competition is in stark contrast to the predictions made by standard market models.

While our way of modeling anecdotal reasoning follows Spiegler, the logic behind our results is novel: Through the introduction of incomplete information we obtain pure equilibria which differ markedly from the mixed equilibria analyzed in Spiegler (2006a) and Szech (2010). Even a tiny amount of uncertainty in the quality realizations allows for pure price strategies where better healers charge higher prices. In this monotone equilibrium, patients who cannot properly distinguish between qualities are naturally driven to low quality healers: Patients pick the healer with the lowest price among all recommended healers. Thus they end up with the worst healer among the recommended ones if prices are monotone in quality.

Let us at this point turn briefly to the political debate about the performance of the US health insurance systems: In light of our results, it is not surprising that recent research revealed that the Veterans Health Administration (VHA) often offers better quality treatments than competitive health insurers in the US:[7] In contrast to most other medical insurers in the US, like Health Maintenance Organizations (HMOs), the VHA does not stand in competition to other insurers. The patients of the VHA typically stick with the institution for the rest of their lives. In contrast, the customers of most other US insurance systems typically switch their health plans on a regular basis. Hence competition plays a big role for most US health plans, but not for the VHA.

Our model gives an intuition for why competition can be detrimental to welfare in

[7]See Brooks (2008) and Longman (2010). Based on 294 indicators of quality, Asch et al. (2004) find that the VHA scores higher than all other sectors of American health care. Patients inside the VHA receive significantly better adjusted overall quality, better chronic disease treatment and preventive care.

the health insurance market: For patients, seeing prices of medical plans is easy. Comparing myriads of different care plans for various health problems is difficult. When choosing their insurance, patients may therefore screen the different medical plans only with regard to coverage of a small sample of conditions. Alternatively, they may rely on recommendations by other insured and their limited experiences. Our analysis shows that in such a market, the choice among many insurers may lead to patients ending up with medical plans of low quality.

That in the complex insurance market consumers may indeed focus on prices too much is underpinned by the recent decision in Germany to legally cut back the price competition among social health insurers to a minimum. The idea behind was to force people to put their attention away from price differences to quality differences.[8]

Our model also contributes to explaining why many therapies that lack evidence of therapeutical advantage compared to simpler therapies[9] or placebos[10] survive in the health market. Even if a doctor has the best intentions, if his therapeutical method does not help patients in the best possible way, patients should better go elsewhere. To maximize overall well-being (welfare), only the best therapies should survive in the market. Our model shows that even strong competition over patients who rely on word of mouth does not drive out therapists of poor quality.

From a theoretical point of view, Ireland (1993) and McAfee (1994) analyze a closely related game with a different interpretation, namely advertising, in mind: Competi-

[8] Compare the following statement by the German Federal Ministry of Health (2009), translated: "The uniform insurance fee ends the unfair competition for the cheapest fee. Instead it opens a fair competition for the best service and additional benefits to the insured." Clearly, such "fairness" considerations would be pointless with perfectly rational patients.

[9] For example, arthroscopic surgery, one of the most often performed surgeries with the aim of lowering pain in arthritic knees, was only recently questioned by Kirkley et al. (2008), who doubt the efficacy of this therapy. Kallmes et al. (2009) find that vertebroplasty, a commonly performed spinal surgery to treat osteoporotic compression fractures, leads to no improvements in pain and pain-related disability.

[10] Compare Fontanarosa et al. (1998).

tion over consumers is modeled as in our study, yet firms know each others' qualities (or rather advertising intensities in their specification) for sure. We add to the analysis of these two papers the equilibrium uniqueness in the pricing stage. The question of equilibrium uniqueness had been pointed out by both authors as an open problem. The uniqueness result stands in an interesting contrast to the multiplicity of equilibria in related models of price dispersion such as Varian's model of sales (1980) or the complete information all-pay auction.[11] Equilibrium uniqueness in the complete information case is a crucial step for characterizing all equilibria of our general game with incomplete information.

Finally, let us point out that the natural definition of welfare is fundamentally different for the market models studied by McAfee and Ireland, as there are no differences in the firms' service-qualities, but only in the firms' advertising activities. Thus, in these models, welfare increases in the number of firms.

The paper is structured as follows: Section 2 presents the model and describes in detail the behavioral $S(1)$ rule our patients follow. In Section 3, we characterize all equilibria of the model where each healer's quality may be drawn from a different distribution function. As a by-product, we show equilibrium uniqueness for the pricing game of Ireland (1993) and McAfee (1994). In Section 4, we assume that the healers' qualities are independently drawn from the same distribution F. We characterize the equilibrium in monotone price setting strategies. We show that as the number of healers gets large, overall welfare goes to zero. As an example, we assume qualities to be uniformly distributed: Welfare starts to decrease (and decreases substantially) in the number of healers as soon as there are more than three healers in the market. Section 5 discusses the robustness of our results. Among others, we discuss in Section 5.1 our choice of equilibrium selection by comparing our

[11] See Baye, Kovenock and de Vries (1992, 1996).

model to a number of related models which make different assumptions on patients' behavior. In Section 5.4 we present some survey evidence of $S(1)$ behavior. Section 6 concludes. Appendix A contains proofs and Appendix B contains the questionnaire for the results presented in Section 5.4.

2 The Model

We consider a market with n rational healers and a continuum of mass 1 of boundedly rational patients. The quality α_i of healer i is drawn from a distribution function F_i. The F_i are commonly known by all healers, but not by the patients. The supports of all F_i are assumed to be subsets of $[0, 1]$. Furthermore, we assume that the expected quality $\overline{\alpha_i}$ of each healer i satisfies $0 < \overline{\alpha_i} < 1$. Without loss of generality, healers are sorted by expected qualities, i.e., $\overline{\alpha}_i \leq \overline{\alpha}_j$ for $i < j$. Let $E[\cdot]$ denote the expectation with respect to the α_i. Initially, the patients are ill. They have a utility of 0 from staying ill, and a utility of 1 from getting cured. A healer with quality α cures each of his patients with probability α independently of the other healers.

The **timing** is as follows:

1. Each healer learns his personal quality realization α_i. This information is private.

2. The healers set their prices P_i simultaneously.

3. The patients decide whether to attend a healer and if so, which one.

4. Patients who consult healer i get cured with probability α_i.

In Step 3, patients decide according to the behavioral rule **S(1)** as introduced by Osborne and Rubinstein (1998), and as utilized in Spiegler (2006a). This rule works as follows:

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

The last point implicitly contains a tie-breaking rule: If a patient has to choose between consulting a recommended healer at a price of one or staying at home, the patient opts for the healer. It can be shown that no equilibrium exists if we depart from this assumption. All other ties can be broken arbitrarily.

Note that patients rely far too much on the signal they get - they over-infer from their sample. The idea behind the $S(1)$ rule is to capture reliance on anecdotal evidence in a simple way: Each patient independently asks some "former" client of each healer.[12] A client of healer i got cured with probability α_i. Thus, with probability α_i, he recommends healer i to the patient. The patient perfectly trusts this report - he either thinks the healer can cure him for sure or not at all.

Note that if a patient consults healer i his utility is $1 - P_i$ with probability α_i and $-P_i$ otherwise. The $S(1)$ rule is supposed to capture the idea that patients are not familiar with how the market works in detail. In particular patients are not aware of the healers' qualities α_i: Patients act as if some healers were always successful and others never.

[12]Of course, this dynamic motivation is only for intuition, as we are in a static model here.

Finally, an alternative interpretation of the $S(1)$ rule is as follows: Assume the healer is a health insurer and the quality of the insurer is given by the proportion of medical problems covered by his insurance plan. Patients just sample each insurer with regard to one random medical condition. Hence with probability α_i they receive the positive signal that the condition considered is covered by the insurance plan. They then think this plan is a good one, in contrast to the plans on which they received a signal of non-coverage. The medical condition a patient faces in the future is independently drawn by nature.

3 Characterization of all Equilibria

In this section, we determine all equilibria of the model. For the analysis, it is helpful to consider the model with deterministic qualities $\overline{\alpha}_1, ..., \overline{\alpha}_n$ as well. This is the special case of our model where the healers hold complete information about each others' qualities. This model has been analyzed by Ireland (1993) and McAfee (1994) in the context of advertising. We add the uniqueness to their characterization of equilibrium. This result is a crucial step towards the characterization of all equilibria of the incomplete information game.

To put the equilibria we find into perspective, note that if the healers know each others' qualities perfectly, and if qualities are strictly between 0 and 1, standard arguments yield that there cannot be an equilibrium in pure strategies: Each healer has the possibility to earn a positive expected payoff as with some probability he is the only recommended healer in the market.[13] Hence each healer chooses a price strictly higher than zero. Thus pure pricing strategies cannot constitute an equilibrium, as there would always be a healer who would like to attract more patients

[13]Recall that patients never attend healers that are not recommended, as they expect a negative utility of $0 - P_i$ from attending them.

by deviating to a slightly lower price. This is why in the game with complete information about qualities, the equilibrium must be in mixed strategies. The unique mixed equilibrium of this game is given by Proposition 1.

Proposition 1 *Define a sequence of prices $p_0, \ldots, p_n$ by*

$$p_i = \frac{(1-\overline{\alpha}_{i+1}) \cdot \ldots \cdot (1-\overline{\alpha}_{n-1})}{(1-\overline{\alpha}_i)^{n-i-1}}$$

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

$$p_0 = \prod_{i=1}^{n-1}(1-\overline{\alpha}_i)$$

and $p_{n-1} = p_n = 1$. Then the unique Nash equilibrium of the complete information game with qualities $\overline{\alpha}_1, ..., \overline{\alpha}_n$ is the following: Each healer i mixes over the interval $[p_0, p_i]$ using the distribution function H_i defined as

$$H_i(p) = \frac{1}{\overline{\alpha}_i}\left(1 - \sqrt[n-j]{\frac{(1-\overline{\alpha}_j) \cdot \ldots \cdot (1-\overline{\alpha}_{n-1})}{p}}\right) \tag{1}$$

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 $H_i = 0$ and, on $[p_i, 1]$, $H_i = 1$. H_n places an atom of size $1 - \frac{\overline{\alpha}_{n-1}}{\overline{\alpha}_n}$ on 1.

The question of uniqueness of equilibrium in the complete information game with qualities $\overline{\alpha}_1, ..., \overline{\alpha}_n$ had been pointed out as an open problem by Ireland (1993) and McAfee (1994).[14] Spiegler (2006a) proves uniqueness of equilibrium for the special cases where all healers offer the same quality and where all but one healers offer the same low quality and one healer a higher quality.

In the complete information game with qualities $\overline{\alpha}_1, ..., \overline{\alpha}_n$, the payoff of healer i from playing some price p while the other healers use the mixed strategies H_j is

[14] It is straightforward to generalize our uniqueness result to the more general demand functions considered in McAfee (1994).

given by

$$\pi_i(p) = p\overline{\alpha}_i \prod_{j \neq i} (1 - \overline{\alpha}_j H_j(p)). \tag{2}$$

The intuition is as follows: In order to attract a patient and earn p, healer i has to be recommended (which happens with probability $\overline{\alpha}_i$) and has to be the cheapest healer among those who are recommended. (The probability that a competitor j is not both recommended and cheaper than p is $1 - \overline{\alpha}_j H_j(p)$.) We insert the explicit formulas for the distribution functions H_j from Proposition 1 into (2). Then we can calculate that the expected equilibrium payoff of healer i is given by

$$\pi_i(p) = \overline{\alpha}_i \prod_{j \neq n} (1 - \overline{\alpha}_j) \text{ for all } p \in [p_0, p_i].$$

From this we can deduce that the distribution functions H_i also form a Nash equilibrium in the incomplete information game:

Proposition 2 *The distribution functions H_i defined in (1) form a Nash equilibrium in the incomplete information game with qualities $\alpha_1, ..., \alpha_n$. In this equilibrium, the payoff of healer i is given by*

$$\pi_i = \alpha_i \prod_{j \neq n} (1 - \overline{\alpha}_j).$$

Note that the equilibrium strategies of the healers do not depend on the realizations of their qualities α_i. Thus the healers do not make use of their private information. The intuition is as follows: Once he got recommended to a patient, it does not matter anymore for a healer how good or bad his quality actually is. For his competitors, the exact quality of the healer plays no role either, as they do not know it: The competitors can base their strategies only on the healer's expected quality. Yet most healers (that is to say all healers $i \neq n$ if $\overline{\alpha}_n > \overline{\alpha}_{n-1}$) do incorporate their expected quality $\overline{\alpha}_i$ into their equilibrium strategies.

The next proposition establishes that the expected healers' payoffs are the same in all equilibria of the incomplete information game. Furthermore, all the equilibria are interchangeable[15]: Any two equilibria A and A' can be combined to form another equilibrium A'' by assigning to the healers their respective strategies from A or A' in an arbitrary way. This is the extent to which the equilibrium uniqueness from the complete information game with qualities $\overline{\alpha}_1, ..., \overline{\alpha}_n$ carries over.

Proposition 3 *A profile of strategies $((G_1^{\alpha_1})_{\alpha_1}, \ldots, (G_n^{\alpha_n})_{\alpha_n})$ is a Nash equilibrium if and only if:*

$$E[\alpha_i G_i^{\alpha_i}(p)] = \overline{\alpha}_i H_i(p) \text{ for all } i \tag{3}$$

and all $G_i^{\alpha_i}$ have their support in $[p_0, p_i]$.

Furthermore, the healers expect the same payoffs in all equilibria.

Obviously, there are infinitely many equilibria since each healer i can make his strategy dependent on α_i in an arbitrary way as long as (3) is satisfied. Notably, if the noise in the α_i is rich enough, pure price strategies are possible in equilibrium.

Proposition 4 *Assume that the distribution functions F_i have continuous and strictly positive densities f_i on $[0,1]$. Furthermore, assume that $\overline{\alpha}_n = \overline{\alpha}_{n-1}$. Then there is a unique pure strategy equilibrium with strictly increasing price setting functions $\bar{P}_i(\alpha_i)$. Moreover,*

$$\bar{P}_i(\alpha_i) = H_i^{-1}\left(\frac{\int_0^{\alpha_i} \beta f_i(\beta) d\beta}{\overline{\alpha}_i}\right).$$

The randomness in the quality realizations allows the healers to remain unpredictable competitors even if they choose a pure pricing strategy. Note that Proposition 4 does not demand for much noise in the sense of a large variance. The main

[15]See Osborne and Rubinstein (1994), p. 23.

requirement in Proposition 4 is that the qualities are drawn from atomless distributions. (An atom would force the healers to mix over prices in equilibrium.) For the sake of brevity, we exclude the case of $\overline{\alpha}_n > \overline{\alpha}_{n-1}$. In that case all our arguments still go through but, because of the atom in H_n, $\bar{P}_n$ reaches the value 1 already for some $\alpha_n < 1$ and then stays constant. Intuition clearly favors the monotone price equilibrium over the other equilibria: It is natural to assume that a healer with a higher quality charges a higher price. The same selection criterion among equilibria is applied in the standard literature on auctions.

4 Welfare

In this section we focus on the symmetric case where the qualities of all healers i are independently drawn from the same distribution $F_i = F$. We assume that F has a continuous and strictly positive density f. Denote by $\overline{\alpha}$ the mean of F. Hence $\overline{\alpha}$ is also the expected quality of a randomly drawn healer. We study welfare in the monotone strategy equilibrium and show that it deteriorates an n gets large. At the end, we compare our welfare results to those of the standard model where patients act rationally, but hold only incomplete information about the healers' qualities.

With the help of the results of the previous section we obtain the following monotone pricing equilibrium:

Proposition 5 *There is a unique equilibrium in monotonically increasing price strategies. In this equilibrium, each healer i uses the price setting function*

$$\bar{P}(\alpha_i) = \left(\frac{1 - \overline{\alpha}}{1 - \int_0^{\alpha_i} \beta f(\beta) d\beta} \right)^{n-1}.$$

The expected equilibrium payoff of healer i is given by

$$\pi_i = \alpha_i(1-\overline{\alpha})^{n-1}.$$

Conditional on his quality realization, each healer plays a pure pricing strategy. From the payoffs we see that each healer i earns in expectation only his expected maxmin payoff – the payoff he can earn for sure no matter what prices his competitors choose: With an expected probability of $(1-\overline{\alpha})^{n-1}$, all the healer i's competitors are not recommended. With a probability of α_i, he himself is recommended. Hence with an expected probability of $\alpha_i(1-\overline{\alpha})^{n-1}$, healer i is the only one who is recommended. Thus by charging a price of 1, healer i can secure an expected payoff of π_i to himself.

Denote by θ_n the healers' aggregate payoffs in a market with n healers:

$$\theta_n = \sum_i E[\pi_i] = n\overline{\alpha}(1-\overline{\alpha})^{n-1}. \tag{4}$$

As one easily sees from (4), the healers' aggregate payoffs may initially increase in n if $\overline{\alpha}$ is not too large but eventually converge to zero as n increases. The intuition is as follows: With few healers and low qualities, competition is soft. Only few patients are attracted by several healers, and many patients do not get recommended to any healer at all. A new healer entering the market may attract most of his patients from the group of patients that would otherwise stay at home. Thus the new healer does not strengthen competition much. Yet if more and more healers enter the market, even with low qualities, more and more patients get recommended to several healers. Then price competition gets more and more severe, driving the healers' payoffs down.

We have found that in expectation the healers' payoffs go to zero as n gets large. But does that mean that the patients are better off the more healers enter the market?

At least, in the limit, patients do not have to pay anything for the healers' services. Yet it turns out that also patients fare badly: Recall that a patient always consults the cheapest healer who is recommended to him because he thinks all recommended healers are of the same high quality and just differ in prices. Yet as the healers apply monotone price-setting strategies, by picking the cheapest the patient also picks the worst of all recommended healers. Since the distribution function of qualities F has support on the whole interval $[0, 1]$, as many healers enter the market there is - with high probability - also a considerable amount of very low quality healers, some of which get recommended. Making this reasoning precise, one can see that as n gets large, overall welfare converges to zero: No patient gets cured in the limit.

Proposition 6 *Denote by γ_n the expected social welfare, i.e. the expected proportion of patients cured, in the monotone equilibrium with n healers. Then*

$$\gamma_n = n \int_0^1 \alpha^2 \left(1 - \int_0^\alpha \beta f(\beta) d\beta\right)^{n-1} f(\alpha) d\alpha \tag{5}$$

and

$$\lim_{n\to\infty} \gamma_n = 0.$$

Like the healers' payoff θ_n, welfare γ_n is an expected value taken over the qualities α_i. The intuition behind formula (5) is the following: Consider the expected quality of the treatment chosen by a patient: A quality α is chosen if a) the healer offering that quality is recommended and b) all his competitors are either not recommended or are charging a higher price. The probability of a) is α. The probability of b) is $\left(1 - \int_0^\alpha \beta f(\beta) d\beta\right)^{n-1}$ as charging a higher price is equivalent to offering a higher quality in the monotone equilibrium.

Welfare may already decrease for a quite low number of healers. This is relevant as it seems much more natural to think of a patient receiving an anecdote on each healer if the market is not too large. However, as pointed out in Spiegler (2006a), n can also be interpreted as the number of healers a patient gets a report on in a market with very many healers. Then n would be a measure of patients' awareness.[16] The following example demonstrates that the effects at work are not only limit results but already play an important role already for moderate numbers of healers. Figure 1 depicts welfare for the special case that qualities are uniformly distributed on $[0, 1]$, i.e. $F(\alpha) = \alpha$:

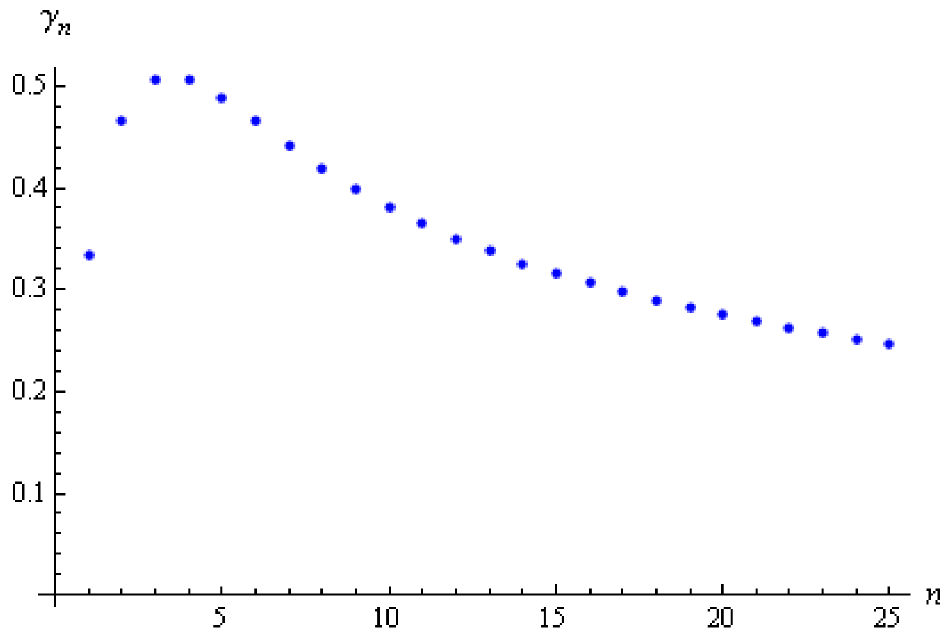

Figure 1: Expected social welfare γ_n for $\alpha_i \sim U[0, 1]$

Figure 1 shows that welfare is maximized for $n = 3$ healers. With three or four healers, the average quality a patient receives is slightly larger than the quality of the average healer $\overline{\alpha} = 1/2$, as better healers are more often recommended. Afterwards, less and less patients get cured, as most patients get recommended to

[16] We discuss this further in Section 5.

several healers and then, led by price-comparison, end up with a low quality healer.

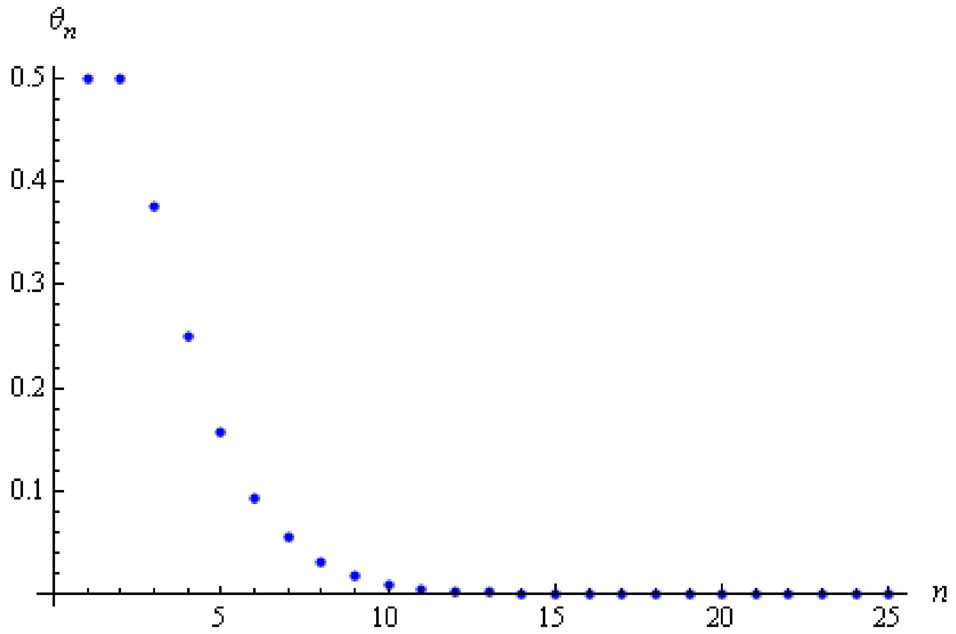

Figure 2: Expected healers' aggregate surplus θ_n for $E[\alpha_i] = \frac{1}{2}$

Figure 2 depicts the sum of the n healers' expected payoffs θ_n. θ_n is maximized with one or two healers where it equals 1/2. Afterwards θ_n decreases quickly.

Figure 3 depicts patients' aggregate surplus which is the difference between overall welfare γ_n and healers' surplus θ_n. The healers' surplus θ_n decreases considerably faster than the proportion of patients cured γ_n. Hence the patients' surplus is largest at an intermediate market size of $n = 8$. As n gets larger, the patients' surplus decreases, driven by the decreasing average quality received γ_n. Note that with one or two healers the patients' surplus is negative: In monopoly, the healer attracts all patients to whom he is recommended. He then charges a price of 1 for a treatment of expected quality $\overline{\alpha} = \frac{1}{2}$. In duopoly, patients are more often recommended to the healer with the higher quality (who offers in expectation a healing probability of $\frac{2}{3}$).

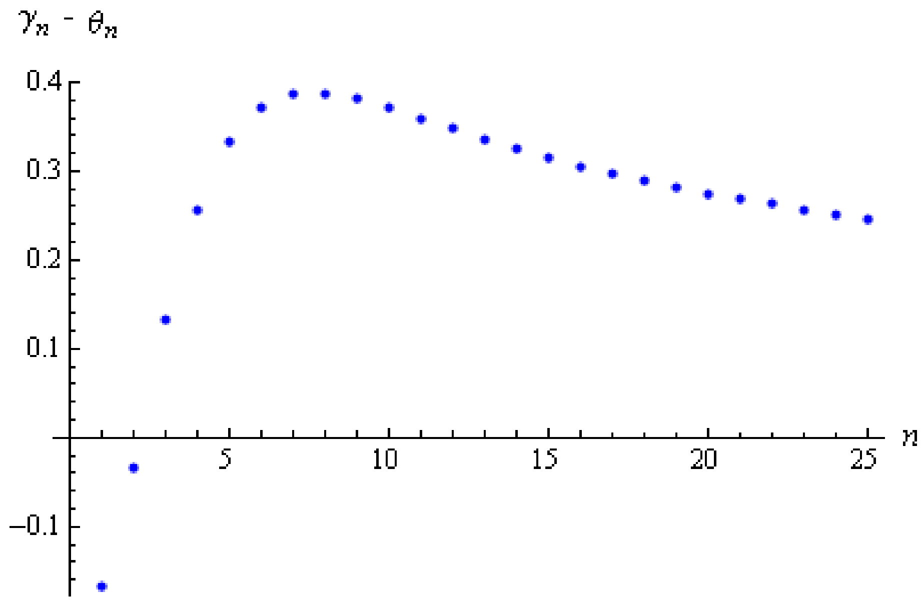

Figure 3: Expected patients' aggregate surplus $\gamma_n - \theta_n$ for $E[\alpha_i] = \frac{1}{2}$

Yet, as competition is weak, prices are still quite high. Patients' surplus increases, but remains negative.

5 Discussion

In this section we discuss several directions for extensions and show robustness of our results.

5.1 Equilibrium Selection

Possibly the aspect of our analysis that has the strongest need for discussion is our focus on the monotone strategy equilibrium in the pricing stage. From Proposition 3 it is clear that for the healers' expected payoffs it does not matter which equilibrium is played. The patients' health, however, varies across equilibria. In the following we give an informal discussion of how our analysis would change in three related models: We consider a model where patients receive several anecdotes, a model where some patients are informed about healers' qualities and a model where patients are rational and hold incomplete information. We will see that the first two models support the focus on the monotone strategy equilibrium. The third model gives support to the equilibrium in which healers mix with the same pricing strategy regardless of their quality.

5.1.1 Patients Receiving Several Anecdotes

A natural generalization of our model would be to assume that patients gather several anecdotes on each healer. While the equilibria of such generalized models usually preclude an explicit solution and would typically require additional assumptions to ensure existence of equilibrium, it is mostly straightforward to see that the main arguments of our analysis carry over: As long as each healer has a certain chance of being perceived as the very best healer by some patients, healers will be unwilling to engage in harsh price competition.

Assume for concreteness that each patient gathers k anecdotes about each of the healers. If a patient receives $l \leq k$ positive reports about a healer he expects to receive a utility of $\frac{l}{k}$ from attending him. This is the so-called $S(k)$-reasoning.[17] As has been pointed out by Spiegler (2003) this generalized model typically does not possess an equilibrium: Healers face a strong discontinuity in their payoffs at the prices $\frac{l}{k}$.[18] Spiegler shows that this existence problem can be solved by assuming

[17]Compare Osborne and Rubinstein (1998).

[18]Spiegler (2003) considers qualities which are known among healers but this does not make a difference with respect to this discontinuity.

that each patient has a willingness to pay v which is a continuously distributed random variable.

We want to demonstrate *heuristically* why we expect the monotone strategy equilibrium to be the unique price equilibrium in the $S(k)$-case. Assume for simplicity that $k = 2$ and fix the strategies of healers $j \neq i$. Healer i's price-setting problem can then be written as the problem of maximizing

$$\pi_i(p) = \alpha_i^2 \pi_{i,2}(p) + 2\alpha_i(1 - \alpha_i)\pi_{i,1}(p)$$

where $\pi_{i,l}(p)$ denotes healer i's expected payoff from patients who have received l positive reports on him (given the remaining healers' strategies).

We want to show that in this case we cannot expect the existence of a mixed equilibrium: Assume a mixed equilibrium existed. Then for each value of α_i we would need that

$$\pi_i'(p) = \alpha_i^2 \pi_{i,2}'(p) + 2\alpha_i(1 - \alpha_i)\pi_{i,1}'(p) = 0.$$

for all p. We cannot expect each summand to be zero separately: The contribution from $\pi_{i,1}(p)$, i.e. from patients who have a rather mixed opinion about healer i, can be sizable for small prices but for high prices it should be low. Conversely, patients with a high opinion still contribute to healer i's payoff at higher prices. Thus $\pi_i'(p) = 0$ can only be ensured if the two contributions cancel out each other. The first order condition is equivalent to

$$\frac{\pi_{i,1}'(p)}{\pi_{i,2}'(p)} = -\frac{2(1 - \alpha_i)}{\alpha_i}$$

where the right hand side is increasing in α_i. This can only be fulfilled for one value of α_i at a fixed p. Hence equilibrium should be pure. Moreover, since the contribution from high valuation patients is most significant for high quality healers, we expect an equilibrium in strictly increasing price strategies. Clearly, this reasoning does not depend on $k = 2$ and thus we expect our welfare results to hold for $k > 2$ as well. This is especially surprising since an $S(k)$-patient with a large k is in a sense not much different from a rational and completely-informed patient.

Finally, note that a higher k does not necessarily imply a better model of patients' actual behavior. With one anecdote per healer, patients face a clear dichotomy: Healers get subdivided into good and bad. This simplicity has been pointed out by Fagerlin et al. (2005) as one of the factors that make anecdotes so attractive to rely on, especially for patients in the health care market.

5.1.2 Well-Informed Patients

Next we consider a model where a fraction of "well-informed" patients perfectly sees the qualities α_i of n healers while the remaining patients use $S(k)$-reasoning. To give a foundation to such a model, one could assume that there are very many healers and that the well-informed patients decide between those n healers whose qualities they know.[19] Healers might also know all qualities but they do not know to which subset of opponents they are compared by a patient. Thus they treat opponents' qualities as independent draws from a common prior.

Consider first the situation where all patients are well-informed. Then the healers essentially play a standard incomplete information first-price auction in which they offer a surplus of $\alpha_i - p_i$ to the patients.[20] Under some technical assumptions on F it can be ensured that healer i's equilibrium price p_i is increasing in α_i. Thus we

[19] Compare also Section 5.2.

[20] See e.g. Krishna (2002). Clearly, bidding in an incomplete information first-price auction does not change when the "seller" (the patients) can observe the "bidders' " (the healers') "valuations" (qualities) as is the case here.

obtain a monotone strategy equilibrium. Consequently, we expect to still obtain a monotone strategy equilibrium if we combine well-informed and boundedly rational patients. Thus while well-informed patients in the market may exert a positive externality on $S(k)$-patients by lowering prices, they cannot keep them away from bad healers.

5.1.3 Incompletely-Informed patients

As a third variant of our model, consider patients who (just like the healers) believe that healers' qualities are independent draws from a common prior F. Patients receive anecdotes about healers and apply Bayesian reasoning to update their beliefs. Accordingly, if healer i has quality α_i, a fraction α_i of patients believes he has quality α^h while the remaining patients believe he has quality α^l for suitable values $0 < \alpha^l < \alpha^h < 1$. Note that – provided that prices do not contain any information about qualities – this is essentially a rescaled version of the $S(2)$-model of Section 5.1.1. Clearly, in this model, no separating equilibrium can exist: For example, under a monotone price strategy, patients would in equilibrium infer qualities from prices. Yet then a healer would want to deviate to higher prices.

To get an impression of the welfare implications of such a model, let us calculate the expected welfare $\widetilde{\gamma}_n$ for the case where all healers apply the same pricing strategy and where patients always attend a recommended healer. Clearly, this is an upper bound on the welfare achievable in a pooling equilibrium.[21] As n increases, $\widetilde{\gamma}_n$ converges quickly to some value strictly above $E[\alpha_i] = \overline{\alpha}$. The intuition is as follows: If all healers play the same price strategy, they all face the same probability that they are the cheapest and thus get attended by the patients - given that they are

[21]Note that $\widetilde{\gamma}_n$ is also the welfare achieved in the $S(1)$ equilibrium of Proposition 2 where (for $F_i = F$) all healers play the same mixed strategy H (independent of their quality realization). Furthermore, note that $\widetilde{\gamma}_n$ also describes the proportion of patients cured in a situation where patients apply anecdotal reasoning but where a fixed price is exogenously prescribed to the healers.

recommended. But since better healers are recommended more often a patient can expect an above average treatment when there are sufficiently many healers in the market. The following remark (proved in the Appendix) provides an explicit expression for $\widetilde{\gamma}_n$:

Remark 1 *If patients randomly attend one of the recommended healers, welfare is given by*

$$\widetilde{\gamma}_n = \frac{E[\alpha^2]}{E[\alpha]}(1 - (1 - E[\alpha])^n) \tag{6}$$

where α has distribution F.

By Jensen's inequality the limit for n to infinity, $E[\alpha^2]/E[\alpha]$, is strictly greater than $E[\alpha] = \overline{\alpha}$ (unless F is deterministic). The second factor in (6), $1 - (1 - E[\alpha])^n$, is the probability that a patient gets at least one recommendation and turns to the market for healers. The first factor of (6), $\frac{E[\alpha^2]}{E[\alpha]}$, is the expected quality a patient receives given he does not stay at home. This factor does not depend on the number of healers n as prices do not reveal anything about healers' qualities. Yet it depends on the variance of F:

$$\frac{E[\alpha^2]}{E[\alpha]} = \frac{Var[\alpha]}{E[\alpha]} + E[\alpha].$$

A higher variance is beneficial as it increases the average quality of a recommended set of healers.

While some basic level of competitions helps to increase welfare substantially, we see that also in this model, stronger competition does not lead to any further significant welfare improvements: $\widetilde{\gamma}_n$ converges exponentially fast to a value which is bounded away from 1 even though more and better healers enter the market.

We thus see that the predictions of the incomplete information model stand in marked contrast to those from the $S(k)$-models. While this can be seen as casting doubt on the validity of the $S(k)$-approach, it should not be overlooked that the opposite implication also holds: As we summarized in the introduction, the recent literature on how patients argue about different therapies shows strong evidence in favor of anecdotal reasoning. There is furthermore little reason to expect that such anecdotal thinking is just a peculiarity of consumers in the health care market. In Section 5.4 we present some survey-data in support of $S(1)$ reasoning: Participants got very influenced by the recommendations when choosing among services of markedly different price, even though they were told these recommendations had (almost) no predictive value.

Since our model leads to considerably different welfare implications than a standard incomplete-information model, it seems problematic to rely only on the latter type of model e.g. for policy recommendations. While trying to connect the two approaches is a challenging field for future research, we believe that our model presents a very interesting (second) benchmark case. See also the discussion in Spiegler (2010).

5.2 Awareness of Patients

Throughout the analysis, we assumed that n is the number of healers in the market. Yet as we outlined before[22], we can reinterpret n as the number of healers a patient samples, hence as a measure of patients' awareness. Especially in a large market, patients might only sample a fraction of all healers − and sampling intensity may be heterogeneous among the patients. We can incorporate this heterogeneity into our model by taking n as an integer-valued random variable. Our results prove to be robust to this extension: For simplicity, focus on the symmetric case $F_i = F$. Let us define

$$\rho = E[(1 - \overline{\alpha})^{n-1}],$$

[22] Compare also Spiegler (2006a).

where the expectation is taken over n. With the same argument as in Section 3, healer i's expected equilibrium payoff is then given by

$$\pi_i = \alpha_i \rho.$$

There exists again a symmetric mixed strategy equilibrium H which does not depend on the realizations of the α_i. The support of H is the interval $[\rho, 1]$ and H is given implicitly as the solution of

$$\rho = pE[(1 - \overline{\alpha}H(p))^{n-1}].$$

From this mixed strategy equilibrium it is straightforward to construct a monotone strategy equilibrium the same way as for deterministic n. Hence again, too much competition turns out to be detrimental to welfare.[23]

5.3 Endogenous Qualities

One assumption of our model which may seem rather strong is that the healers' qualities are exogenous random variables. It is plausible that while healers may not be able to fully control their qualities, they can influence them at least to some extent. With the application to advertising in mind, Ireland (1993) and McAfee (1994) analyze extended games where the healers choose their qualities themselves before pricing takes place. They focus on the case of complete information about qualities. Szech (2010) adds the welfare analysis under the $S(1)$ interpretation.[24]

[23] As an aside, from applying Jensen's inequality we see that $\rho > (1 - \overline{\alpha})^{E[n]}$. The healers' equilibrium payoff is thus higher than in the model with a deterministic number of $E[n]$ healers. In this sense, the heterogeneity in n is beneficial for the healers.

[24] Welfare results differ markedly from welfare in the advertising interpretation. Under the interpretation of advertising, firms differ only in advertising activities, but not in service qualities. α_i describes solely the probability with which a consumer gets aware of firm i. Hence any consumer

It is straightforward to combine the analysis of endogenous qualities for the complete information case with our incomplete information model: Recall that ex ante, a healer's expected payoff only depends on his expected quality (and not on any further properties of the distribution function). Thus the analysis of the complete information case transfers immediately to a model where healers choose between several distributions from which their quality is drawn (in the sense of, e.g., choosing between different specializations): When choosing between different distribution functions, healers only take into account the expected quality.

From the analysis of the complete information case it is known that healers typically choose much lower than socially optimal qualities, to make competition softer and hence raise revenues.[25] Also, if the best possible qualities are not too low, there will be much difference in the qualities offered by the best and the worst healers. Hence endogenous quality choice creates the situations of varying qualities to which our bad welfare results apply.

5.4 Survey Evidence of $S(1)$ Behavior

To gather some preliminary empirical evidence in favor of our $S(1)$ assumption, we conducted a short survey in a lab experiment asking 93 subjects[26] which service-provider they would choose among six different options in the following hypothetical setting: The service was a professional prophylactic tooth cleaning. Subjects were informed about the prices of the different providers (dentists). Additionally, they had information from a newly opened up internet platform providing exactly one

ending up at a firm receives the same gross utility.

[25]See Ireland (1993) and McAfee (1994) for the explicit form of the pure quality-setting equilibria, respectively, with and without costs of quality-setting. Szech (2010) adds a more detailed analysis of mixed strategy equilibria and shows that welfare decreases for larger numbers of healers under the interpretation of anecdotal reasoning.

[26]Subjects aged between 19 and 44 and also participated in an unrelated other experiment when invited to the lab. They mostly consisted of students at the University of Bonn.

Table 1: Choice of Services

Dentist	Thumbs	Price	# Subjects
A	Down	120	0
B	Up	80	51
C	Down	100	0
D	Up	120	0
E	Down	60	5
F	Up	100	15
G	Down	80	0

rating (thumbs up or thumbs down) from respectively one former patient on each provider. Clearly, such a signal carries basically no or very little information, as was also pointed out by many subjects in an answer sheet (where subjects could explain their choice of provider). Service prices ranged from 60 to 120 Euro, where the cheapest provider was without recommendation as seen in Table 1.[27]

We find that 74% of subjects choose Provider B and thus among the recommended providers the provider with the lowest price. We take this as a hint that indeed $S(1)$ reasoning seems to capture a way people make selections in markets of little quality information. The second option that was chosen by a sizable group of patients was the second-cheapest recommended option while only few chose the cheapest one.[28]

6 Conclusion

"It is unwise to pay too much, but it's worse to pay too little. When you pay too little, you sometimes lose everything because the thing you bought was incapable of doing the thing you bought it to do." This recommendation is attributed to the social thinker John Ruskin (1819-1900). Indeed, in markets where qualities are not easy to grasp, competition among firms may lead to consumers ending up with poor

[27]To control for order effects, part of the subjects received the options in a different order, see Appendix B.

[28]A translation of the German instructions and more details are found in Appendix B.

qualities, as they focus too much on price differences. Our model shows that even if patients try to get an idea about the qualities in a market, they likely end up with a bad quality.

The assumption that consumers rely on word of mouth captures empirical findings from the psychological and economic literature. Recently, also medical research puts a lot of attention to the phenomenon that lay people tend to prefer to rely on anecdotes even if statistical evidence is available and presented in an appealing, easy-to-grasp way. This fact has even led to recommendations of incorporating personal stories into evidence-based results, such that patients may be more willing to adhere to statistical recommendations.[29]

Assuming that consumers apply anecdotal reasoning, our model generates very different predictions than those made by standard market models. Stronger competition turns out to be detrimental to welfare. Recent surprising results from the US medical system support this conclusion, showing that the non-competing Veterans Health Administration often provides higher quality services than the competitive health systems prevalent in the US.

Generally, we believe that more research is needed to explore the interplay between perfectly rational firms and boundedly rational consumers following behavioral, possibly market-specific rules instead of perfectly rational thinking.

[29] Compare e.g. Glenton et al. (2006).

A Proofs

Proof of Proposition 1

In McAfee (1994) it is shown that the vector of strategies in Proposition 1 is indeed an equilibrium. McAfee (1994) also shows *payoff uniqueness*, i.e., for all i, healer i's equilibrium payoff is given by

$$\pi_i = \alpha_i C \text{ where } C := \prod_{j=1}^{n-1}(1-\alpha_i) > 0. \tag{7}$$

Observe that the first equality states that all healers' expected equilibrium payoffs *conditional on being recommended* must be identical.

We hence take these results as given and show how to obtain equilibrium uniqueness from this point on. We start with a number of preliminary observations:

- *In any equilibrium healers do not place atoms on prices except for possible atoms on* 0 *or* 1*:* If healer i sets an atom in $p \in (0,1)$, other healers playing prices right above p would want to shift their probability mass to prices marginally below p in order to substantially increase their winning probability (while only marginally decreasing prices). If no other healer played prices right above p, healer i could profitably shift the atom upwards.

- *In addition, at most one healer sets an atom on* 1 *in equilibrium:* If two or more healers played an atom in 1, this would result in a positive probability of ties. Thus at least one healer could profitably deviate by shifting his atom marginally downwards.

- *The union of the healers' strategy supports must go up to* 1*:* Playing higher prices is dominated. If the union of supports went only up to a lower price $p_H < 1$, any healer mixing up to p_H could profitably deviate to playing 1.

- Due to the positive equilibrium payoffs, the union of strategy supports must be bounded away from 0. Denote by $p_L > 0$ the infimum of the union of equilibrium supports.

- *The union of supports must be an interval* $[p_L, 1]$*, i.e. there cannot be any gaps in the union of supports:* If there was an interval $[\underline{p}, \overline{p}] \subset [p_L, 1]$ where no healer was active, a healer who would be playing prices right below $\underline{p}$ could deviate by shifting probability mass from a small interval below $\underline{p}$ to $\overline{p}$, yielding a substantially better price at a marginally lower probability of winning.

- *Furthermore, there cannot be a subset* $[\underline{p}, \overline{p}] \subset [p_L, 1]$ *where only one healer is active:* Such a healer could profitably deviate by concentrating all probability mass of the interval in an atom at $\overline{p}$. He would then receive a higher price at the same probability of winning.

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

1) *In any equilibrium, the strategy support of each healer must go down to the same* $p_L > 0$*. Furthermore, in any equilibrium,* $p_L = C$*.*

Proof of 1): Consider two healers 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 $k = i, j$. Then, with positive probability, healer i plays a price from $[p_L^i, p_L^j]$. Healer j's payoff from playing p_L^j must equal his equilibrium payoff $\alpha_j C > 0$. Yet this implies that healer i can earn more than his equilibrium payoff of $\alpha_i C$ by playing p_L^j: Since - unlike healer j - healer i does not compete against healer i (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 a contradiction to (7). Hence the support of every healer must go down to the same lowest price p_L.

To see that $p_L = C$, note that for all j, healer j's payoff from playing p_L must be $\alpha_j p_L$: The other healers charge higher prices with probability 1. Thus healer j attracts all the patients to whom he is recommended and receives p_L from all of them. This leads to a payoff of $\alpha_j p_L$, which is only consistent with (7) if $p_L = C$.

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

2) *Let $\mathcal{D} \subset \{1, \ldots, n\}$ denote the set of healers who are active on some interval $I = (\underline{p}, \overline{p})$ in some arbitrary but fixed equilibrium. Assume all healers $j \in \mathcal{D}$ are active at any $p \in I$ and let $m = \#\mathcal{D}$. (Note that from our preliminary observations it follows that $m \neq 1$.) Then for all $j \in \mathcal{D}$ any equilibrium distribution function $H_j(p)$ must satisfy for all $p \in I$*

$$H_j(p) = \frac{1}{\alpha_j}\left(1 - \sqrt[m-1]{\frac{L}{p}}\right) \tag{8}$$

where the constant $L > 0$ is independent of p and j. Moreover,

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

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

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

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

$$\prod_{k \in \mathcal{D} \setminus \{j\}} (1 - \alpha_k H_k(p)) = \frac{L}{p}. \tag{9}$$

Now consider (9) for two different healers $i, j \in \mathcal{D}$. Taking the quotient of (9) for i and (9) for j yields that for all $p \in I$

$$1 = \frac{1 - \alpha_j H_j(p)}{1 - \alpha_i H_i(p)}$$

which implies that there is a function $h(p)$ such that $h(p) = \alpha_k H_k(p)$ for all $k \in \mathcal{D}$. Substituting $h(p)$ for $\alpha_k H_k(p)$ on the left hand side of (9) and then solving for h yields

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

and thus

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

as required.

The last main step shows that no healer has a gap inside his equilibrium price interval, i.e. no healer is inactive over some range of prices (above p_L) while putting positive probability mass on prices above that range:

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

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

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

Defining the constant factor from the other inactive healers as

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

and making use of (8) from the last step, we can express $\pi_j(p)$ as

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

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

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

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

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

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

To conclude the proof, we still have to show that the vector of strategies defined in the proposition is actually the only candidate for an equilibrium. We have seen that all supports start at $p_L = C$ and since healers do not set atoms or leave gaps in their supports, all healers remain active up to the price p_1 where the first healer(s) j have used up their probability mass, i.e. where $H_j(p_1) = 1$. Note that on any interval $[p_L, \overline{p}]$ where all healers are active, all distribution functions are uniquely determined by Step 2. Likewise, p_1 and the set of healers with $H_j(p_1) = 1$ are uniquely pinned down by this. Above p_1, all healers who still have probability mass to spend must remain active. By Step 2, distribution functions above p_1 are again uniquely determined, pinning down in turn 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 calculate that this unique candidate is actually the vector of strategies stated in the proposition, and that this unique candidate is indeed an equilibrium.

□

Proof of Proposition 2

The payoff of healer i from playing p while the other healers play H_j is given by

$$\pi_i(p) = p\alpha_i E\left[\prod_{j\neq i}(1-\alpha_j H_j(p))\right] = p\alpha_i \prod_{j\neq i}(1-\overline{\alpha}_j H_j(p)) \tag{10}$$

by the independence of the α_j. This differs from (2) only by a factor of $\overline{\alpha}_i/\alpha_i$ which is independent from p. Thus $H_i(p)$ must be a best response for healer i in the incomplete information game with qualities $\alpha_1, ..., \alpha_n$ as well. (Otherwise the H_i would not form a Nash equilibrium in the complete information game with $\overline{\alpha}_1, ..., \overline{\alpha}_n$.) □

Proof of Proposition 3

"⇐" follows almost immediately: If a strategy profile satisfies (3), the expected payoff of healer i from playing p while the other healers play $G_j^{\alpha_j}$ is the same as in the equilibrium studied in Proposition 2:[30]

$$\pi_i(p) = p\alpha_i \prod_{j \neq i}(1 - E[\alpha_j G_j^{\alpha_j}(p)]) = p\alpha_i \prod_{j \neq i}(1 - \overline{\alpha}_j H_j(p)). \tag{11}$$

Hence it does not make any difference for healer i whether his competitors play H_j or $G_j^{\alpha_j}$. Healer i then does not have an incentive to deviate from $G_i^{\alpha_i}$ because this strategy has support in $[p_0, p_i]$, the support of H_i. Thus all healers playing $G_i^{\alpha_i}$ is an equilibrium.

For "⇒", we first verify that if the $G_j^{\alpha_j}$ are equilibrium strategies, we do not have to worry about atoms. This is needed to justify the expression (12) for the expected payoffs below. Note first that it is inconsistent with equilibrium behavior for a healer j to play a price $\widetilde{p} < 1$ with positive probability in expectation over α_j:[31] If other healers had probability mass on prices marginally above $\widetilde{p}$, they would shift this mass downwards. If no other healers had probability mass on prices marginally above $\widetilde{p}$, healer j could earn more by shifting his probability mass from $\widetilde{p}$ upwards. Additionally, at most one healer j plays a price of 1 with positive probability in expectation over α_j: If several healers did so, at least one of them would have an incentive to shift probability mass downwards to escape tie-breaking.

Since there are no atoms (except possibly one in 1) we can write the payoff of healer i from playing p while the other healers play $G_j^{\alpha_j}$ as

$$\pi_i(p) = p\alpha_i \prod_{j \neq i}(1 - E[\alpha_j G_j^{\alpha_j}(p)]). \tag{12}$$

[30]Note that (3) implies that for all $p < 1$ in expectation over α_j the probability that healer j plays p is zero. We thus do not have to worry about atoms.

[31]Note that we do not rule out in the following that for fixed α_j the distribution function $G_j^{\alpha_j}$ contains atoms. We only show that there are no atoms the other healers can anticipate, i.e., atoms in $E[G_j^{\alpha_j}]$.

Define for each healer i a distribution function

$$\widetilde{H}_i(p) = \frac{1}{\overline{\alpha}_i} E\left[\alpha_i G_i^{\alpha_i}(p)\right].$$

It must hold that $\widetilde{H}_i(p) = H_i(p)$: Clearly, with a similar reasoning as before, if all healers playing $G_i^{\alpha_i}$ is an equilibrium, all healers playing $\widetilde{H}_i(p)$ must be an equilibrium as well, as the expected payoffs from playing any price p are identical in both situations:

$$p\alpha_i \prod_{j \neq i} (1 - \overline{\alpha}_j \widetilde{H}_j(p)) = p\alpha_i \prod_{j \neq i} (1 - E[\alpha_j G_j^{\alpha_j}(p)]). \tag{13}$$

Note that $\widetilde{H}_i(p)$ does not depend on healer i's private information. From comparing the left hand side of (13) with (2) we see that all healers playing $\widetilde{H}_i(p)$ is also an equilibrium of the complete information game with qualities $\overline{\alpha}_1, ..., \overline{\alpha}_n$ since payoffs differ only by a constant factor between the two games. From Proposition 1 we know that $(H_1, \ldots, H_n)$ is the unique equilibrium of the complete information game. This yields $\widetilde{H}_i = H_i$ which by the definition of $\widetilde{H}_i$ implies (3).

Finally, we have to show that the support of $G_i^{\alpha_i}$ lies in $[p_0, p_i]$ for all values of α_i.[32] Note that it does not make any difference for healer i whether his competitors play $(G_j^{\alpha_j})_{\alpha_j}$ or H_j. But playing prices outside $[p_0, p_i]$ against competitors who play H_j is strictly dominated. (This is an easy calculation similar to Step 2 in the proof of Proposition 1). Thus, if $G_i^{\alpha_i}$ is a best response to $(G_j^{\alpha_j})_{\alpha_j}$, its support must be included in $[p_0, p_i]$.

That healers' expected payoffs are the same in all equilibria is a direct consequence of our result that (3) must hold in all equilibria. □

[32]Note that (3) implies already that this holds for F_i-almost all values of α_i.

Proof of Proposition 4

We first start with a definition: Define the function $K_i(\alpha_i)$ as

$$K_i(\alpha_i) = \int_0^{\alpha_i} \beta f_i(\beta) d\beta.$$

Note that K_i is strictly increasing since the integrand is positive for $\alpha_i > 0$. Furthermore, $K_i(0) = 0$ and $K_i(1) = \overline{\alpha}_i$.

Among others, we have to show that there exists a family $(G_i^{\alpha_i})_{\alpha_i}$ which satisfies the sufficient conditions for a Nash equilibrium from Proposition 3 and which consists only of distributions $G_i^{\alpha_i}(p)$ that put all mass on one price. For this purpose, note that a price setting function $\bar{P}_i(\alpha_i)$ translates into a family of distribution functions via

$$G_i^{\alpha_i}(p) = 1_{\{p \geq \bar{P}_i(\alpha_i)\}}.$$

Thus, for a price setting function, (3) becomes

$$\int_0^1 \alpha_i 1_{\{p \geq \bar{P}_i(\alpha_i)\}} f(\alpha_i) d\alpha_i = \overline{\alpha}_i H_i(p). \tag{14}$$

Now, to prove the proposition, we have to show that there exists a unique equilibrium in strictly increasing price setting functions. Define a strictly increasing price setting function via

$$\bar{P}_i(\alpha_i) = H_i^{-1}\left(\frac{K_i(\alpha_i)}{\overline{\alpha}_i}\right).$$

H_i^{-1} denotes the inverse of the restriction of H_i to $[p_0, p_i]$. Thus H_i^{-1} is a bijection from $[0, 1]$ to $[p_0, p_i]$. To see that $\bar{P}_i$ is a well-defined bijection from $[0, 1]$ to $[p_0, p_i]$ note also that $K_i(\alpha_i)/\overline{\alpha}_i$ is a bijection from $[0, 1]$ to $[0, 1]$. That $\bar{P}_i$ is strictly increasing follows because K_i and H_i are strictly increasing. Considering the inverse of $\bar{P}_i$, we see that $\bar{P}_i$ satisfies (3):

$$\begin{aligned} & \bar{P}_i^{-1}(p) = K_i^{-1}(\overline{\alpha}_i H_i(p)) \\ \Leftrightarrow \quad & K_i(\bar{P}_i^{-1}(p)) = \overline{\alpha}_i H_i(p) \\ \Leftrightarrow \quad & \int_0^{\bar{P}_i^{-1}(p)} \alpha_i f_i(\alpha_i) d\alpha_i = \overline{\alpha}_i H_i(p). \end{aligned} \tag{15}$$

As $\bar{P}_i$ is strictly increasing, the final equality is equivalent to (14) and thus to (3). Since $\bar{P}_i$ only takes values in $[p_0, p_i]$, we have hence shown (making use of Proposition 3) that the functions $\bar{P}_i$ form a Nash equilibrium. Furthermore, from (15) it is evident that $\bar{P}_i(\alpha_i)$ is the unique monotonically increasing equilibrium price setting function. □

Proof of Proposition 5

This proposition is an immediate corollary of results derived in Section 3 applied to the symmetric case $F_i = F$. In Proposition 3 we show that the healers' expected payoffs are identical in all equilibria. In Proposition 2 we prove that there is an equilibrium where the expected payoff of healer i is given by

$$\pi_i = \alpha_i(1 - \overline{\alpha})^{n-1}.$$

Proposition 4 shows that there is a unique monotonically increasing strategy equilibrium, given by the price setting function

$$\bar{P}(\alpha_i) = H^{-1}\left(\frac{\int_0^{\alpha_i} \beta f(\beta) d\beta}{\overline{\alpha}}\right)$$

where H is defined in Proposition 1 as

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

with support $[(1-\overline{\alpha})^{n-1}, 1]$. To verify that this is the price setting function stated in Proposition 5 we just have to calculate that

$$H^{-1}(k) = \left(\frac{1-\overline{\alpha}}{1-\overline{\alpha}k}\right)^{n-1}.$$

Inserting $k = \int_0^\alpha \beta f(\beta) d\beta / \overline{\alpha}$ yields the desired result. □

Proof of Proposition 6

Denote by $\alpha_{i:n}$ the i^{th} *lowest* of the values α_1,...,α_n. In the following, we make use of three well-known facts:

First, the density of $\alpha_{i:n}$ is given by[33]

$$f_{i:n}(\alpha) = n\binom{n-1}{i-1} F(\alpha)^{i-1}(1-F(\alpha))^{n-i} f(\alpha).$$

Second, recall the Binomial Theorem: For all $a, b > 0$ and $m \in \mathbb{N}$

$$\sum_{i=0}^{m} \binom{m}{i} a^i b^{m-i} = (a+b)^m. \tag{16}$$

Finally, we make use of the fact that

$$E[1-\widetilde{\alpha}|\widetilde{\alpha} < \alpha] \; = \; \int_0^\alpha (1-\beta)\frac{f(\beta)}{F(\alpha)} d\beta \; = \; 1 - \frac{1}{F(\alpha)} \int_0^\alpha \beta f(\beta) d\beta. \tag{17}$$

We now calculate γ_n in order to verify (5). Recall that each patient consults the worst healer who is recommended to him. Hence the probability that a patient consults healer i equals the probability that i is recommended and that all healers worse than i are not recommended. Thus

$$\gamma_n = E\left[\sum_{i=1}^{n} \alpha_{i:n}^2 \prod_{j=1}^{i-1} (1-\alpha_{j:n})\right].$$

[33] Compare for instance David (1970).

Note that this calculation takes into account that if no healer is recommended, the patient receives a quality of zero. Observe that conditional on $\alpha_{i:n}$ taking some value α, $\prod_{j=1}^{i-1}(1-\alpha_{j:n})$ has the same distribution as a product of $i-1$ independent random variables:

$$\begin{aligned}\gamma_n &= \int_0^1 \sum_{i=1}^n \alpha^2 E\left[\prod_{j=1}^{i-1}(1-\widetilde{\alpha}_j)\middle|\widetilde{\alpha}_1,\ldots,\widetilde{\alpha}_{i-1}<\alpha\right] f_{i:n}(\alpha)d\alpha \\ &= \int_0^1 \sum_{i=1}^n \alpha^2 E[1-\widetilde{\alpha}|\widetilde{\alpha}<\alpha]^{i-1} f_{i:n}(\alpha)d\alpha \qquad (18)\end{aligned}$$

where the $\widetilde{\alpha}_j$ and $\widetilde{\alpha}$ are independent and distributed according to F. Plugging the definition of $f_{i:n}$ into (18) and rearranging yields

$$\gamma_n = n\int_0^1 \alpha^2 \left[\sum_{i=1}^n \binom{n-1}{i-1}\left(F(\alpha)E[1-\widetilde{\alpha}|\widetilde{\alpha}<\alpha]\right)^{i-1}(1-F(\alpha))^{n-i}\right] f(\alpha)d\alpha.$$

After shifting the summation index and inserting (17), this becomes

$$\gamma_n = n\int_0^1 \alpha^2 \left[\sum_{i=0}^{n-1} \binom{n-1}{i}\left(F(\alpha)-\int_0^\alpha \beta f(\beta)d\beta\right)^{i}(1-F(\alpha))^{(n-1)-i}\right] f(\alpha)d\alpha.$$

Applying (16), we obtain

$$\gamma_n = n\int_0^1 \alpha^2 \left[1-\int_0^\alpha \beta f(\beta)d\beta\right]^{n-1} f(\alpha)d\alpha$$

as we wanted to show.

We still have to prove that

$$\lim_{n\to\infty}\gamma_n = 0.$$

Consider the random variable Γ_n which is the quality of the treatment one fixed patient receives in equilibrium. Note that there are two levels of randomness in Γ_n, the randomness in the α_i and the randomness in the recommendations the patient gets. Clearly

$$E[\Gamma_n] = \gamma_n.$$

We first show that Γ_n converges to zero in probability, i.e., for every $\delta > 0$

$$\lim_{n\to\infty} \text{Prob}\,(\Gamma_n < \delta) = 1.$$

Fix $\delta > 0$. The idea is that for any k we can choose n large enough such that with high probability at least k healers have qualities in $(\frac{\delta}{2}, \delta)$. Each of these is recommended with a probability larger than $\frac{\delta}{2}$. Hence if k is large enough it is very likely that one of them is recommended. Precisely, we have to show that for every $\epsilon > 0$ there is some $n(\epsilon)$ such that

$$\text{Prob}\,(\Gamma_n < \delta) \geq 1 - \epsilon$$

for all $n > n(\epsilon)$. Denote by the random variable D^n_δ the number of healers whose qualities lie in the interval $(\frac{\delta}{2}, \delta)$. Denote by R^n_δ the number of healers with qualities in the interval $(\frac{\delta}{2}, \delta)$ who are recommended.

For any $k \leq n$

$$\begin{aligned}
\text{Prob}\,(\Gamma_n < \delta) &\geq \text{Prob}\,(R^n_\delta \geq 1) \\
&\geq \text{Prob}\,(D^n_\delta \geq k)\,\text{Prob}\,(R^n_\delta \geq 1\,|D^n_\delta \geq k) \\
&\geq \text{Prob}\,(D^n_\delta \geq k)\left(1 - (1 - \frac{\delta}{2})^k\right). \qquad (19)
\end{aligned}$$

Choose $k(\epsilon)$ large enough such that

$$\left(1 - (1 - \frac{\delta}{2})^{k(\epsilon)}\right) > \sqrt{1 - \epsilon}.$$

Then choose $n(\epsilon)$ $(= \widetilde{n}(\epsilon, k(\epsilon)))$ large enough such that

$$\text{Prob}\left(D_\delta^{n(\epsilon)} \geq k(\epsilon)\right) > \sqrt{1-\epsilon}.$$

This is possible because we have assumed $f > 0$ which implies that, independently, each healer has with positive probability a quality in $(\frac{\delta}{2}, \delta)$. Then by (19)

$$\text{Prob}\left(\Gamma_n < \delta\right) > \sqrt{1-\epsilon}^2 = 1 - \epsilon$$

for all $n > n(\epsilon)$. Thus we have shown that Γ_n converges to zero in probability. Since Γ_n is bounded, this implies that Γ_n converges to zero in mean (see, for instance, Grimmett and Stirzaker (1992)):

$$\lim_{n\to\infty} \gamma_n = \lim_{n\to\infty} E[\Gamma_n] = 0.$$

□

Proof of Remark 1

Let $\sum_A$ denote $\sum_{A\subset\{1,\dots,n\},\, A\neq\emptyset}$ and let $\#A$ denote the number of healers in A. Since price setting does not depend on the realizations of the α_i, all recommended healers have the same chance of offering the lowest price. Thus the expected quality a patient receives is the expectation of the average quality of the recommended healers:

$$\widetilde{\gamma}_n = E\left[\sum_A \frac{\sum_{i\in A}\alpha_i}{\#A} Prob[A \text{ is the set of recommended healers}]\right].$$

Obviously it holds that

$$Prob[A \text{ is the set of recommended healers}] = \prod_{j \in A} \alpha_j \prod_{k \in A^c} (1 - \alpha_k).$$

Hence

$$\widetilde{\gamma}_n = E\left[\sum_A \frac{1}{\#A} \sum_{i \in A} \left(\alpha_i^2 \prod_{j \in A \setminus \{i\}} \alpha_j \prod_{k \in A^c} (1 - \alpha_k)\right)\right].$$

By the independence of the α_i, this implies

$$\widetilde{\gamma}_n = \sum_A E[\alpha^2] E[\alpha]^{\#A-1} (1 - E[\alpha])^{n-\#A}$$

where α is a random variable with distribution F. Since $\{1, \ldots, n\}$ has $\binom{n}{k}$ subsets with k elements we can rewrite this to

$$\widetilde{\gamma}_n = \frac{E[\alpha^2]}{E[\alpha]} \sum_{k=1}^{n} \binom{n}{k} E[\alpha]^k (1 - E[\alpha])^{n-k}.$$

By the binomial theorem (16) this is the same as

$$\widetilde{\gamma}_n = \frac{E[\alpha^2]}{E[\alpha]} (1 - (1 - E[\alpha])^n).$$

Hence we are done. □

B Survey Questionnaire

Instructions for subjects consisted of the content of Table 1 preceded by the following text:[34]

Please imagine the following hypothetical situation:

[34]translated from German

Assume you want to attend a dentist for professional prophylactic tooth cleaning. In your city, the dentists listed in the following table are all equally easy to reach for you. On a quite new internet portal, you find for each of the dentists exactly one rating by one former patient of the tooth cleaning offered. This rating is either positive ("thumbs up") or negative ("thumbs down").

In the table on the following page you find for each dentist this rating as well as the price of the tooth cleaning. Please mark the dentist who you would prefer to attend for the tooth cleaning! Please mark only one dentist!

The table was followed by the following control question: *How did you decide for the dentist you chose? Please give a short explanation:*

To control for order effects, 46 of the 93 subjects received the treatments of Table 1 in the order "GFADEBC" relabelled "ABCDEFG". This had little effect on the proportion of subjects choosing the cheapest recommended option: In the original order, respectively, 32, 11 and 4 subjects chose options B, F and E. After reordering, 37, 5 and 4 subjects chose the corresponding options.

References

[1] Asch, Steven M., Elizabeth A. McGlynn, Mary M. Hogan, Rodney A. Hayward, Paul Shekelle, Lisa Rubenstein, Joan Keesey, John Adams, and Eve A. Kerr. Comparison of Quality of Care for Patients in the Veterans Health Administration and Patients in a National Sample. Annals of Internal Medicine, 141(12): 938-945, 2004.

[2] 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.

[3] Baye, Michael, Dan Kovenock and Casper de Vries. The all-pay-auction with complete information. Economic Theory, 8: 291-305, 1996.

[4] Borgida, Eugene, and Richard E. Nisbett. The differential impact of abstract versus concrete information on decisions. Journal of Applied Social Psychology, 7: 258- 71, 1977.

[5] Brooks, Janet R. Veterans' health system blazing trails. Canadian Medical Association Journal, 179: 518-519, 2008.

[6] David, Herbert A. Order Statistics, Wiley, New York, 1970.

[7] 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.

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

[9] Enkin, Murray W., and Alejandro R. Jadad. Using anecdotal information in evidence-based health care: Heresy or necessity? Annals of Oncology, 9: 963-966, 1998.

[10] Fagerlin, Angela, Catherine Wang, and Peter A. Ubel. Reducing the influence of anecdotal reasoning on people's health care decisions: is a picture worth a thousand statistics? Medical Decision Making, 25(4): 398-405, 2005.

[11] Finucane, Melissa L., Ali Alhakami, Paul Slovic, and Stephen M. Johnson. The affect heuristic in judgments of risks and benefits. Journal of Behavioral Decision Making, 13: 1-17, 2000.

[12] Fontanarosa, Phil B., and George D. Lundberg. Alternative Medicine Meets Science. Journal of the American Medical Association, 280: 1618-1619, 1998.

[13] Garfield, Joan, and Andrew Ahlgren. Difficulties in Learning Basic Concepts in Probability and Statistics: Implications for Research. Journal for Research in Mathematics Education, 19(1): 44-63, 1988.

[14] Gattellari, Melina, Jeanette E. Ward, and Michael J. Solomon. Randomized, controlled trials in surgery. Diseases of the Colon & Rectum, 44(10): 1413-1420, 2001.

[15] German Federal Ministry of Health, Gesundheitsfonds fördert Wettbewerb, News Release, 08 January 2009, http://www.bundesregierung.de/nn_774/Content/DE/Artikel/2009/01/2009-01-08-gesundheitsfonds.html

[16] Glenton, Claire, Elin S. Nilsen, and Benedicte Carlsen. Lay perceptions of evidence-based information - a qualitative evaluation of a website for back pain sufferers. BMC Health Services Research, 6: 34, 2006.

[17] Grimmett, Geoffrey R., and David R. Stirzaker. Probability and Random Processes, 2nd Edition, Clarendon Press, Oxford, 1992.

[18] Hamill, Ruth, Timothy D. Wilson, and Richard E. Nisbett. Insensitivity to sample bias: generalizing from atypical cases. Journal of Personality and Social Psychology, 39: 578-89, 1980.

[19] Ireland, Norman. The provision of consumer information in a Bertrand oligopoly. Journal of Industrial Economics, 41, 61-76, 1993.

[20] Jenni, Karen E., and George F. Loewenstein. Explaining the "identifiable victim effect". Journal of Risk and Uncertainty, 14: 235-237, 1997.

[21] Kahneman, Daniel, and Amos Tversky. On the psychology of prediction. Psychological Review, 80: 237-51, 1973.

[22] Kallmes, David F., Bryan A. Comstock, Patrick J. Heagerty, Judith A. Turner, David J. Wilson, Terry H. Diamond, Richard Edwards, Leigh A. Gray, Lydia Stout, Sara Owen, William Hollingworth, Basavaraj Ghdoke, Deborah J. Annesley-Williams, Stuart H. Ralston, and Jeffrey G. Jarvik. A Randomized Trial of Vertebroplasty for Osteoporotic Spinal Fractures. The New England Journal of Medicine, 361: 569-579, 2009.

[23] Kirkley, Alexandra, Trevor B. Birmingham, Robert B. Litchfield, J. Robert Giffin, Kevin R. Willits, Cindy J. Wong, Brian G. Feagan, Allan Donner, Sharon H. Griffin, Linda M. D'Ascanio, Janet E. Pope, and Peter J. Fowler. A Randomized Trial of Arthroscopic Surgery for Osteoarthritis of the Knee. The New England Journal of Medicine, 359: 1097-1107, 2008.

[24] Krishna, Vijay. Auction Theory. Academic Press, San Diego, 2002.

[25] Lipkus, Isaac M., Gregory Samsa, and Barbara K. Rimer. General performance on a numeracy scale among highly educated samples. Medical Decision Making, 21(1): 37-44, 2001.

[26] Loewenstein, George F., Elke U. Weber, Christopher K. Hsee, and Ned Welch. Risk as feelings. Psychological Bulletin, 127(2): 267-86, 2001.

[27] Longman, Phillip. The Best Care Anywhere: Why VA Health Care Is Better Than Yours. Second Edition, Polipoint Press, Sausalito, 2010.

[28] McAfee, R. Preston. Endogenous Availability, Cartels and Merger in an Equilibrium Price Dispersion. Journal of Economic Theory 62(1): 24-47, 1994.

[29] McCulloch, Peter, Irving Taylor, Mitsuru Sasako, Bryony Lovett, Damian Griffin. Randomised trials in surgery: problems and possible solutions. BMJ, 324: 1448-1451, 2002.

[30] Osborne, Martin J. and Ariel Rubinstein. A course in game theory. MIT Press, Cambridge, 1994.

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

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

[33] Schwartz, Lisa M., Steven Woloshin, William C. Black, and H. Gilbert Welch. The role of numeracy in understanding the benefit of screening mammography. Annals of Internal Medicine, 127: 966-972, 1997.

[34] 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.

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

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

[37] Spiegler, Ran. But Can't we Get the Same Thing with a Standard Model? Rationalizing Bounded-Rationality Models. Economics and Philosophy, forthcoming, 2010.

[38] Szech, Nora. Becoming a Bad Doctor. Working Paper, Bonn Graduate School of Economics, 2010.

[39] Tversky, Amos, and Daniel Kahneman. Belief in the law of small numbers. Psychological Bulletin. 76(2): 105-110, 1971.

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

[41] Weinstein, Neil D. What does it mean to understand a risk? Evaluating risk comprehension. Journal of the National Cancer Institute Monographs, 25: 15-20, 1999.

[42] Weinstein, Neil D., Kathy Atwood, Elaine Puleo, Robert Fletcher, Graham Colditz, and Karen M. Emmons. Colon cancer: risk perceptions and risk communication. Journal of Health Communication, 9(1): 53-65, 2004.

[43] Yamagishi, Kimihiko. When a 12.86% mortality is more dangerous than 24.14%: implications for risk communication. Applied Cognitive Psychology. 11: 495-506, 1997.