Competitive Experts: A Solution to the Credence Goods Problem

Winand Emons*

University of Bern and CEPR

July 1999

Abstract

This paper is about a market for credence goods. With a credence good consumers are never sure about the extent of the good they actually need. Therefore, sellers act as experts determining the customers' requirements. This information asymmetry between buyers and sellers obviously creates strong incentives for sellers to cheat on services. We analyze whether the market mechanism may induce non-fraudulent seller behavior. From the observation of market data such as prices and market shares consumers attempt to infer the quality of the sellers' services. We show that a market equilibrium resulting in non-fraudulent behavior does indeed exist. This equilibrium is efficient and consumers get the entire surplus.

Keywords: credence goods, expert services, incentives, Bertrand-Edgeworth competition.
Journal of Economic Literature Classification Numbers: D45, D82, L15.

*Universität Bern, Volkswirtschaftliches Institut, Abteilung für Wirtschaftstheorie, Gesellschaftsstraße 49, CH-3012 Bern, Switzerland, 'winand.emons@vwi.unibe.ch', 'http://www-vwi.unibe.ch/theory/emons03.htm'.

1. Introduction

The term credence goods refers to goods and services whose sellers not only provide the good but who also act as experts determining the customers' needs, simply because consumers are unfamiliar with the good in question. This peculiarity often occurs in medical, legal, and financial advice services, as well as in a wide variety of repair professions. It is aggravated by the fact that typically consumers never discover the veracity of a diagnosis, or even whether the treatments they authorized were really performed. Therefore, such services have been termed by Darby and Karni (1973) credence goods.

This information asymmetry creates obvious incentives for opportunistic seller behavior. On the one hand, if there is plenty of money in repair, sellers may recommend unnecessary treatments; this problem has been dubbed 'demand inducement' in the health economics literature. On the other hand, sellers may not perform urgently needed repairs if other activities are more profitable. Since chances of consumers finding out about such fraudulent behavior are slim, how should they view the recommendations of an expert who has a vested interest in providing treatment?

To give a couple of anecdotes where fraud was covered up: In Switzerland patients with the minimum level of schooling are twice as likely to have their womb or gall-stones removed than patients with a university degree; for hip-joint operations the probability is even 150% higher. Ordinary children are 80% more likely to have their tonsils out than children of medical doctors (**K**tip 05/22/1996). In the Canton of Ticino the population average had 33% more of the seven most important operations than medical doctors and their families. Interestingly enough, lawyers and their beloved have about the same operation frequency as the families of medical doctors (Domenighetti et al. (1993)). Perhaps a third of current health-care spending in the US goes on irrelevant tests, unproven procedures, and unnecessarily pricey drugs and devices (Economist 02/13/1999). Even in China overprescription is routine; hospitals use these profits to subsidize underfunded operations (Economist 11/07/1998). Further empirical evidence from the market for physician services suggests that fee-for-service doctors tend to overprescribe while salaried doctors tend to shirk (Gaynor (1994)). In the auto-repair business a survey of 62 automobile repair shops conducted by DOT found that 53% of the service charges were for needless repairs (New York Times 5/8/1979). Even more strikingly, unnecessary repairs were recommended to car owners by employees of Sears Automotive Centers in 90% of the test cases (Wall Street Journal 6/23/92). Other examples include the life-insurance industry where a New York investigation found the sale of unsuitable policies, high-pressure selling, and unbridled sales expenses (Newsweek 2/7/1994), as well as the market for legal advice where the anecdotal evidence is perhaps best summarized by the joke of the longevity study which found that the average lawyer lives twice as long as the average school teacher: Life span for lawyers was computed using billing hours.

Apparently, there is a need for mechanisms to discipline fraudulent experts. One often mentioned mechanism is reputation. This mechanism requires 'watchdog agencies' to verify service quality and repeated relationships to penalize cheating experts. Another mechanism ensuring honest services is the separation of diagnosis and treatment. Unless there is collusion, the diagnosing expert has no incentive to recommend unnecessary treatments and the repairing expert may only fix what has been diagnosed by her colleague. An example of this simple yet effective mechanism is the often encountered separation of the prescription and the preparation of drugs.

This 'separation' mechanism, however, fails to do a good job when it is cheaper to provide diagnosis and repair jointly rather than separately. It is, for example, cheaper to repair any damage while the transmission or belly is open for diagnosis than to put everything back together and repeat the process for the actual repair elsewhere. Apparently, such economies of scope between diagnosis and repair also make the related mechanism of consulting several experts unattractive.

In this paper we want to analyze whether the market may solve the fraudulent expert problem when there are 1. one-shot relationships so that reputation cannot work for lack of punishment possibilities and 2. profound economies of scope between diagnosis and treatment. In our set-up repair is possible only after diagnosis. A customer choosing the services of a second expert, therefore, automatically incurs the cost of a further diagnosis which makes the 'separation' as well as the 'second opinion' mechanisms unattractive. In fact, we assume a sufficiently high cost of switching experts so that a consumer visits just one expert. From the observation of market data such as quoted prices and market shares consumers attempt to infer the sellers' incentives to provide honest/fraudulent services. We show that a market equilibrium inducing non-fraudulent behavior does indeed exist.

We consider experts who are capacity constrained: an active expert may have to ration her clientele due to insufficient capacity or she may also end up with idle capacity. Experts charge separate prices for diagnosis and repair. Competition is thus of the Bertrand-Edgeworth type. First we analyze how an expert's incentives depend on the interplay of prices, capacity, and size of her clientele. If, say, the expert does not have enough customers, she may carry out unnecessary repairs to utilize her otherwise idle capacity; with too many customers she may repair inefficiently little if diagnosis is more profitable than treatment. We show that if the experts charge what we call minimum average cost prices, they are indifferent between diagnosis and repair given enough customers to allow them to work at full capacity: with these prices diagnosis and repair generate the same profit at the margin. Therefore, experts are honest with minimum average cost prices given a clientele permitting them to work at full capacity.

In a second step we analyze whether minimum average cost prices give rise to a non-fraudulent equilibrium. With minimum average cost prices an expert ends up making zero profits. Unilaterally increasing prices above minimum average costs has two effects: On the one hand, the deviating expert loses customers to her colleagues. On the other hand, she makes more money per remaining customer which may, however, be insufficient to cover her overheads. If too many customers leave, such a price increase is unattractive.

If an expert unilaterally increases prices, consumers face the following trade-off: Either they stay with the deviating expensive expert who provides service for sure because she does not have enough customers. Or they go to one of the non-deviating cheap experts where there is, however, a risk of being rationed because they have too many customers; in case of being rationed the consumer ends up with no service at all because the cost of switching to another expert is too high. If the expected loss incurred in going to a non-deviating expert is not too large, too many customers leave the deviating expert and minimum average cost prices do indeed lead to a non-fraudulent market equilibrium. A non-fraudulent equilibrium exists if, essentially, an expert's capacity is sufficiently small relative to the market size. The condition ensuring existence allows for simple comparative statics results which are in line with empirical observations.

The extent of the theoretical literature on fraudulent experts is fairly small. In a classic article Darby and Karni (1973) discuss how reputation, market conditions, and technological factors affect the amount of fraud. Their paper relies heavily on verbal arguments and anecdotes. Yet it contains some of the ideas we formalize in the paper at hand. Pitchik and Schotter (1987) describe a mixed-strategy equilibrium in an expert-customer game. The expert randomizes between either reporting truthfully or not; the customer randomizes between acceptance and rejection of a treatment recommendation. Demski and Sappington (1987) focus on the problem of inducing an expert to acquire a costly expertise. Whereas in our model diagnosis is necessary prior to repair, 'blind treatment' is possible in Demski and Sappington; repair is assumed to be costless. In this set-up they study optimal contracts between a principal and an expert (agent). Wolinsky (1993) examines customer search for multiple opinions and reputation considerations. In his specialization equilibrium some experts exclusively provide diagnosis while the other experts engage in either activity. Consumers first visit a 'diagnosis-only' expert. If she recommends treatment, consumers visit a 'two-activity' expert for a second diagnosis and the actual repair. Our analysis differs from Wolinsky's in several respects. In particular, while our non-fraudulent equilibrium is efficient, Wolinsky's equilibria are not. Taylor (1995) considers experts who, unlike our experts, incur no cost for unnecessary treatments. Unnecessary repairs are thus not inefficient in his set-up. His experts never diagnose a product as healthy; moreover, ex post contracting, free diagnostic checks, consumer procrastination in obtaining checkups, and long-term maintenance agreements may occur in Taylor's equilibria.

Closest to this analysis are our papers Emons (1997) and (1999 a). In Emons (1997) we consider basically the same set-up as in this paper. There we derive an equilibrium where experts are honest and make expected zero profits. Since the experts' entry decisions are in mixed strategies, the equilibrium is, however, inefficient: there are either too few experts in the market charging monopoly prices, or there are too many experts charging loss yielding marginal cost prices. In the paper at hand we show that if consumers incur a sufficiently high cost to switch experts, then in equilibrium the efficient number of experts enters and charges minimum average cost prices. Accordingly, the equilibrium in this paper yields non-fraudulent services *and* efficient entry.

In Emons (1999 a) we consider a credence good monopolist. The monopolist's capacity is determined endogenously. From the observation of i) capacity and prices or ii) just prices consumers attempt to infer the quality of the seller's services. Moreover, we distinguish between the cases of observable and unobservable expert diagnosis and repair services. We show that for three out of the four possible constellations the monopolist always chooses non-fraudulent behavior. Only when capacity and services are non-observable does no trade take place. The three papers are related in their basic result: if consumers rationally process ex ante information, the market mechanism can solve the fraudulent expert problem. Experts are honest in order to maximize the consumers' surplus. In the competitive set-ups honesty is necessary in order to survive; in the monopoly case non-fraudulent service generates the highest profit for the credence goods monopolist. The papers are complements: while two papers show for one informational scenario that competitive markets can solve the fraudulent expert problem, the monopoly paper shows that competition is not necessary and, more importantly, checks the robustness of these results concerning the amount of observable signals.[1)]

The remainder of the paper is organized as follows. In the next section we describe the model. In section 3 we derive the experts' incentive structure. In the subsequent section we establish the existence of the efficient non-fraudulent equilibrium. In section 5 we discuss the welfare properties, perform comparative statics exercises and conclude. Proofs are relegated to the Appendix.

2. The Model

We consider a durable good that is up for diagnosis and potential repair. The good has a remaining capacity which we normalize to 1 monetary unit. During its remaining life, our 'one-hoss shay type' durable either makes available total remaining services 1 or delivers services 0.

The product can be either in good or in bad shape. If the product is in good condition, it makes available services 1 with probability $q_h \in (0, 1)$; when it is in bad shape, the corresponding probability is q_ℓ with $0 < q_\ell < q_h$. In both cases the product may fail but it is more likely to do so when it is in bad shape. Let p denote the probability that the product is in bad shape. The consumer does not know in which of the two conditions his product is, nor can he infer this ex post since both types may break down.

Experts, however, are able to detect the product's condition. By diagnosing the product, an expert finds out whether it is in good or in bad shape. When the product is in bad shape the expert can fix it so that it is in good shape afterwards. Let $d > 0$ be the total resource cost of diagnosing a product; the total resource cost of a repair is $r > 0$.

The timing of the production decisions, however, is such that these costs are not experienced as genuine marginal costs. An expert has to make a prior decision on entry. The expert has L units of time (say, hours) available. If she does not enter the market for expertise, she can work L hours in an alternative activity. If she does enter the expert business, she allocates her L units of time to diagnosis and repair; d is the time an expert needs per diagnosis and r the time per repair. An expert's time cost, however, is sunk. Once she has entered the market, she can only use her time for diagnosis and repair; she can no longer work in the alternative job.

The experts' reservation wage is normalized to 1. Accordingly, L is the sunk cost of becoming active; d and r measure the minimum average costs of diagnosis and repair if, say, the expert performs either activity exclusively. Note that marginal costs are different from average costs. An active expert has fixed capacity the cost of which is sunk. Therefore, her marginal costs are 0 except for the capacity margin where marginal costs are "$+\infty$". When, in the following, we talk about minimum average costs we mean d and r.

There is a continuum of identical consumers with total measure 1.[2)] Consumers are risk neutral and care only about monetary flows. Accordingly, given that we have normalized the product's remaining capacity to 1 monetary unit, without diagnosis and repair a consumer's expected utility is $\bar{U} = (1-p)q_h + pq_\ell$. With (honest) diagnosis and repair priced at minimum average costs the consumer's expected utility amounts to $q_h - d - pr$. The consumer incurs the cost of diagnosis in any case. With probability p the product is in bad shape and needs treatment. In return the consumer has a product that is in good shape for sure.

It is efficient to check the product and fix it if necessary, meaning $q_h - d - pr > \bar{U}$ or $p(q_h - q_\ell) > d + pr$. Fixing a bad product increases the consumer's utility by $(q_h - q_\ell)$. With probability p the product is in bad shape. Accordingly, the expected benefit from diagnosing and repairing the product is $p(q_h - q_\ell)$. As an indicator of the surplus the experts' services may generate, define $W := p(q_h - q_\ell)/(d + pr)$. This surplus measure will play an important role in what follows.

There are I identical experts indexed by $i = 1, \ldots, I$. An expert either enters the market with capacity L or she does not enter at all; we call the former an active and the latter an inactive expert. We assume that repair is possible only after diagnosis.[3)] Given non-fraudulent behavior, an expert's capacity L in units of time thus translates into the capacity $L/(d+pr)$ in terms of customers.

Let $\kappa = (d+pr)/L$. Accordingly, κ non-fraudulent experts are sufficient to serve the entire market. The number of experts, however, has to be an integer. Let k be the smallest integer greater or equal to κ. The efficient number of experts then is either $(k-1)$ or k. If $\kappa = k$, the efficient number of active experts is obviously k and aggregate welfare is $q_h - d - pr$. If $k-1 < \kappa < k$, with $(k-1)$ honest active experts, the fraction $(\kappa - k + 1)/\kappa$ of consumers obtains no service. If k honest experts are active, there is excess capacity. We will assume that $\kappa = k$. Under this assumption experts choose pure strategies in equilibrium. In section 5 we describe a pure strategy equilibrium for $k-1 < \kappa < k$, albeit for a slightly modified game.

Our experts are not 'too large' relative to the market; more specifically, an expert cannot serve more than half of the market given honest behavior. Accordingly, k experts, $k \geq 2, k \in \mathbf{N}$, are sufficient to serve the entire market. Moreover, to ensure competitive behavior there are more experts than necessary to serve the whole market, i.e., $k < I$.

Let us now describe how experts may defraud consumers. The consumer does not know which condition the product is in. Later, when consuming the remaining services he learns whether his product will work or fail. Yet, a good product may break down and a bad product may work satisfactorily. Accordingly, the consumer cannot use the information about his product's later performance to infer its condition at the time when it was up for diagnosis and repair.

After diagnosis the expert knows which condition the product is in. When the product is in bad shape, she can repair it, i.e., turn it into good shape. Yet she can also 'repair' a good product; in this case the expert unnecessarily works r units of time on the product — leaving it at least in good shape.[4)] Alternatively, when the product is in good condition, the expert can recommend not to fix it. Nevertheless, she can make the same recommendation when the product is in bad shape. Ex post the consumer has no way of finding out whether his product was repaired unnecessarily or whether it needed treatment that was not provided. The expert's services thus constitute 'credence' goods as distinct from search and experience goods — from ex post observations the consumer can never be certain of the quality of the services he has purchased. The only possibility for the consumer not to be defrauded is to infer the expert's incentives to be honest from ex ante observable variables such as the quoted prices and market shares.[5)]

An expert picks prices D and R that she charges for diagnosis and repair. Moreover, she chooses a repair policy conditional on the product's condition. We identify this policy by the probability of repair. Let α denote the probability of repair given that the product is in good shape and β the probability of treatment if the product is in bad shape. These two conditional probabilities determine the unconditional ex ante probability of repair $\gamma = (1-p)\alpha + p\beta$ which is quite useful for later purposes.

With this notation we may distinguish three scenarios. If $\alpha = 0,\ \beta = 1$, and thus $\gamma = p$ we talk of *efficient repair*. The expert fixes all bad and no good products; thereafter a product is certainly in good shape. A consumer's expected utility with this honest repair policy is $q_h - D - pR$.

If $\alpha > 0,\ \beta = 1$, and thus $\gamma = (1-p)\alpha + p$ we talk of *too much repair*. The expert not only fixes all bad but also good products. With this fraudulent repair policy a consumer's expected utility amounts to $q_h - D - \gamma R$. Obviously, the consumer prefers efficient repair to too much repair.

Finally, if $\alpha = 0,\ \beta < 1$, and thus $\gamma = \beta p$ we will talk of *too little repair*. The expert fixes no good and not all bad products. With this deceitful repair policy a product may be in bad shape and the consumer's expected utility is $(1-p+\gamma)q_h + (p-\gamma)q_\ell - D - \gamma R$. The consumer prefers efficient to too little repair if $(q_h - q_\ell) \geq R$ which must be satisfied since $(q_h - q_\ell)$ is the consumer's reservation utility for repair if the product is in bad shape. If the expert is indifferent between honest and fraudulent behavior, she behaves honestly. Note that the expert's repair policy defines her capacity in terms of customers $L/(d + \gamma r)$.

Consumers incur no cost to visit one expert. There is, however, a cost to switch from one expert to another. We assume that this switching cost exceeds $p(q_h - q_\ell)$. This switching cost implies that a consumer visits only one expert.[6)] If the consumer is rationed by, say, the expert offering the most attractive terms, he does not go to another expert. The rationed consumer ends up with no service and, accordingly, has utility $\bar{U}$. This simplifying assumption distinguishes our set-up from the standard Bertrand-Edgeworth literature where consumers have no switching cost. In those models customers who are not served by the cheapest seller go to the firm charging the second lowest price etc. until eventually all customers have been served. See, e.g., Allen and Hellwig (1986).[7)]

Let $\eta_i(\cdot)$ denote the probability that a consumer goes to expert $i, i = 1, \ldots, I$. Since consumers have total mass 1, η_i also measures expert i's clientele and her market share. If $\eta_i \leq L/(d + \gamma_i r)$, expert i has enough capacity to treat her entire clientele. Let ζ_i denote the probability of being served by expert i so that, in this case, $\zeta_i = 1$. If $\eta_i > L/(d + \gamma_i r)$, expert i has more customers than she can handle with her repair policy. Therefore, she has to ration her customers randomly and $\zeta_i = L/(d + \gamma_i r)\eta_i$. The number of customers treated by the expert is thus $\min\{\eta_i; L/(d + \gamma_i r)\}$; her expected profit amounts to $\min\{\eta_i; L/(d + \gamma_i r)\}(D_i + \gamma_i R_i) - L$ if she is active and zero otherwise.

Let us now turn to the formulation of the game which we have set up as a three stage game. In the first stage of the game expert $i, i = 1, \ldots, I$, picks prices (D_i, R_i);

expert i stays inactive by quoting $D_i = R_i =$ " $+\infty$". In the second stage consumers observe the quoted prices $(D_i, R_i)_{i=1}^I$. Then each consumer chooses which expert to go to, i.e., each consumer picks $\eta_i \in [0,1]$, $i = 0, \dots, I$, $\sum_{i=0}^I \eta_i = 1$, where η_0 denotes the probability of going to no expert. In the third stage expert i chooses her repair policy $\alpha_i(D_i, R_i, \eta_i)$ and $\beta_i(D_i, R_i, \eta_i)$, $i = 1, \dots, I$.[8)]

In stage two consumers have beliefs $(\hat{\alpha}_i, \hat{\beta}_i)_{i=1}^I$ about the experts' stage three repair policies. Consumers evaluate the expected utility $U(D_i, R_i, \hat{\alpha}_i, \hat{\beta}_i, \zeta_i)$ with each expert i, $i = 1, \dots, I$, according to these beliefs. Each consumer chooses η_i so as to maximize his expected utility; if the offers of several experts are utility maximizing, the consumer gives equal weight to all of them. We confine our attention to symmetric consumer strategies. Experts choose prices and repair policies so as to maximize expected profits.

We focus on subgame perfect equilibria. This means, in particular, that each decision maker acts in a sequentially rational fashion, following a strategy from each point forward that maximizes the expected payoff given the current information and beliefs. In our set-up this implies that the experts' repair policies are indeed optimal once consumers arrive. In equilibrium the consumers' beliefs are borne out: what consumers expect is what experts actually choose to do.

3. Experts' Incentive Structure

Let us begin the analysis by studying the experts' incentives in stage three which are embedded in the functions $\alpha_i(\cdot)$ and $\beta_i(\cdot)$. Recall that an expert enters the market with a capacity of L units of time having a sunk cost L. In terms of customers the expert has capacity $L/(d+pr)$ given honest behavior. Apparently, expert i's behavior depends on the size of her clientele η_i relative to her capacity $L/(d+pr)$. According to whether $\eta_i \gtreqless L/(d+pr)$ we will say that expert i has too many/enough/not enough customers given non-fraudulent behavior. If, say, the expert does not have enough customers, she may start 'repairing' good products to utilize her otherwise idle capacities. If she has too many customers, she may, e.g., be tempted not to fix all bad products given that diagnosis is more profitable than repair.

The last example indicates that the expert's incentives also depend on the relative profitability of diagnosis to repair which, in turn, is determined by her prices D_i and R_i. If the expert has too many customers, the only constraint she faces (at the margin) is her precious time. To maximize profits, she compares the profit per hour repair $(R_i - r)/r$ with the profit per hour diagnosis $(D_i - d)/d$. If the former exceeds the latter, she will repair too much and vice versa if diagnosis is more profitable than treatment. We specify these ideas more precisely in the following Lemma.

Lemma 1:

i) If $\eta_i > L/(d + pr)$*, expert* i *is honest if and only if* $R_i = rD_i/d$*;*

ii) if $\eta_i = L/(d + pr)$*, expert* i *is honest if and only if* $R_i \leq rD_i/d$*;*

iii) if $\eta_i < L/(d + pr)$*, the expert is honest if and only if* $R_i = 0,\ i = 1, \ldots, I$*.*

< insert Figure 1 about here >

The message of Lemma 1 can be seen in Figure 1. Consider the line $R_i = rD_i/d$ along which $(R_i - r)/r = (D_i - d)/d$. Accordingly, on this line the expert is indifferent between diagnosis and treatment so that with too many customers she opts for efficient repair.[9)] In region (I) where $R_i > rD_i/d$ the expert prefers repair to diagnosis. Whatever the number of customers, she will 'fix' anything she diagnoses, i.e., repair too much. In region (II) in which $R_i < rD_i/d$ the expert prefers diagnosis to repair so that she wishes to increase the number of diagnoses at the expense of repairs. With enough customers, however, she cannot diagnose more products; she repairs efficiently to make some money out of her otherwise unused capacity. Along the D_i-axis the expert has proper incentives if she does not have enough customers. She does not repair too much to utilize her idle capacity because there is no money in treatment.[10)]

Subgame perfection implies that the consumers' beliefs $(\hat{\alpha}_i, \hat{\beta}_i)_{i=1}^{I}$ reflect the experts' incentive structure we have just derived. Note that it is possible to pin down the experts' incentives even further once we incorporate the entire list of prices $(D_i, R_i)_{i=1}^{I}$. We will do this in the next section. The most important aspects of the experts' incentives, however, are summarized by Lemma 1.

4. Equilibrium

Let us now tackle the question whether a market equilibrium exists in which experts play pure strategies. As a shortcut we will call such a constellation a pure strategy equilibrium (even though consumers play mixed strategies). Lemma 1 together with the nature of competition has some immediate consequences for the equilibrium. Recall that, given honest behavior, there are more experts than necessary to serve the whole market; competition between experts is of the Bertrand-Edgeworth type; and consumers can deduce non-fraudulent services from the observation of capacity and prices.

These facts have the following implications. Since there are more experts than necessary, in equilibrium either some experts are inactive, or some experts have spare capacity, or all experts exhaust their capacity by over-treatment. Accordingly, there is capacity waiting for 'better' uses. This 'unused' capacity implies that in a pure-strategy equilibrium experts cannot make positive profits. Moreover, this spare capacity ensures non-fraudulent service. Suppose, on the contrary, these observations are not true so that the consumer's utility is less than $q_h - d - pr$. Then an expert can profitably undercut with either an offer on the line $R_i = rD_i/d$ or an offer on the D_i-axis. With too many (not enough) customers, the undercutting expert has proper incentives with the first (second) offer. Thus, whatever the size of her clientele, the undercutting expert can convince consumers of her honest services at a lower cost.[11)] Zero profits together with non-fraudulent services imply that consumers appropriate the entire surplus $p(q_h - q_\ell) - d - pr$, which, in turn, is generated if and only if exactly k experts are active.

Accordingly, if pure strategy equilibria exist, k experts will be active, provide honest service, and make zero profits. We should, therefore, look for non-fraudulent equilibria on the zero-profit line given honest behavior and enough customers $R_i = (d + pr - D_i)/p$. See Figure 1. From Lemma 1 we know already that in region (I) experts repair too much whatever the number of customers. Accordingly, only the lower part of the zero-profit line, i.e., those prices with $D_i \geq d$, contains equilibrium candidates. For reasons that will become clear later, we will pick minimum average cost prices from the incentive compatible part of the zero-profit line and check under what conditions these prices result in a non-fraudulent equilibrium. More specifically, we consider the situation in which the first k experts are active and charge $D_i = d$ and $R_i = r$. With these prices, the experts are honest and make zero profits given enough customers; all consumers are served, i.e., $\zeta_i = 1,\ i = 1, \ldots, k$. The other $I - k$ experts are inactive. Since the proof that such a situation may indeed be an equilibrium is rather long, we proceed in two steps. We first show that, given that the active experts charge minimum average cost prices, no inactive expert wishes to enter. Then we analyze under which conditions active experts do not deviate from minimum average cost prices.

Lemma 2: *Suppose $D_i = d,\ R_i = r,\ i = 1, \ldots, k,\ D_i = R_i =$ "$+\infty$", $i = k+1, \ldots, I$, and the experts' repair policies and the consumers' beliefs are as described in Lemma 1. Then no inactive expert deviates.*

The message of Lemma 2 is by no means surprising (perhaps the only surprising thing about the Lemma is the mess it takes to show that it is really true). The active experts' capacity just matches the consumers' needs so that nobody is rationed, the active experts are honest, and, furthermore, they charge minimum average cost prices so that they end up making zero profits. We, therefore, have an efficient allocation in which all the surplus goes to consumers. Basic economic intuition suggests that there should be no room for an additional expert to enter; this is indeed the case.

The proof of Lemma 2 is along the following lines. For each set of prices and the corresponding repair policy as described by Lemma 1, we assume that the deviating expert has enough customers so as to fully utilize her capacity. Under this assumption she may charge the lowest prices which would allow her to break even. Then we show that even if she charges these most customer-friendly prices, consumers do worse with her than with an established minimum average cost expert. Consequently, for all potentially profitable prices, the entering expert has no customers and makes a loss L.

This last observation is important when we now consider the incentives of an active expert, say expert 1, to deviate from minimum average cost prices. We know from Lemma 2 that for any potentially profitable deviation, a consumer served by expert 1 does worse than a consumer served by a non-deviating expert, say expert 2. Accordingly, with any potentially profitable deviation, expert 1 will lose customers to the other established experts.

This observation does not yet imply that expert 1 loses all customers which would make any deviation from minimum average cost prices unattractive. If some of expert 1's customers go to the other active experts, the latter have a larger clientele than they can handle. They must, therefore, ration their clientele so that, e.g., $\zeta_2 < 1$. This means that one of expert 2's customers now runs the risk of not being served at all which, in turn, implies that his ex ante expected utility of going to expert 2 falls: with probability ζ_2 he is served and has utility $q_h - d - pr$; with probability $(1 - \zeta_2)$ he gets no service and has utility $\bar{U}$.

Consequently, consumers face the following trade-off: Either they go to the 'expensive' expert 1 who, however, provides service for sure because she does not have enough customers; or they go to one of the 'cheap' experts where there is, however, a risk of getting no service at all because they have too many customers. In the following Proposition we show that if the expected loss from being rationed is not too big, an established expert cannot exploit this trade-off to unilaterally deviate from minimum average cost prices so that minimum average cost prices do indeed lead to a non-fraudulent equilibrium.[12)]

Proposition 1: *The strategies $D_i = d,\ R_i = r,\ i = 1, \ldots, k,\ D_i = R_i =$ "$+\infty$", $i = k{+}1, \ldots, I$, together with the experts' repair policies and the consumers' beliefs as described in Lemma 1 constitute an equilibrium if and only if $k \geq W$.*

The proof of Proposition 1 is along similar lines to that of Lemma 2. As we have just described, Lemma 2 implies that with any potentially profitable deviation expert 1 loses customers. We can, therefore, confine our attention to prices promising to be profitable for fewer than $1/k$ customers. For all these deviations and the corresponding repair policies as described by Lemma 1 we assume that expert 1 charges zero-profit prices. Obviously, the zero-profit prices depend on the size of the remaining clientele: the more customers leave, the more the deviating expert has to charge to cover her overheads.[13)]

Then we show that even if expert 1 charges these most customer-friendly break-even prices, consumers do worse with her than with the non-deviating active experts provided that $k \geq W$. If consumers do worse with break-even prices, they will do worse with all potentially profitable prices. Expert 1 will have no customers and, therefore, she will not deviate from minimum average cost prices.

Let us now explain why the condition $k \geq W$ ensures the existence of a non-fraudulent equilibrium. Recall that $W = p(q_h - q_\ell)/(d + pr)$ measures the surplus a consumer gains from non-fraudulent service priced at minimum average cost. Accordingly, W is an indicator of a consumer's loss if he goes to a non-deviating expert and obtains no service. Now suppose the deviating expert loses her entire clientele of size $1/k$ to the other active experts. Then the probability of a consumer being rationed by a non-deviating expert is $1/k$. The term W/k thus measures the risk of going to a non-deviating expert.

Now consider first the special case where expert 1 charges $R_1 = 0$ and $D_1 = d+pr+1$. With these prices expert 1 has proper incentives with not enough customers because there is no money in repair. This deviation is, therefore, the most attractive one for consumers: the deviating expert is more expensive but she is still honest. A consumer going to expert 1 gets non-fraudulent service but pays 1 monetary unit in excess of the expected minimum average costs. If $1 \geq W/k \Leftrightarrow k \geq W$, the expected loss of going to a non-deviating expert is smaller than the loss of consulting the deviating expert. Expert 1 will have no customers and, therefore, she will not deviate.

For all the other possible deviations the expert does not have proper incentives; see Lemma 1. Accordingly, we should expect that if all customers leave the expert with 'honest' deviations, they will certainly desert her with fraudulent deviations. It turns out that this is indeed the case. Once we have taken care of the fraudulent repair policy, the consumers' decision problem basically has the same structure as the trade-off we have just described: even if the expert raises prices just by a little bit, she loses so many customers that her remaining consumers paying the higher prices generate insufficient revenue to cover overheads L.

5. Discussion and Conclusions

The condition $k \geq W$ which is necessary and sufficient for the existence of a non-fraudulent market equilibrium permits some simple comparative statics exercises. The larger k, the smaller the risk of being rationed by a non-deviating expert. The smaller the risk of being rationed, the more attractive non-deviating experts, the less scope there is for increasing prices. We may, therefore, conclude that the larger the number of experts necessary to serve the entire market with honest behavior, the more likely market institutions are to solve the fraudulent expert problem.

Increasing W means that the experts' services generate more surplus for consumers; being rationed becomes less of a joke. If obtaining no service implies, say, losing life and limb, there is plenty of scope for increasing prices. Accordingly, the more valuable the experts' services, the less likely are simple market institutions as described in this paper to solve the fraudulent expert problem.

These comparative statics results are in line with the observation that certain repair services such as, e.g., carpentry, plumbing, and bicycle repair tend to be provided by (free) markets, whereas dental, medical, and legal services are often highly regulated. Suppose with the medical profession that the repair services of surgeons create more surplus per customer than the repair services provided by, say, cobblers. Then our model predicts that botching cobblers are better candidates to be left to the market than gouging surgeons (which also explains the apparent contradiction between the examples of fraud given in the Introduction and the main result of this paper). Furthermore, expert services without extensive regulation, such as many skilled trades, are typically provided by a large number of small-scale experts; this observation is consistent with our comparative statics result concerning k.

The welfare analysis is straightforward. It is obvious that our equilibrium is efficient. Experts are honest and there is neither excess capacity nor rationing. All the surplus goes to consumers. In contrast, any situation with fraud is inefficient: there is either over- or under-treatment compared to the social optimum. Moreover, in our set-up each product is efficiently diagnosed only once. We have, therefore, no welfare losses as in Wolinsky's (1993) specialization equilibrium in which consumers with bad products inefficiently consult two experts.

Let us now explain why we focus on minimum average cost prices as an equilibrium candidate. With all prices on the zero-profit line $R_i = (d + pr - D_i)/p$, experts break even given enough customers; see Figure 1. If $D_i \geq d$, we know from Lemma 1 that with enough customers experts have proper incentives. Furthermore, all prices on the zero-profit line with $D_i \geq d$ and the corresponding honest repair policy give rise to the same expected utility $q_h - d - pr$. Accordingly, we might conjecture that each set of prices on the lower part of the zero-profit line is as good an equilibrium candidate as minimum average cost prices are. This is, however, not the case.

Recall that for any potentially profitable deviation expert 1 loses customers to the other active experts. In our minimum average cost pricing set-up this implies that the non-deviating active experts have to ration their customers. Yet with minimum average cost prices they stay honest given too many customers, see Lemma 1. Thus the deviating expert can only exploit the fact that the other experts are less attractive due to rationing.

For prices on the zero-profit line with $D_i > d$ the incentives of a non-deviating expert are completely different — diagnosis is more attractive than repair. With too many customers, the non-deviating expert will diagnose all products and, accordingly, engage in too little repair. The expert spends less time on repair and thus fixes fewer products than her rationing yet honest minimum average cost colleague does. It follows immediately that consumers do worse with such cheating than with non-fraudulent rationing. Since for zero-profit prices with $D_i > d$ the non-deviating experts are less attractive than our minimum average cost experts, there is more scope to profitably deviate. Consequently, we need stronger conditions than $k \geq W$ to support these prices as equilibria. By focusing on minimum average cost prices we have thus identified the weakest condition ensuring the existence of a non-fraudulent equilibrium. This discussion shows that the equilibrium we have identified is not unique. Other prices on the zero-profit line may also support an equilibrium — though under stronger conditions. Yet recall that in any pure strategy equilibrium experts make zero profit, offer honest services, and consumers get the entire surplus.

If $k - 1 < \kappa < k$, there is another formulation of the game with an efficient equilibrium. Suppose experts pick capacity $l \in [0, L]$ sequentially under perfect information. Expert 1 starts and expert I finishes; the game then continues as our previous game. In the equilibrium of this game the first $(k-1)$ experts choose $l_i = L$, expert k picks $l_k = d + pr - (k-1)L$, and the other experts stay out of the business; the active experts charge $D_i = d$ and $R_i = r$ and are honest. In this pure strategy equilibrium aggregate capacity equals aggregate demand so that it is efficient.

A few remarks concerning the switching costs are in order. First, our identical consumers have unit demand.[14)] Accordingly, all consumers buy as long as the price does not exceed their willingness-to-pay. Suppose there were no switching cost. Then experts can charge up to the reservation price without losing any customers. Thus, unlike Bertrand-Edgeworth models with downward sloping demand, here firms are not penalized for charging high prices. Apparently, we need such a penalty and the switching cost does just this — a customer who goes to a colleague never comes back. Second, recall that the switching cost has to exceed the consumer's benefit from the expert's service. To have existence, this benefit has to be bounded. Accordingly, we only need a 'small' switching cost to support or results.

The relation of our set-up to the literature on Bertrand-Edgeworth competition using the same rationing rule as we do is as follows. In Peters' (1984) model the number of active firms and their capacities are given. He derives a mixed strategy equilibrium in prices. Deneckere and Peck (1995) consider firms choosing prices and capacities.[15)] Their firms may thus cut the price and expand capacity simultaneously, leading to discontinuities in demand. They show that with sufficiently many firms, an equilibrium in pure strategies exists. In a certain sense our approach is between these models. Like Peters' firms our experts have fixed capacity once they are in the market. But our experts decide strategically about entry so that the number of active firms, and thus aggregate capacity, is endogenous. Nevertheless, we do not determine capacity endogenously at the firm level as Deneckere and Peck do.

To conclude, a few remarks for the empirically-inclined reader. Empirical tests of the theoretical results are extremely difficult due to the very nature of the problem: it is fraud that we are looking for. Nevertheless, Marty (1998) shows using 8000 bills of Swiss general practitioners that busy doctors charge significantly less per patient than doctors with insufficient demand, indicating that there is indeed demand inducement. Keeler and Fok (1996) study the impact of an insurance reform in California that, after higher reimbursements for cesarean deliveries, equalized fees for vaginal and cesarean delivery, a relative price shift of 21%. They found a 0.7% nonsignificant drop of cesarean rates. This result, which doesn't appear consistent with Lemma 1, can be explained by other incentive devices such as medical malpractice suits which certainly discipline medical doctors in California. Interestingly enough, despite their empirical result, Keeler and Fok (1996) recommend the equalization of fees as called for by our results 'because it need not hurt providers and may improve patient trust'.

In a simple framework we were able to work out conditions under which the market mechanism can solve the fraudulent expert problem. For a lot of skilled trades offering services of credence quality the market mechanism actually seems to do a fairly good job just as our model predicts; at least we couldn't find any anecdotes of, say, cheating plumbers, electricians, or cobblers.[16)] In other professions, as the examples in the Introduction suggest, there is, however, fraud. The majority of these examples is from the medical profession where the market certainly does not operate in such an unhampered way as is assumed in our model; prices are often set by a regulator rather than the seller, insurers pay for the services, distorting consumers' incentives to gather and process the necessary information, etc. Accordingly, these examples of fraud do not contradict our analysis. Perhaps our results may help to find out what goes wrong in these professions so that better mechanism can be designed to induce honest services. Since expert services are often subject to licensing and regulation, a more thorough understanding of these markets will be helpful for public policy purposes. For credence goods sellers the following strategy recommendations follow from our analysis: With the cost structure given in the paper it is possible to convince rational consumers of the quality of your services. Therefore, try to mimic this cost structure by setting up, e.g., a partnership.

Endnotes

1) Other related theoretical papers include Milgrom and Roberts (1986), Glazer and McGuire (1991), Pitchik and Schotter (1993), Dana and Spier (1993), Wolinsky (1995), Pesendorfer and Wolinsky (1998), Dulleck (1998), Richardson (1999) and Emons (1999 b). For an experimental study mimicking a market for expertise, see Plott and Wilde (1982). There is also a small empirical literature on credence goods. Ekelund et al. (1995) and Mixon (1995) show that sellers of credence goods provide more informational cues such as certification and licensing in the *Yellow Pages* than do sellers of search goods.

2) We make the continuum assumption not only for notational convenience. With a finite number of consumers we run into the following problem. Suppose an expert expects a clientele with $(1-p)$ good and p bad products. With a finite number of customers, however, the actual realization of her clientele will be different from the expected one. Accordingly, at the end of the day she will realize that she has too little or excess capacity and she will start behaving fraudulently. With a continuum of customers we do not encounter this difficulty which would complicate the analysis substantially.

3) This is the standard assumption made in literature; see, e.g., Nitzan and Tzur (1991), Wolinsky (1993), or Taylor (1995). It captures in a straightforward manner the idea that it is cheaper to provide diagnosis and repair jointly rather than separately. An exception is the paper by Demski and Sappington (1987).

4) We assume that diagnosis and repair are verifiable. This assumption is appropriate for, say, dentists whose customers, willy-nilly, suffer any (un-)necessary drilling. It is not appropriate for, e.g., a customer bringing his car to the shop in the morning and picking it up in the evening without being able to tell whether the mechanic has worked on the vehicle. Here the expert has yet another possibility to defraud her customers. She can claim to have fixed the car without having touched it, thus collecting repair fees from an unlimited number of customers. This related problem is dealt with in Emons (1999 a).

5) The fraudulent expert problem may disappear if consumers purchase long term insurance contracts that fully cover all repairs *and* forgone services during the entire product life; such covenants are commonly known as service or health maintenance plans. With these contracts experts have correct incentives since they bear all marginal costs. Yet such long term insurance contracts are particularly prone to consumer moral hazard so that in equilibrium consumers may purchase no service maintenance plans. The problem of too little repair may be solved by a short term warranty for lost services: if the product fails, the expert pays the consumer a sufficiently large amount of money. An honest expert may offer such a warranty at a lower cost than an expert who, say, doesn't repair at all. Such warranties provide experts with an incentive not to cheat. Yet, they may easily fail to do the job when there is consumer moral hazard in the last stage of product life. See, Emons (1988, 1989). For an analysis of performance contracts

(contingent fees) in expertise problems, see Emons (1999 b).

6) High information and search costs to consumers do definitely exist with credence goods. An example is Chadwick's analysis of funeral provisions in England in the 19th century, when there were about 600-700 undertakers in London to provide 120 funerals per day. Chadwick argues that supply-side competitiveness was thwarted by demand-side characteristics such as high search costs and led to monopoly-like conditions over each funeral service; see, e.g., Ekelund and Hébert (1997, p. 217-218). We interpret these obstacles as switching costs.

7) See Carlton (1978), Peters (1984), and Deneckere and Peck (1995) for models on Bertrand-Edgeworth competition in which consumers, as in our set-up, visit only one firm. The switching cost literature also makes basically the same assumption about consumer search. See, e.g., Klemperer (1987).

8) A few remarks for those readers who feel that games should be written down properly: Since players choose simultaneously in stages one and two, we have a game of 'almost perfect' instead of 'perfect' information; see, e.g., Tirole (1988), 431-432. After stage three payoffs are determined as follows. First nature chooses whether the product is in good or bad shape. Then players follow their plans of stages one to three. Finally, nature decides whether the product works or fails and the actual payoffs are realized.

9) In the principal-agent literature a related result is known as the equal compensation principle. See, e.g., Milgrom and Roberts (1992), 228-232.

10) Darby and Karni (1973) also point out that the sellers' incentives depend on the state of demand. When there is 'no customer waiting for service', sellers have an incentive to oversell their services to utilize idle resources; this incentive to oversell disappears when 'the length of the queue of customers waiting for service is positive'. Darby and Karni do not discuss that the sellers' incentives also depend on prices.

11) Note that we have a network externality here. If inert customers do not go to the undercutting expert, she does not have enough customers. With prices $R_i = rD_i/d$, she will not be honest so that it is indeed optimal for consumers not to consult her. If consumers have enough momentum to co-ordinate on the undercutting expert, it is optimal to visit her. If consumers do not have enough momentum, an expert can undercut with prices $D_i > 0,\ R_i = 0$. With these prices she has proper incentives with not enough customers.

12) Without switching cost the rationed consumers would return to expert 1 as long as her prices do not exceed their willingness-to-pay so that she would make positive profits. The switching cost is thus necessary to support the equilibrium. Without this assumption there would be no penalty for raising prices.

13) If the expert could decrease her capacity alongside her price increase, this effect would be absent. It is essential for our argument that capacities are given. See Deneckere and Peck (1995) for a discussion of what happens if firms simultaneously change prices and capacities.

14) See Richardson (1999) for an analysis of a credence good market with non-identical

consumers and a downward sloping demand curve.

15) Carlton (1978) analyzes the same type of competition in a non-strategic context.

16) See also Plott and Wilde (1982, p. 99) who were 'amazed' by how well the market did in their experiments. They conclude that markets as social control devices cannot be dismissed a priori.

Appendix

Proof of Lemma 1: i) If $\eta_i > L/(d + pr)$, the expert has more customers than she can handle with honest behavior. Given her time constraint, she is only interested in the profit per hour repair $(R_i - r)/r$ compared to the profit per hour diagnosis $(D_i - d)/d$. If $R_i = rD_i/d$, which implies $(R_i - r)/r = (D_i - d)/d$, she is indifferent between diagnosis and repair and, therefore, behaves honestly. If $R_i > rD_i/d$, she prefers repair to diagnosis and thus repairs too much and vice versa if $R_i < rD_i/d$.

ii) If $\eta_i = L/(d + pr)$ the expert fully utilizes her capacity with non-fraudulent behavior. If $R_i < rD_i/d$, she strictly prefers diagnosis to repair; yet she makes diagnoses for her entire clientele. She has to repair to use up her remaining time $L - \eta_1 d$; honestly fixing the bad products of her clientele just exhausts her capacity. If $R_i = rD_i/d$, the argument is along similar lines as i). If $R_i > rD_i/d$, the expert strongly prefers repair to diagnosis. Hence, she will repair all products she diagnoses and treat fewer than η_i customers.

iii) If $\eta_i < L/(d + pr)$ expert i has unused capacity with non-fraudulent behavior. As long as $R_i > 0$, she makes money by repairing some more products to use her idle capacity. Only when $R_i = 0$ the incentive for too much repair disappears.

Q.E.D.

Proof of Lemma 2: We have to show that no inactive expert, say expert I, wants to become active. To do this we have to distinguish whether expert I enters with prices $R_I \gtreqless rD_I/d$.

If $R_I = rD_I/d$ and $\eta_I \geq L/(d + pr)$, expert I has proper incentives. Nevertheless, with $D_I > d, \eta_I = 0$ because the k established active experts are cheaper. With $D_I = d$, $\eta_I = 1/(k + 1)$, and $\gamma_I = (L - \eta_I d)/\eta_I r$ which entails no positive profits. Prices with $D_I < d$ entail losses because they are below average costs.

If $R_I > rD_I/d$, for all $\eta_I > 0$ the expert sets $\alpha_I = \beta_I = \gamma_I = 1$, i.e., anything she diagnoses will be repaired. With this repair policy and $\eta_i \geq L/(d + r)$ customers, the expert works at full capacity which in turn allows for the lowest prices to break even. With $\eta_I \geq L/(d + r)$ customers and prices $R_I = d + r - D_I$, expert I makes zero profits. If she charges these full-capacity, zero-profit prices, the expected utility of a consumer of hers amounts to $U(D_I, R_I, \hat{\alpha}_I, \hat{\beta}_I, \zeta_I) \leq q_h - D_I - R_I = q_h - d - r < q_h - d - pr = U(d, r, \hat{\alpha}_1, \hat{\beta}_1, \zeta_1)$ where the first inequality follows from $\zeta_I \leq 1$ and the last expression is the consumer's expected utility with an honest established expert, say expert 1. Even if expert I enters with the lowest possible (i.e., full-capacity, zero-profit) prices, a consumer does better with an established expert. Consequently, for all prices $R_I > rD_I/d$ that allow for non-negative profits, expert I will have no customers and make losses.

If $R_I < rD_I/d$, expert I earns more from diagnosis than from repair. Accordingly, she sets the ex ante probability of repair γ_I to the lowest possible value that still permits her to work at full capacity, i.e.,

$$\gamma_I = \begin{cases} 1, & \text{if } \eta_I \le L/(d+r); \\ (L-\eta_I d)/\eta_I r, & \text{if } L/(d+r) < \eta_I < L/d; \\ 0, & \text{otherwise.} \end{cases}$$

If $\eta_I \le L/(d+r)$, expert I's clientele is so small that she fixes anything she can get hold of, i.e., $\alpha_I = \beta_I = \gamma_I = 1$. She makes zero profits with this repair policy if she charges $D_I = L/\eta_I - R_I$. With these prices a consumer's expected utility amounts to $U(D_I, R_I, \hat{\alpha}_I, \hat{\beta}_I, \zeta_I) = q_h - D_I - R_I = q_h - L/\eta_I \le q_h - d - r < q_h - d - pr = U(d, r, \hat{\alpha}_1, \hat{\beta}_1, \zeta_1)$ where the first inequality follows from $\eta_i \le L/(d+r)$. Even if expert I charges the most favorable prices for consumers, they do better with the established experts. Consequently, with all prices allowing for non-negative profits, expert I has no customers and makes losses.

Next consider $L/(d+r) < \eta_I < L/d$ so that $\gamma_I \in (0,1)$. With this repair policy, expert I makes zero profits by charging $R_I = r(L - \eta_I D_I)/(L - \eta_I d)$. We now have to distinguish whether $\gamma_I \gtrless p$.

If $L/(d+r) < \eta_I < L/(d+pr)$, $\gamma_I > p$. This implies that expert I will fix all bad and some good products, i.e., $\alpha_I > 0, \beta_I = 1$. With this repair policy and zero-profit prices the consumer's expected utility is $U(D_I, R_I, \hat{\alpha}_I, \hat{\beta}_I, \zeta_I) = q_h - D_I - \gamma_I R_I = q_h - L/\eta_I < q_h - d - pr = U(d, r, \hat{\alpha}_1, \hat{\beta}_1, \zeta_1)$ since $\eta_I < L/(d+pr)$. Hence, expert I will attract no customers.

For $\eta_I = L/(d+pr)$, if expert I charges zero-profit prices, the consumers' expected utility with her is the same as with the established experts. Accordingly, she ends up with $1/(k+1)$ customers. But $1/(k+1) < L/(d+pr)$, so that she makes losses.

If $\eta_I > L/(d+pr)$, $\gamma_I < p$. Expert I fixes no good and not all bad products, i.e., $\alpha_I = 0, \beta_I < 1$. With zero profit prices, the consumers' expected utility is $U(D_I, R_I, \hat{\alpha}_I, \hat{\beta}_I, \zeta_I) = q_h - (p - \gamma_I)(q_h - q_\ell) - D_I - \gamma_I R_I = q_h - (p - \gamma_I)(q_h - q_\ell) - L/\eta_I < q_h - r(p - \gamma_I) - L/\eta_I = q_h - d - pr = U(d, r, \hat{\alpha}_1, \hat{\beta}_1, \zeta_1)$. Consequently, expert I attracts no customers.

Finally, if $\eta_I \ge L/d$, expert I has so many customers that she exhausts her capacity only by diagnosing, i.e., $\alpha_I = \beta_I = \gamma_I = 0$. With this repair policy she makes zero profits by charging $D_I = d$. This gives rise to an expected utility $U(D_I, R_I, \hat{\alpha}_I, \hat{\beta}_I, \zeta_I) \le (1-p)q_h + pq_\ell - d < q_h - d - pr = U(d, r, \hat{\alpha}_1, \hat{\beta}_1, \zeta_1)$ because $r < q_h - q_\ell$. Therefore, consumers do better with one of the established experts. Overall then, no profitable entry is possible for an inactive expert.

Q.E.D.

Proof of Proposition 1: From Lemma 2 we know that no inactive expert enters. It remains to be shown that no active expert, say expert 1, deviates. Lemma 2 implies that with any potentially profitable deviation, expert 1 loses customers to the other active experts who then have to ration their clientele randomly. Accordingly, after any deviation, expert 1 will have less than $1/k$ and the other established experts more than $1/k$ customers. We can, therefore, confine our attention to deviations which are potentially profitable for fewer than $1/k$ customers.

If $R_1 > rD_1/d$, expert 1 sets $\alpha_1 = \beta_1 = \gamma_1 = 1$, i.e., anything she diagnoses will be treated. With this repair policy, $\eta_1 \leq L/(d+r)$ customers and prices (D_1, R_i) satisfying $D_1 + R_1 = L/\eta_1$ she breaks even (with more than $L/(d+r)$ customers, she charges $D_1 + R_1 = d + r$ and rations accordingly). If she follows this repair policy and charges zero-profit prices, the utility of one of her consumers amounts to $U(D_1, R_1, \hat{\alpha}_1, \hat{\beta}_1, \hat{\zeta}_1) \leq q_h - L/\eta_1$.

The non-deviating expert 2 has more customers than she can handle so that the probability of being served by her $\zeta_2 < 1$. Accordingly, the expected utility of one of her customers is $U(d, r, \hat{\alpha}_2, \hat{\beta}_2, \zeta_2) = \zeta_2 \left[q_h - d - pr\right] + (1 - \zeta_2)\left[(1-p)q_h + pq_\ell\right]$. With zero-profit prices, expert 1 will have no customers if $U(D_1, R_1, \hat{\alpha}_1, \hat{\beta}_1, \zeta_1) < U(d, r, \hat{\alpha}_2, \hat{\beta}_2, \zeta_2)$ or

$$(1 - \zeta_2)\left[p(q_h - q_\ell) - d - pr\right] < L/\eta_1 - d - pr. \qquad (\dagger)$$

For the probability of being rationed by expert 2 we have

$$1 - \zeta_2 = 1 - \frac{1}{1 - \eta_1} \cdot \frac{k-1}{k} = \frac{\eta_1(L/\eta_1 - d - pr)}{(1 - \eta_1)(d + pr)}.$$

Accordingly, $(\dagger)$ reduces to $W < 1/\eta_1$ which is satisfied since $\eta_1 < 1/k$. Consequently, for prices $R_1 > rD_1/d$, expert 1 has no customers.

Now consider prices with $0 < R_1 \leq rD_1/d$. In this region expert 1 sets γ_1 to the lowest possible value that still permits her to work at full capacity, i.e.,

$$\gamma_1 = \begin{cases} 1, & \text{if } \eta_1 \leq L/(d+r); \\ (L - \eta_1 d)/\eta_1 r, & \text{if } L/(d+r) < \eta_1 < L/(d+pr). \end{cases}$$

With this repair policy the expert makes zero profits by charging $R_1 = r(L - \eta_1 D_1)/(L - \eta_1 d)$. With these prices and repair policy the consumer's expected utility is $U(D_1, R_1, \hat{\alpha}_1, \hat{\beta}_1, \zeta_1) = q_h - L/\eta_1 < U(D_2, R_2, \hat{\alpha}_2, \hat{\beta}_2, \zeta_2)$ as we have just shown. Thus, expert 1 will have no customers with prices $0 < R_1 \leq rD_1/d$.

Consider now those prices above the zero-profit line with $R_1 = 0$. With not enough customers, expert 1 engages in efficient repair, i.e., $\alpha_1 = 0$, $\beta_1 = 1$, and $\gamma_1 = p$, because there is no money in treating products. To break even expert 1 has to charge $D_1 = L/\eta_1$. If the expert charges zero-profit prices, a consumer served by her thus has utility $U(D_1, R_1, \hat{\alpha}_1, \hat{\beta}_1, \zeta_1) = q_h - L/\eta_1$ which is less than the utility with expert 2. Overall then, if $k \geq W$ expert 1 has no incentive to deviate from minimum average cost prices so that we do indeed have an equilibrium.

Conversely, expert 1 does not deviate if her customers' utility from any deviation is less than the utility consumers have with a non-deviating expert. Expert 1 maximizes the utility of her clientele with zero-profit prizes and, as we have just seen, $\max U(D_1, R_1, \hat{\alpha}_1, \hat{\beta}_1, \hat{\zeta}_1) = q_h - L/\eta_1 \leq U(d, r, \hat{\alpha}_2, \hat{\beta}_2, \zeta_2)$ implies $k \geq W$.

Q.E.D.

References

Allen, B. and M. Hellwig, Price-Setting Firms and the Oligopolistic Foundations of Perfect Competition, *American Economic Review, Papers and Proceedings* 76 (1986), 387-392.

Carlton, D., Market Behavior with Demand Uncertainty and Price Inflexibility, *American Economic Review* 68 (1978), 571-588.

Dana, J. and K. Spier, Expertise and Contingent Fees: The Role of Asymmetric Information in Attorney Compensation, *Journal of Law, Economics, and Organization* 9 (1993), 349-367.

Darby, M. R. and E. Karni, Free Competition and the Optimal Amount of Fraud, *Journal of Law and Economics* 16 (1973), 67-88.

Demski, J. and D. Sappington, Delegated Expertise, *Journal of Accounting Research* 25 (1987), 68-89.

Deneckere, R. and J. Peck, Competition over Price and Service Rate when Demand is Stochastic: A Strategic Analysis, *Rand Journal of Economics* 26 (1995), 148-162.

Domenighetti, G., Casabianca, A., Gutzwiller, F., and S. Martinoli, Revisiting the most Informed Consumer of Surgical Services, *International Journal of Technology Assessment in Health Care* 9 (1993), 505-513.

Dulleck, U., The Credence Goods Monopoly Problem, Dept. of Econ. Discussion paper, Humboldt University, Berlin (1998).

Ekelund, R. and R. Hébert, A History of Economic Theory and Method, 4th ed., McGraw-Hill (1997).

Ekelund, R., F. Mixon, and R. Ressler, Advertising and Information: An Empirical Study of Search, Experience and Credence Goods, *Journal of Economic Studies* 22 (1995), 33-43.

Emons, W., Warranties, Moral Hazard, and the Lemons Problem, *Journal of Economic Theory* 46 (1988), 16-33.

Emons, W., On the Limitation of Warranty Duration, *Journal of Industrial Economics* 37 (1989), 287-301.

EMONS, W., Credence Goods and Fraudulent Experts, *Rand Journal of Economics* 28 (1997), 107-119.

EMONS, W., Credence Goods Monopolists, *International Journal of Industrial Organization,* (1999 a) forthcoming.

EMONS, W., Expertise, Contingent Fees, and Insufficient Attorney Effort, *International Review of Law and Economics* (1999 b), forthcoming.

GAYNOR, M., Issues in the Industrial Organization of The Market for Physician Services, *Journal of Economics and Management Strategy* 3 (1994), 211-255.

GLAZER, J. AND T. MCGUIRE, The Economics of Referrals, IIBR Working Paper No. 31/93, Tel Aviv University.

KEELER, E. AND T. FOK, Equalizing Physician Fees Had Little Effect on Cesarean Rates, *Medical Care Research and Review* 53 (1996), 465-471.

KLEMPERER, P., Markets with Consumer Switching Costs, *Quarterly Journal of Economics* 102 (1987), 375-394

MARTY, F., Capacity as a Determinant of the Supply for Physicians' Services, Discussion Paper 05-98, Department of Economics., University of Bern (1998), ftp://www-vwi.unibe.ch/wpapers/marty/phy_0598.pdf

MILGROM, P. AND J. ROBERTS, Relying on the Information of Interested Parties, *Rand Journal of Economics* 17 (1986), 18-32.

MILGROM, P. AND J. ROBERTS, Economics, Organization and Management, Prentice-Hall, 1992.

MIXON, G., Advertising as Information: Further Evidence, *Southern Economic Journal* 61 (1995), 1213-1218.

NITZAN, S. AND J. TZUR, Costly Diagnosis and Price Dispersion, *Economics Letters* 36 (1991), 245-251.

PESENDORFER, W. AND A. WOLINSKY, Second Opinions and Price Competition: Inefficiency in the Market for Expert Advice, Discussion paper, Dept. of Econ., Princeton University (1998), http://www.princeton.edu/~pesendor/asher.pdf

PETERS, M., Bertrand Equilibrium with Capacity Constraints and Restricted Mobility, *Econometrica*, 52 (1984), 1117-1127.

Pitchik, C. and A. Schotter, Honesty in a Model of Strategic Information Transmission, *American Economic Review* 77 (1987), 1032-1036; Errata, *AER* 78 (1988), 1164.

Pitchik, C. and A. Schotter, Information Transmission in Regulated Markets, *Canadian Journal of Economics* 26 (1993), 815-829.

Plott, C. and L. Wilde, Professional Diagnosis vs. Self-Diagnosis: An Experimental Examination of some special Features of Markets with Uncertainty, *in* Research in Experimental Economics, Vol. 2 (V. Smith, ed.), JAI Press (1991), 63-112.

Richardson, H., The Credence Good Problem and the Organization of Health Care Markets, Dept. of Econ. Discussion paper, Texas A&M University (1999).

Taylor, C., The Economics of Breakdowns, Checkups, and Cures, *Journal of Political Economy* 103 (1995), 53-74.

Tirole, J., The Theory of Industrial Organization, MIT Press, 1988.

Wolinsky, A., Competition in a Market for Informed Expert Services, *Rand Journal of Economics* 24 (1993), 380-398.

Wolinsky, A., Competition in Markets for Credence Goods, *Journal of Institutional and Theoretical Economics* 151 (1995), 117-131.

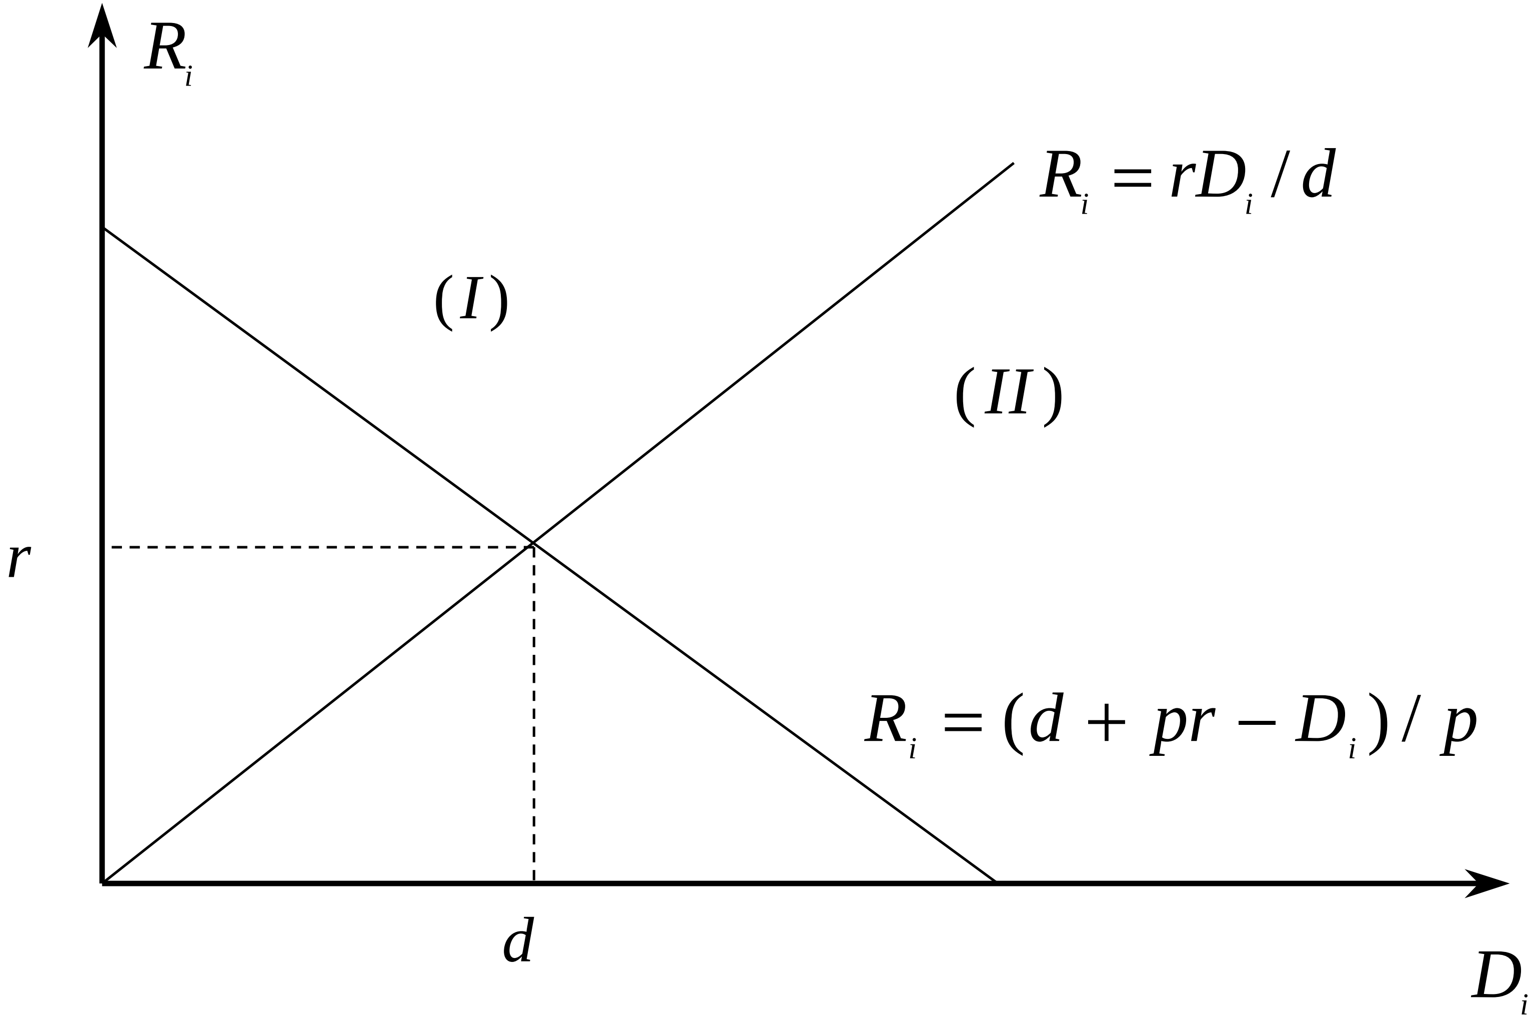

Figure 1: The equal compensation and the zero-profit lines