

A Theory of Fraud and Over-Treatment in Experts Markets

Ingela Alger* and François Salanié**

June 2006

Abstract

Consumers often rely on an expert's diagnosis to assess their needs. If the expert is also the seller of services, he may use his informational advantage to induce over-treatment, which is a pervasive phenomenon in experts markets. We offer and discuss conditions leading to equilibrium over-treatment in an otherwise purely competitive model. This market failure results from consumers' ability to turn down an expert's recommendation: experts defraud consumers in order to keep them uninformed, as this deters them from seeking a better price elsewhere.

Keywords: experts, fraud, over-treatment.

JEL: D41, D82, L11.

*Boston College. ingela.alger@bc.edu

** Université de Toulouse (LERNA-INRA). salanie@toulouse.inra.fr

Acknowledgement: We wish to thank Hideo Konishi, Levis Koshin, Ching-to Albert Ma, and Chris Snyder for helpful comments and discussions, as well as participants at seminars at Boston College, Boston University, Université de Cergy-Pontoise, in the Harvard/MIT/BU Health Seminar series, at Iowa State, Northeastern, and Tufts Universities, the University of Washington, as well as the EARIE meetings in Lausanne, the SETI conference in Houston, and the International Industrial Organization Conference in Boston. We are very grateful to the co-editor, and to two anonymous referees for their detailed and helpful comments. Ingela Alger is grateful for financial support from the Richard W. Tresch Fund for Junior Faculty Research at Boston College. The usual disclaimer applies.

1 Introduction

Economists often assume that consumers know which goods or services they want. In fact, in many situations consumers have to rely on an expert's advice to assess their needs. Examples include all sorts of repairs (car repairs, plumbing), health care, legal and tax services; firms face the same difficulty when choosing computers and software. The informational advantage of an expert suggests that he may have an incentive to make false recommendations, especially if he is also the seller of the services. Emons (1997) cites a Swiss study showing that the average population had 33% more of seven important surgical interventions than physicians and their families. In the late 1970's the Department of Transportation estimated that 53% of auto repair charges represented unnecessary repairs (see Wolinsky, 1993). Together with anecdotal evidence,¹ these observations indicate that we need a better understanding of fraud and inefficient over-treatment in experts markets. Is over-treatment associated with a fundamental market failure? Which mechanism may explain its stability, in spite of competition? And what is its impact on market organization and on welfare?

The impact of competition on over-treatment has been much debated in health economics. In that literature,² the physician's market power over the consumer is a central factor behind the over-treatment phenomenon (also known as the supplier-induced demand hypothesis). The moral hazard arising from the consumer's insurance coverage is believed to exacerbate the problem. However, fraud may well play a role even in the absence of insurance and market power. In this paper, we propose a simple model with risk-neutral, uninsured consumers to investigate whether a competitive experts market performs efficiently. We identify a set of conditions under which the market involves equilibrium fraud and over-treatment.

¹See, for instance, "The Mechanic from Hell" in the New York Times, August 4, 1996.

²McGuire (2000) offers a survey of both theoretical and empirical results, supporting the hypothesis that physicians induce consumers to utilize more health care than they would have chosen if well informed.

The following reasoning shows how competition may in fact favor inefficient over-treatment. The key problem in experts markets is that the consumer knows that there is a loss, but only an expert can determine which treatment is needed. Moreover, once the loss is fixed there is no evidence in favor of one treatment or another. Taken together, these features define a credence good, a term coined by Darby and Karni (1973). Because there typically exist some economies of scope between diagnosis and treatment,³ it is costly for the consumer to get a second opinion. The expert thus enjoys an important strategic advantage over the consumer. Notice however that these features do not create *per se* any incentive to induce over-treatment. They simply make over-charging more likely: upon making the diagnosis the expert may increase the price without fearing that the customer rejects his offer.⁴

To avoid such a hold-up problem, consumers often ask for a commitment to prices. In practice experts commit to a tariff, including prices for different inputs such as spare parts, drugs, or labor. Once the diagnosis is made, the expert provides a bill listing the inputs that he claims are needed for the repair, and computes the repair price accordingly. This system requires the expert to justify the use of additional inputs to increase the repair price. Still, it does not create over-treatment: if the customer cannot observe whether these inputs were actually used or not, the expert will not have to use them. The expert thus defrauds the customer since he makes false claims about inputs used, but over-treatment does not appear.

Suppose now that the customer is able to verify which inputs were actually used. The expert then has to incur an additional “fraud cost” when he lies about the needed

³These may derive from several sources. Performing a diagnosis may require to strip down an engine, thus transferring some of the repair costs to the diagnosis stage. Once the diagnosis is obtained, asking a second expert to repair the loss may involve additional transportation costs; in the health care case, changing doctors may represent losing trust capital developed with the first physician. Finally it may be difficult to transmit precise information about a diagnosis to another party.

⁴In the literature this hold-up problem is associated with the work of Diamond (1971), who considers an exogenous cost of switching from one seller to another; it may be so severe to make the monopoly price prevail as the unique equilibrium price.

treatment, equal to the cost of unnecessary but verifiable inputs. Now, lying increases the price, but also increases the repair cost. Thus an interesting effect appears: avoiding fraud and over-treatment requires the expert to reduce the difference between repair prices (in the limit, if the prices do not differ at all, there is no fraud incentive); in other words, cheap “minor” repairs must subsidize expensive, “major” ones. But such cross-subsidies are difficult to sustain if the customer is allowed to turn down an expert’s offer of a minor repair to find a better price elsewhere. It may therefore well be that competition, and in particular more competition on minor repairs through cream-skimming, favors fraud and over-treatment.

In order to test these intuitions, we set up a model with the following key elements: informational advantage for the expert, economies of scope between diagnosis and repair, commitment to prices, verifiability of some inputs by the customer (fraud cost), opportunistic customers, and competition on minor repairs. The model is inspired by Wolinsky (1993). This is a simple framework in which the loss may require either a minor or a major repair. Experts compete in tariffs, which include repair prices and a price for the diagnosis. The model incorporates an optimal search by the customer among experts. Compared to the literature, we innovate by taking into account the fraud cost originating in the use of unnecessary inputs when an expert imitates a major repair. We also allow for a varying strength of competition on minor repairs, by introducing an outside agent who offers to fix minor losses at a given price.

We first analyze Perfect Bayesian Equilibria when this outside agent offers competitive prices. Equilibria without fraud can be classified into two types. When both economies of scope and the fraud cost are high enough, there exists an efficient equilibrium in which the first visited expert is truthful and fixes the loss. Each repair price is set close enough to marginal cost to deter cream-skimming, without inducing fraud. Otherwise, when the fraud cost is high relative to economies of scope we show, as in Wolinsky (1993), the existence of a “specialization equilibrium” without fraud. Here,

the first visited expert is truthful because the customer rejects his offer of an expensive repair, and gets the loss fixed by another expert. Because the first visited expert only performs minor repairs, he is a “minor specialist” (the outside agent is in fact such a minor specialist). Specialization is inefficient because the diagnosis cost is sometimes incurred twice.

Finally, when the fraud cost is small relative to the economies of scope (and the external agent offers competitive prices), we devise a pure-strategy equilibrium in which experts claim that a costly repair is needed regardless of the customer’s true needs. Thus the consumer does not learn anything and accepts to pay a high repair price, despite knowing that with some probability unnecessary inputs are used; fraud and over-treatment happen in equilibrium. Our theory thus relates fraud and over-treatment to information transmission and cream-skimming: the expert pools information in order to deter the consumer from seeking a better price elsewhere.

By contrast, if the outside agent is removed so that competition on minor repairs is reduced, there exist multiple equilibria. The fraud and specialization equilibria introduced above are still equilibria, for the same parameter values. However, now there always exist equilibria without fraud. More competition on minor repairs, represented by the introduction of an outside agent, indeed favors fraud by making it the only possible equilibrium outcome for some parameter values.

Previous papers have looked at two extreme cases. First, some have assumed that fraud is costless.⁵ Then lying is not inefficient *per se*, and the only source of inefficiency is the cost of getting a second opinion. By contrast we allow for fraud costs. This allows us to properly define over-treatment, confirm its empirical relevance since it appears as

⁵Pitchik and Schotter (1987), Wolinsky (1993), Taylor (1995), Dulleck and Kerschbamer (2006), and Fong (2005). Sometimes the focus is on other manipulations: Taylor (1995) adopts a framework where he can analyze inefficiencies arising in the level of maintenance of the durable good; Emons (1997, 2001) and Fong (2005) allow for under-treatment; Pesendorfer and Wolinsky (2003), and Emons (2001) analyze potential inefficiencies in the amount of effort provided by an expert to make a diagnosis; Wolinsky (1993) allows for diagnosis errors.

an equilibrium phenomenon, and discuss its welfare effects in a meaningful way since now fraud has a welfare cost.⁶ In line with the intuition it turns out that fraud occurs in equilibrium when the fraud cost is low enough, and the diagnosis cost is high enough to deter the customer from seeking a second advice. Conversely, a high fraud cost means that the consumer has full control — it is as if the type of the repair were verifiable. Then our model shows that equilibria are efficient, thus confirming the results obtained by a second group of papers, where repair is verifiable (Dulleck and Kerschbamer, 2006, Emons, 1997 and 2001); they find no fraudulent over-treatment in equilibrium.

Our model allows for an endogenous diagnosis price, often neglected in the literature.⁷ Because experts are allowed to reduce the diagnosis price below its cost, they may redistribute repair profits to customers. Hence all our equilibria are zero-profit equilibria, contrary to some equilibria in Wolinsky (1993) and Dulleck and Kerschbamer (2006). Consequently the equilibrium diagnosis price is set below the diagnosis cost, and can sometimes be set to zero, in line with the evidence. We also verify that customers with minor losses subsidize consumers with major losses.⁸

Our model is also relevant for the health economics literature. This literature typically features a monopolistic competition setup. No strategic interaction occurs between physicians. Each physician selects both price and quantity of treatment, taking the reservation utility of the consumer (defined by other physicians' prices and quantities) as given. Over-treatment occurs when the patient would like to consume less treatment, given the unit price; this is reminiscent of the second-degree price discrimination

⁶Notice that the fraud cost is also equal to the social cost of over-treatment. Hence a higher fraud cost may have ambiguous effects on total welfare, since on the one hand it makes fraud less likely due to better monitoring, while on the other hand it increases its social cost when fraud indeed happens in equilibrium. Caution may therefore be called for when estimating the welfare effects of a policy aiming at deterring fraud.

⁷Pitchik and Schotter (1987) take all prices as given. In Wolinsky (1993), Dulleck and Kerschbamer (2006), and Fong (2005), the diagnosis price is exogenously set at the diagnosis cost.

⁸In our model these cross-subsidies aim at reducing incentives to defraud customers. There are several other theories for why firms in general choose to resort to cross-subsidies between goods or services; in particular similar results arise in theories of loss-leader pricing (see for instance Tirole, 1988, and Lal and Matutes, 1994).

literature (see Farley, 1986).

Some papers have subsequently incorporated an asymmetry of information between the physician and the patient; this enables the physician to affect the consumer's incentives to monitor the doctor, or to end the relationship before treatment. In Dranove (1988) increased competition leads the physician to set a lower price; this in turn causes the patient to accept a treatment offer more readily, so that fraud occurs more often. We also argue that more competition may favor fraud; but our argument relies on the fact that competition renders cross-subsidies less sustainable. Finally De Jaegher and Jegers (2001) provide a precise discussion of how the credence good literature may be applied to the analysis of the supplier-induced demand hypothesis; they also offer a simple model similar to Pitchik and Schotter (1987) in which prices are exogenous.

2 The Model

Our model features a consumer (or, equivalently, a continuum of *ex ante* identical consumers) and experts. Initially the consumer discovers a loss which must be repaired: for instance, the car does not work properly, the house roof leaks, or a tooth aches. The loss may be minor (state m , probability $\mu \in (0, 1)$) or major (state M , probability $1 - \mu$).⁹ Whereas an expert may distinguish between the two states, the consumer cannot.

There are $n \geq 4$ identical experts $i = 1, \dots, n$. An expert may observe the consumer's state at a diagnosis cost $d \geq 0$. The expert may then repair the loss, using some inputs at their given prices. The minimum cost to repair the loss is $\underline{c} \geq 0$ if the loss is minor, and $\bar{c} > \underline{c}$ if the loss is major. For further use, we define C as the minimum expected cost of getting the loss fixed by an expert:

$$C \equiv d + \mu \underline{c} + (1 - \mu) \bar{c}.$$

⁹Using the vocabulary introduced by Dulleck and Kerschbamer (2006), consumers are homogenous. Dulleck and Kerschbamer also study the case where consumers differ in the probability distribution over the two states. In Fong (2005) consumers may differ according to the size of the loss that they incur if the loss is untreated, or in the cost that would be required to fix the loss; the monopolistic expert then uses fraud in the form of over-charging as a substitute for price discrimination.

An expert is needed to make a diagnosis and to make a repair. Moreover, we assume that an expert can repair a loss only if he has made the diagnosis himself. This creates some economies of scope between diagnosis and repair which are measured by d .¹⁰

As is common in the literature we assume that whether the loss has been fixed or not is observable by the consumer. In some cases it may even be verifiable.¹¹ We depart from the rest of the literature by assuming that some inputs are necessary only for a major repair, and that the use of some of these “characteristic inputs” is observable. Let $f \in [0, \bar{c} - \underline{c}]$ denote the cost of these inputs. To illustrate, many illnesses may be treated either through drugs, or through surgery; the latter (the major treatment) is clearly observable. In fact, for deontological and practical reasons the physician must inform the patient and get his consent before the treatment.

The same features arise in many other cases. A successful major repair of a firm’s computer system hinges on the compatibility of new parts with existing devices and usage requirements. The expert would typically then have to reveal at least some of the inputs that he intends to use; by contrast this may not be necessary for a minor debugging operation. Similarly, a household investing in a major house repair would likely want to choose the attributes, such as colors, location and materials, of some inputs. Notice also that in some countries an expert is required by law to commit to a tariff, including prices for different inputs such as spare parts, drugs, or labor. Once the diagnosis is made, the expert provides a bill listing the inputs that he claims are needed for the repair, and computes the repair price accordingly.

We therefore adopt the following assumption:

Assumption 1 *Prior to using characteristic inputs in a major repair, the expert must inform the consumer and get his consent.*

¹⁰Including a switching cost incurred each time a consumer visits an expert would in fact be equivalent to raising the diagnosis cost by that amount. Hence d can be interpreted as a diagnosis-and-switching cost.

¹¹Using the vocabulary introduced by Dulleck and Kerschbamer (2006), verifiability amounts to “liability.” Wolinsky (1993) assumes verifiability.

This assumption can be interpreted as expressing that major repairs are highly complex, and the final selection of inputs depends on the consumer's preferences, which are difficult to fully describe *ex ante*. The assumption may be too strong for simpler repairs: for instance, a consumer would probably not care much about which inputs are used to repair his or her defective DVD player. In Section 5.1 we discuss the consequences of relaxing Assumption 1, in relation with warranty contracts and adverse selection.

The key effect of Assumption 1 is that the consumer is able to identify whether the intended treatment is minor or major, and allows him to break the relationship should he be able to get a better price elsewhere for that treatment.¹² The consumer's ability to observe some inputs which are used only for a major repair also has important consequences: it forces the expert to incur the cost f should he try to charge the consumer for a major repair when in fact a minor repair would be sufficient. To illustrate, suppose that the roof leaks and that the rooftop needs to be replaced (the minor intervention). Suppose also that the roofer (falsely) claims that the rafters are damaged and need to be changed as well (the major intervention). The consumer may easily observe *ex post* whether the rafters have been replaced. In this example f is the cost of replacing the rafters, and it must be added to the cost of the minor intervention. Thus, f represents a fraud cost for the expert: the cost of a minor repair increases by f if the expert claims that the repair is major.¹³ Whenever the expert defrauds the consumer by claiming that the loss is major although it is minor, there is *over-treatment* because of the over-consumption of inputs, as represented by f . Fraud and over-treatment therefore go hand in hand throughout the paper. Finally, an expert cannot claim to make a minor repair if the loss is major, since the expert would need to use the characteristic inputs to fix the loss.

¹²A warranty contract whereby the consumer would pay a price upfront against the expert's promise to fix the loss would not be feasible. If the expert recommended a treatment which is offered at a better price elsewhere, the consumer and the expert would have an incentive to renegotiate. By contrast warranty contracts are feasible in situations where Assumption 1 does not apply, as will be discussed in Section 5.1.

¹³Alternatively, one could interpret the positive fraud cost as the expected penalty associated with being caught by an external auditor, or as a moral cost.

These assumptions imply that experts may set two different repair prices, depending on whether or not characteristic inputs are used. The tariff of an expert i thus consists of a diagnosis price $p_i \geq 0$, a repair price \bar{p}_i to be applied whenever the characteristic inputs are used, and a repair price \underline{p}_i to be applied whenever they are not used. The game proceeds as follows. In the first stage each expert $i = 1, \dots, n$ publicly announces a tariff $(p_i, \underline{p}_i, \bar{p}_i)$, to which he is able to commit.¹⁴ In the second stage the customer sequentially visits experts until his loss is repaired. A visit to an expert i consists of several substages; if in any substage either the expert rejects the consumer or the consumer rejects the expert's offer, the consumer visits another expert. In the first substage the expert accepts or rejects the customer; if the consumer is accepted he pays the diagnosis price p_i , and the expert observes the state of nature at cost d . In the second substage the expert either rejects the customer, or offers a price to repair the loss. If he offers to repair the loss he may choose between price \underline{p}_i and \bar{p}_i if the loss is minor, but he can only offer price \bar{p}_i if the loss is major. Finally the consumer decides whether to reject or accept the offer. If he accepts the expert fixes the loss, and the repair price offered by the expert is paid.

Several remarks are noteworthy here. We assume that an expert may always turn down a customer, either before or after the diagnosis. Indeed it seems realistic to think that an expert may always invoke unexpected delays to avoid serving a customer, and will do so if he expects his profits to be negative. Similarly, the consumer may turn down the expert's offer. This enables the consumer to reconsider his choice of expert upon having updated his beliefs about the state. Notice that it also justifies the constraint $p_i \geq 0$ (otherwise consumers would visit experts simply to collect a negative diagnosis price).

¹⁴The commitment assumption enables us to avoid the hold-up problem mentioned in the introduction. It is equivalent to assuming that experts may commit to input prices, and that consumers know which inputs are needed for each type of treatment. As already observed, repair shops often post their input prices in a place which is visible by the customers. In some countries this is compulsory.

We focus on pure-strategy Perfect Bayesian Equilibria (PBE). The consumer's (behavioral) strategy specifies which experts to visit and the order in which they should be visited, and whether or not to accept the recommendation of a visited expert, depending on the family of tariffs announced in the first stage of the game, and the recommendations that he may have received in the past. We assume that experts cannot observe a consumer's history. Hence expert i 's strategy includes a tariff $(p_i, \underline{p}_i, \bar{p}_i)$, and also a rejection strategy (accept a consumer or not) and a recommendation strategy (reject, offer the price \underline{p}_i , or offer the price \bar{p}_i) which both depend on the observed state of the loss, and on the family of announced tariffs. A given expert may or may not behave truthfully, where truthful behavior consists in offering price \underline{p}_i if the loss is minor.

Given these strategies, the consumer's payoff is the benefit derived from getting the loss fixed minus the sum of all prices paid to the experts he visits (only the diagnosis price if he turns down the expert's offer, but both the diagnosis price and the repair price if he accepts the offer). The consumer's expected payoff is computed using the initial beliefs μ , and the usual Bayesian updating rule, to be applied whenever the consumer receives an offer. The payoff of an expert equals his revenues, which are simply the prices paid by consumers who visited him, minus the costs. The expert's expected payoff is computed using the actual distribution of losses among consumers who decide to visit him. This distribution may differ from the initial distribution $(\mu, 1 - \mu)$ because the customer may visit different experts depending on his history.

We impose three tie-breaking rules which are intended to tilt the balance in favor of more efficient equilibria. First, we assume that if an expert is indifferent between accepting or rejecting a consumer, he accepts; similarly, if he is indifferent between offering \underline{p}_i or \bar{p}_i to a consumer in state m , he offers \underline{p}_i ; finally, if a consumer is indifferent between accepting and rejecting an offer, he accepts it. These tie-breaking rules allow us to focus on deterministic recommendation and acceptance strategies. Some inactive players with a suitably chosen tariff may support an equilibrium. We discard such equilibria by imposing that in each set of experts with the same recommendation strategy:

i) there is at least one expert who is visited by the consumer with positive probability in equilibrium; ii) experts set the same tariff; iii) there are at least two experts.¹⁵ Finally, as we mentioned in the introduction, intuition suggests that competition may favor fraud. Therefore, we begin the analysis by imposing:

Assumption 2 *There are outside agents who offer the “minor specialist” tariff $(d, \underline{c}, \underline{c})$.*

This outside agent repairs minor losses but rejects major losses. A consumer who knows his state may get the loss fixed at the minimum cost (notice that an outside agent’s behavior and tariffs are given, and that he cannot make losses anyway). In Section 5.2 we relax this assumption to study a setup with less competition.

3 Equilibria

As mentioned earlier, our model is close to the one proposed by Wolinsky (1993), in which equilibria are either efficient, or involve inefficient specialization. We first determine conditions under which these types of equilibria exist. We then show that a new type of equilibrium emerges in our setting, namely, equilibria involving costly fraud.

3.1 Efficient Equilibria

Definition 1 *An equilibrium is efficient if the first visited expert makes a truthful recommendation, and the consumer accepts the recommendation.*

Supposing that such an efficient equilibrium exists, let us characterize the tariff $(p, \underline{p}, \bar{p})$ of the first visited expert, taking into account the presence of minor specialists $(d, \underline{c}, \underline{c})$. First, it must be that the expert accepts to repair upon having observed the state:

$$\underline{p} \geq \underline{c} \quad \text{and} \quad \bar{p} \geq \bar{c}.$$

¹⁵While i) indeed discards some equilibria, ii) and iii) could be relaxed and are only introduced to shorten tedious proofs. Condition ii) is a symmetry requirement, and iii) conveys the idea that there are many experts.

By assumption the customer accepts both offers. To deter fraud when a minor loss is diagnosed, the expert's profits when offering \underline{p} must exceed the profits when offering \bar{p} :

$$\underline{p} - \underline{c} \geq \bar{p} - \underline{c} - f.$$

Finally, by receiving the offer \underline{p} the consumer learns that his loss is minor. Since he may reject this offer to visit a minor specialist instead, in an efficient equilibrium one must have

$$\underline{p} \leq \underline{c} + d.$$

Taken together, these inequalities imply

$$f \geq \bar{c} - \underline{c} - d \equiv f^*.$$

In the appendix we show that $f \geq f^*$ is also a sufficient condition for an efficient equilibrium to exist.¹⁶

Proposition 1 *There exists an efficient equilibrium if and only if $f \geq f^* = \bar{c} - \underline{c} - d$. In any such equilibrium experts make zero expected profits. For any $f \geq f^*$ the tariff $(\underline{p}, \bar{p}) = (d - \mu(\bar{c} - \underline{c} - f), \bar{c} - f, \bar{c})$ is an equilibrium tariff.*

This result highlights a fundamental tension between, on the one hand, the incentives to induce the expert to do the repair (\bar{p} must be high enough) and do it efficiently ($\bar{p} - \underline{p}$ must be small enough to deter fraud), and, on the other hand, the consumer's incentive to accept the offer (\underline{p} must be small enough). The last requirement is due to the consumer's opportunism, since he may find a better price elsewhere. We next examine which kinds of inefficiencies may arise when f cannot sustain an efficient equilibrium.

3.2 Specialization Equilibria

Definition 2 *An equilibrium is a specialization equilibrium if the consumer first visits a minor specialist, who makes a truthful recommendation. If the minor specialist offers*

¹⁶Other equilibrium tariffs may exist. We will discuss general properties of equilibrium tariffs in more detail in subsection 4.1.

to do a minor repair the consumer accepts; otherwise the consumer visits another expert, whose recommendation to do a major repair he accepts.

These equilibria formalize the idea that consumers sometimes resort to a second opinion.¹⁷ This phenomenon originates from the fraud potential, even if it does not occur in equilibrium. For the following result, which is proved in the appendix, we define

$$f^{**} = \frac{1 - \mu}{\mu^2} [d - \mu(\bar{c} - \underline{c})].$$

Proposition 2 *There exists a specialization equilibrium if and only if $f \in [f^{**}, f^*]$. For any such f the tariffs $(d, \underline{c}, \underline{c})$ (for minor specialists) and $(d, \underline{c}, \bar{c})$ (for experts repairing only major losses) are equilibrium tariffs. In such an equilibrium experts make zero expected profits.*

A necessary condition for a specialization equilibrium to exist is that $f^{**} < f^*$, which requires

$$d < \frac{\mu}{1 - \mu(1 - \mu)} (\bar{c} - \underline{c}).$$

This condition is intuitive: since specialization involves inefficient diagnoses, the diagnosis cost must be small enough for equilibrium specialization. Moreover, the right-hand side is increasing in μ , the probability that the loss requires a minor repair: the larger is this probability, the larger is the expected benefit from specialization, which implies that an inefficient specialization equilibrium may be sustained for larger values of the diagnosis cost d . Similarly, a larger difference between the repair costs, $\bar{c} - \underline{c}$, also implies a larger benefit from specialization.

As long as there are some economies of scope ($d > 0$) specialization may mean a costly second visit to an expert for a consumer needing a major intervention. As a result,

¹⁷A minor specialist and a consumer in state M would have an incentive to renegotiate, so as to avoid a costly second visit to an expert. But of course, should such a renegotiation be allowed and anticipated, a minor specialist would recommend a major repair in both states. Prohibiting renegotiation, as we do, allows for a stronger competition on minor repairs. Wolinsky (1993) discusses how minor specialists may commit to repair only minor losses, for example by not investing in devices used in major repairs.

if the diagnosis cost and the probability of a major intervention were large, the consumer would prefer to get the loss fixed by one single expert even if that would involve over-treatment: a relevant threat to the specialization equilibrium is an expert who always recommends the major intervention. There is therefore a fundamental trade-off between costly repeated diagnosis and costly over-treatment. For this reason over-treatment appears as an equilibrium phenomenon when the fraud cost is not too large, as we show next.

3.3 Fraud Equilibria

Definition 3 *An equilibrium is a fraud equilibrium if the first visited expert always recommends a major intervention, and the consumer accepts the recommendation.*

An expert who always recommends a major treatment has to incur the fraud cost f whenever the loss is minor. Total repair cost is therefore $C + \mu f$, where the second term measures the inefficiency of a fraud equilibrium. The proof of the following result is in the Appendix.

Proposition 3 *There exists a fraud equilibrium if and only if $f < f^*$ and $f \leq f^{**}$. In such an equilibrium experts make zero expected profits. For any such f the tariff $(p, \underline{p}, \bar{p}) = (d - \mu(\bar{c} - \underline{c} - f), \bar{c}, \bar{c})$ is an equilibrium tariff.*

This result shows that fraudulent over-treatment may appear as an equilibrium even in a competitive model. The reason is that if experts offered both interventions at reasonable prices and recommended interventions honestly, they would lose the profitable customers, i.e., those who needed a minor repair. Fraud avoids consumer defection: the consumer never learns what type of intervention he really needs. But of course, it implies an additional expected cost of μf for the consumer.¹⁸

¹⁸One may extend the model by allowing experts to mimic a minor repair when they in fact proceed to a major one. The associated fraud cost f' would be defined as the cost of hiding the use of inputs that are characteristic of a major repair. Then one could build equilibria with fraud and over-treatment in which experts always claim that a minor repair is needed. The associated inefficiency would be $(1 - \mu)f'$. Since in most cases f' is high compared to f , such equilibria may only emerge when μ is high.

In a fraud equilibrium the repair price is high and independent of the true loss. For a fraud equilibrium to exist the diagnosis cost must therefore be large enough to discourage the consumer from visiting a minor specialist first. This is reflected in the condition $d > \mu(\bar{c} - \underline{c})$, which is necessary for $f^{**} > 0$.

4 Properties of equilibria

So far, we have determined necessary and sufficient conditions for the existence of efficient, specialization, and fraud equilibria.¹⁹ The next proposition states that these are the only type of equilibria that may exist, implying that fraud and over-treatment *must* arise as equilibrium phenomena when there are minor specialists.²⁰

Proposition 4 *If the diagnosis cost d is positive, an equilibrium is either an efficient, a specialization, or a fraud equilibrium.*

Figure 1 provides a graphical representation of our results, which we may compare to those in the closest set of papers. Wolinsky (1993) assumes that $f = 0$; there then exists two regimes, as in our model. If the diagnosis cost d is small there is specialization; if it is large equilibria are efficient. Indeed, when $f = 0$ fraudulent recommendations are costless and therefore involve no inefficiency.²¹ At the other end of the spectrum Dulleck and Kerschbamer (2006) consider the case where $f = \bar{c} - \underline{c}$; as in our model

¹⁹These results were derived under the assumption that the loss had to be fixed. When the customer may replace the damaged good at a price r , similar results (with minor changes in the threshold values) would hold if $r \geq \bar{c}$. By contrast, if $r < \bar{c}$ only the specialization equilibrium would be able to survive. However, the consumer would replace the good without a diagnosis if the diagnosis cost were too large. By contrast to the case where $r \geq \bar{c}$, the market would always perform efficiently, for the consumer would resort to a diagnosis precisely in those cases where the diagnosis cost is sufficiently low to make it worthwhile from a welfare point of view.

²⁰In the case $d = 0$, one may moreover build equilibria with diagnosis providers who provide a truthful diagnosis free of charge.

²¹The threshold value for d below which specialization occurs in our model, $(\bar{c} - \underline{c})\mu$, is smaller than the one obtained by Wolinsky, which is $(\bar{c} - \underline{c})\mu/(1 - \mu)$. This difference arises because Wolinsky does not allow for an endogenous diagnosis price, which implies that in his model experts make strictly positive profits in the equilibrium where consumers visit one single expert.

they find that equilibria are then efficient.

[FIGURE 1 ABOUT HERE]

4.1 Equilibrium tariffs and cross-subsidies

Equilibrium tariffs are typically not unique. For instance, if $f = \bar{c} - \underline{c}$ marginal cost pricing $(d, \underline{c}, \bar{c})$ is an equilibrium tariff; but another equilibrium tariff involves a free diagnosis: $(0, \underline{c} + d, \bar{c} + d)$. The following proposition describes equilibrium tariffs in a more systematic manner (it is a direct consequence of Propositions 1, 2, and 3, and its proof is therefore omitted). While the parameter values for which the different kinds of equilibria exist depend on Assumption 2—see further Proposition 6 in Section 5.2—the qualitative features of equilibrium tariffs as described in the following proposition do not.

Proposition 5 *Marginal cost pricing, where the consumer pays d for a diagnosis, \underline{c} for a minor repair, and \bar{c} for a major repair, occurs only in specialization equilibria (it may also occur in an efficient equilibrium if $f = \bar{c} - \underline{c}$). Otherwise, the diagnosis price may take any values in $[0, d - \mu(\bar{c} - \underline{c} - f)]$, each customer visits only one expert, and experts make positive profits on customers with minor losses, and negative profits on customers with major losses.*

Proposition 5 says that diagnosis prices may sometimes be set below cost, and even at zero.²² It also shows that the largest possible diagnosis price is increasing in the fraud cost f . Indeed the smaller is f , the larger is the profit made on a minor intervention, and the smaller is the diagnosis price. This might explain why car and home repair services typically go together with low diagnosis prices, whereas health care providers in

²²Taylor (1995) finds that there exist equilibria with a zero diagnosis price; however, only the diagnosis *cum* treatment price is uniquely determined, and the diagnosis price may be as well set at diagnosis cost. In Emons (1997), the diagnosis cost is sunk once an expert has entered the market; the zero diagnosis price that his model predicts under certain circumstances thus amounts to marginal cost pricing.

general charge significant amounts for a diagnosis: for repair services consumers observe very few inputs, which suggests a small fraud cost; with health care services it seems reasonable to assume that the opposite is true.

Our model also explains why experts may resort to cross-subsidization among their services. Making repair prices closer sometimes avoids fraud. When fraud cannot be avoided, customers always pay the same repair price, and once again customers with minor losses subsidize customers with major losses. Specialization emerges when these cross-subsidies are not sustainable.

4.2 Welfare

Since demand is perfectly inelastic, welfare is measured by the expected cost of getting the consumer's loss fixed. This expected cost is C for efficient equilibria, $C + (1 - \mu)d$ for specialization equilibria, and $C + \mu f$ for equilibria with fraud. Here we discuss how welfare is affected by changes in the economies of scope parameter d , and by changes in the fraud cost parameter f . Figure 1 provides visual aid for the following remarks.

First, given the number of experts visited in an equilibrium, welfare is decreasing in the diagnosis cost d . Second, it is straightforward to verify that welfare increases in a discontinuous manner when equilibrium switches from the specialization regime to the fraud regime. In the fraud regime welfare decreases less rapidly with the diagnosis cost, since the consumer then visits only one expert in equilibrium. Welfare again increases in a discontinuous manner when reaching the efficient regime.²³ This may be seen in Figure 2, where the values d^* and d^{**} correspond to the critical values of d for which equilibria switch from one kind to another.

[FIGURE 2 ABOUT HERE]

²³The fact that welfare changes in a discrete manner is likely a result of the discreteness of the model. Further research may seek to determine whether the same is true for the non-monotonicity of welfare which is exhibited here.

For a given diagnosis cost d welfare is maximized if the fraud cost f is sufficiently large for an efficient equilibrium to exist, or if equilibrium involves costless fraudulent recommendations ($f = 0$). For values in between, equilibria are inefficient, and welfare decreases in the fraud cost f within the fraud regime. Interestingly, increasing the fraud cost so as to eliminate fraud has a detrimental effect on welfare if specialization arises instead.

[FIGURE 3 ABOUT HERE]

Fraud in experts markets is generally viewed as a problem. Our model suggests a more cautious evaluation: in our setup fraud is an issue from a welfare point of view only if fraud is costly ($f > 0$). If the fraud cost is nil, an equilibrium involving fraudulent recommendations is efficient. Receiving a truthful recommendation is not valuable *per se*; equilibrium recommendation and visit strategies matter only for the costs that they entail.

5 Discussion

In this section we clarify the role played by some of our assumptions, and we discuss some extensions.

5.1 No repair descriptions

Assumption 1 says that prior to using characteristic inputs in a major repair an expert must inform the consumer. This assumption can also be relaxed to better fit simpler repairs. Since then an expert would no longer reveal whether he intends to use a characteristic input, repair prices would no longer be conditional on whether or not characteristic inputs are used, as they were above. As a result warranty contracts would be feasible: an expert could sell a contract whereby he promises to fix any loss at a specific price. Competition would drive the price of such a contract to C .

Clearly such a contract is vulnerable to adverse selection²⁴. Indeed a consumer may prefer to first visit a minor specialist, and in case his loss is major he would then visit another expert proposing a warranty contract. Such an opportunistic behavior can only be deterred if

$$d + \mu \underline{c} + (1 - \mu)C > C,$$

or, equivalently,

$$d > \mu(\bar{c} - \underline{c}).$$

This is reminiscent of the analysis in Wolinsky (1993), where prices are not conditional on any specific inputs to be used either. The market outcome is efficient if the diagnosis cost is large, but otherwise specialization arises. This discussion highlights the main point of this paper, which is that inefficient over-treatment arises if a major repair requires an expert to inform the consumer that some observable inputs will be used. Indeed, equilibrium over-treatment occurs precisely because it enables an expert to avoid revealing information about the true loss.

5.2 The role of minor specialists

In the above analysis efficient equilibria fail to exist when $f < f^*$, because the required mark-up above the minor repair cost cannot be sustained if some experts offer minor repairs at cost. If the existence of minor specialists were not imposed by assumption, the conditions stated in Propositions 1-3 would still be sufficient for the respective equilibria to exist if one allowed for inactive players: a minor specialist with a tariff $(d, \underline{c}, \underline{c})$ makes zero profits in any case, and this tariff always represents a best response by an expert. All the arguments used to prove Propositions 1-3 thus remain valid. However, there would also exist other equilibria, where no expert player offers a specialist tariff. In

²⁴Warranty contracts signed prior to a loss also suffers from the drawback of linking the consumer to a particular expert, which may imply high transportation costs. Admittedly, experts may form networks to mitigate this problem, but the analysis of such networks would require a very different setup.

particular, it turns out that there would exist equilibria without fraud for all parameter values. The following is shown in appendix.

Proposition 6 *If assumption 2 is relaxed there exists an efficient equilibrium if and only if: $f \geq f^*$ or $f \leq f^{**}$.*

For the parameter values for which only fraud equilibria exist under Assumption 2, there now also exists an efficient equilibrium. This equilibrium is described as follows. Suppose that all experts set the same tariff $(p, \underline{p}, \bar{p})$. To deter fraud, the price for a minor repair must be set high enough, and when $f < f^*$ it must be set above $\underline{c} + d$. The most relevant deviation is thus for one of the experts to offer a minor specialist tariff. But this would lead the non-deviating experts to understand that offering to do a minor treatment would lead to consumer defection. Hence non-deviating experts would choose to defraud the customer by always recommending a major treatment. When $f < f^{**}$, the customer indeed prefers to visit a non-deviating expert and accept the (sometimes fraudulent) offer, than to first visit the minor specialist. The deviation is not profitable because it causes experts to change their recommendation behavior.

In real experts markets, this type of equilibrium may be relevant when experts are able to rapidly react to the opening of a “cheap” repair store offering basic repairs at a low price; then the store would attract no customers and ultimately close. The situation is different if a minor specialist appears for another reason: some consumers may be informed about their loss; do-it-yourself kits for minor repairs may be available. This is the situation that our Assumption 2 depicts, and now efficient equilibria disappear for a sufficiently small fraud cost. Overall, competition by minor specialists may thus be an important force behind fraudulent over-treatment in experts markets.

5.3 Mixed strategies

In the above analysis we focused on pure strategy equilibria. Deriving results in a setting where experts and consumers may use mixed strategies happens to be quite difficult.²⁵

²⁵To the best of our knowledge only Pitchik and Schotter (1987) and Wolinsky (1993) allow for mixed strategies in a model of a competitive experts market. However, they impose

The complexity arises from the fact that an expert may *deviate* with a tariff which is followed by a mixed recommendation strategy. We conjecture that fraud equilibria should exist even when such deviations are possible. However, mixed strategy equilibria without fraud may also exist. In particular there may exist equilibria in which experts are truthful because consumers reject their recommendation of a major treatment with a probability large enough to make fraud unprofitable.²⁶ Such equilibria would be slightly more efficient than specialization equilibria. However, as in an efficient equilibrium cross-subsidies between minor and major repairs would be necessary, and therefore the threat of consumer defection would still be relevant. Consider a situation where a small markup above cost for a minor repair successfully prevents consumer defection. Then the expert would have a strong incentive to defraud the consumer and recommend a major repair, unless the consumer very likely rejected a major repair. The fundamental trade-off between costly repeated advice and costly over-treatment would again be at work. For values of f sufficiently small relative to d the consumer should prefer to get a fraudulent recommendation than having to incur the diagnosis cost d twice.

6 Concluding remarks

In this paper, we have proposed a theory of fraud in experts markets which is based on the strategic content of information: an expert may want to lie to the customer in order to keep him uninformed, thereby preventing him from seeking a better price elsewhere. Because the customer is able to observe some of the inputs used, fraud comes at the cost of using unnecessary inputs and is inefficient. This argument applies to various industries, such as car repairs or health care, where information is scarce and each visit to an expert is costly. The model we have solved provides several additional insights.

First, fraud is problematic because it entails inefficient over-treatment. We find that such over-treatment may arise in equilibrium, even if the market is competitive.

exogenously given repair prices.

²⁶Fong (2005) exhibits a similar equilibrium in a monopolistic setting. We were able to verify that in our setting these candidates are not equilibria for $f < \min\{f^{**}, f^*\}$.

This finding is new in the literature, and it may offer a sound and structural basis for applied research on physician-induced demand. In particular, the argument given above supports the so-called physician-induced demand hypothesis, which states in a provocative manner that more competition could lead to more over-treatment. As we have explained, this is due to the opportunism of both experts and customers. Creating commitments through more stringent contracts in which the customer would commit to accept a repair may restore efficiency; in health care the rise of managed health care centers may be explained along these lines. However, such “warranty contracts” are likely to suffer from other important problems, such as adverse selection, and the provision of incentives for the expert to fix the loss once it arises. Further research should address these issues in a systematic manner.

Second, the mere threat of fraud may create welfare costs, even in the absence of equilibrium fraud, by requiring the customers to seek a costly double advice. This is what happens in our model in the specialization equilibrium. Furthermore, welfare may be lower in this regime; eradicating equilibrium fraud may thus be a misleading objective.

Third, our model may explain why diagnosis prices are often set below diagnosis costs. As in most models with switching costs, by doing so experts want to attract consumers in the first place; but another important rationale is that this enables experts to transfer to customers the profits originating from treatment prices being above marginal cost. Our analysis suggests that such mark-ups may be pervasive in experts markets, either as an instrument to deter fraud, or as a direct result of fraud.

Finally, in our framework there is no intrinsic value in obtaining accurate information from the expert. However, such information may matter for third parties; in particular, the performance of insurance markets depends on whether insurance contracts may be contingent on the true state or not. For this reason, and also because insurance is a highly relevant aspect of many experts markets, it would be desirable to extend our model to allow for insurance.

Appendix

In the proofs below we say that an expert is a first visited expert if with some positive probability he is chosen by a consumer for the first visit. For any set of tariffs announced in the first stage of the game, a continuation equilibrium is a PBE of the subgame following these announcements. Let \underline{e} (resp. \bar{e}) denote the best total price for fixing a minor (resp. major) loss in the considered continuation equilibrium. Notice that $\underline{e} \geq \underline{c}$ and $\bar{e} \geq \bar{c}$, as experts can always choose to reject a consumer after the diagnosis stage. We will use ε to refer to an arbitrarily small positive number. Finally, we will use the tie-breaking rules spelled out at the end of section 2, without necessarily referring to them each time.

We begin by proving a lemma that will be used in the proofs to the propositions.

Lemma 1 *Denote $(p, \underline{p}, \bar{p})$ the tariff of a first visited expert in a continuation equilibrium. Under Assumption 2 this expert is either*

- a) a truthful expert who does not make any repair (and then $d = 0$)*
- b) an expert who only repairs major losses (and then $f = \bar{c} - \underline{c}$ and $\bar{p} = \bar{c}$)*
- c) a truthful expert who repairs only minor losses and earns zero-profits*
- d) a truthful expert who repairs all losses (and then $f \geq f^*$)*
- e) an expert who always offers a major treatment and repairs all losses, for a total price $p + \bar{p} \in [C + \mu f, \frac{d}{\mu} + \underline{c}]$ (and then $f \leq f^{**}$).*

Proof of Lemma 1. Suppose first that no losses are repaired at the customer's first visit. Then this visit would be useless unless the consumer could deduce his true type from the expert's offer after the diagnosis stage. Moreover, by assumption we know that at the first visit the customer would reject all the repair offers that the expert may make. Therefore any consumer, at any stage of the continuation game, would also reject the repair offers; this shows that the expert never repairs any losses. To avoid making non-negative profits, it must then be that $p \geq d$. After his first visit, the consumer is informed about his loss, and he therefore ends up paying $p + \mu \underline{e} + (1 - \mu)\bar{e}$. By

visiting a minor specialist first, the consumer would have paid $d + \mu \underline{c} + (1 - \mu)\bar{e}$, so that $p + \mu \underline{e} \leq d + \mu \underline{c}$. Since $\underline{e} \geq \underline{c}$ and $p \geq d$, this is only possible if $p = d$ and $\underline{e} = \underline{c}$. The latter equality implies that there exists an expert with a tariff $(0, \underline{c}, \bar{p}')$, who accepts consumers and proposes \underline{c} to consumers with minor losses. Since this expert accepts consumers, to avoid getting negative profits he must expect to sometimes repair major losses through an offer \bar{p}' that must be accepted, so that $\bar{p}' \leq \bar{e}$. But then the consumer could visit this expert first, and pay $\mu \underline{c} + (1 - \mu)\bar{p}'$; as a result we must have $\mu \underline{c} + (1 - \mu)\bar{p}' \geq p + \mu \underline{e} + (1 - \mu)\bar{e}$, which is possible only if $p = 0$. Since we know that $p \geq d$, we get $d = 0$, as announced in case a of the lemma.

Second, suppose that only major losses are repaired at the customer's first visit (assumption A). Then the expert must propose \bar{p} to major losses, and this offer is accepted by the customer at his first visit. Moreover, \bar{p} must never be proposed to minor losses, because this would be accepted by the customer at his first visit, thus contradicting assumption A. Assumption A also implies that either the expert rejects minor losses, or his offer \underline{p} is rejected by the customer at his first visit (and would therefore also be rejected by any customer at any stage of the continuation game). As a result the expert earns zero profits on the treatment of minor losses. But by proposing \bar{p} to minor losses, the expert could earn $\bar{p} - \underline{c} - f$ (since \bar{p} is accepted by the customer at his first visit). It must therefore be that $\bar{p} - \underline{c} - f \leq 0$, but since $\bar{p} \geq \bar{c}$ and $f \leq \bar{c} - \underline{c}$ this implies that $f = \bar{c} - \underline{c}$ and $\bar{p} = \bar{c}$, as announced in case b of the lemma.

Third, suppose that only minor losses are repaired at the customer's first visit (assumption A'). Let us first show that the expert never repairs major losses, for any customer at any stage of the continuation game. This is clear if $\bar{p} < \bar{c}$. If $\bar{p} \geq \bar{c}$ our tie-breaking rules imply that the expert must propose \bar{p} to major losses. Under assumption A' \bar{p} must be rejected by a customer at his first visit, so that the expert repairs minor losses through the offer \underline{p} . Therefore from an offer \bar{p} any consumer can infer that his loss is major. As a result, if the offer \bar{p} is rejected by the customer at his first visit, it would also be rejected by a customer who is not at his first visit. We may conclude that the

expert never repairs major losses. Since under Assumption 2 the consumer can visit a minor specialist first, our expert must make zero profits, as announced in case c.

Fourth, suppose that all losses are repaired at the customer's first visit. If the expert is truthful, then it must be under Assumption 2 that $f \geq f^*$, as we have argued in section 3.1; so we get case d.

Finally, if the expert always recommends a major treatment, the consumer ends up paying $p + \bar{p}$. But under Assumption 2 the consumer could have visited a minor specialist first, so that we must have

$$p + \bar{p} \leq d + \mu \underline{c} + (1 - \mu)(p + \bar{p})$$

or, equivalently, $\mu(p + \bar{p}) \leq d + \mu \underline{c}$. Moreover, $p + \bar{p} \geq C + \mu f$ because otherwise the expert would make negative profits.²⁷ Together these inequalities yield $f \leq f^{**}$, as announced in item e. ■

Proof of Proposition 1. The condition $f \geq f^*$ was shown to be necessary in the text. Conversely let us show that if $f \geq f^*$ there exists an efficient equilibrium in which all players offer the tariff $(p, \underline{p}, \bar{p}) = (d - \mu(\bar{c} - \underline{c} - f), \bar{c} - f, \bar{c})$. Note that the diagnosis price is positive, and that the assumption $f \in [0, \bar{c} - \underline{c}]$ implies that \underline{p} is above \underline{c} . With this tariff the expert has no incentive to make a false recommendation given that the consumer accepts, and the consumer accepts both recommendations because turning either of them down would only imply a larger cost to get the loss fixed (indeed $f \geq f^*$ implies that $\underline{p} \leq \underline{c} + d$). The experts make zero expected profits, and the consumer's expected cost for getting the loss fixed is C . We now prove that the condition $f \geq f^*$ guarantees that there exists no profitable deviation in tariff.

²⁷Because experts do not randomize, at any stage of a continuation game consumers are either uninformed (with beliefs μ) or informed of their loss type. Moreover $f \in [0, \bar{c} - \underline{c}]$ implies that $\underline{c} + d \leq C + \mu f \leq \bar{c} + d$. Therefore $p + \bar{p} < C + \mu f$ would imply that the expert gets negative profits on both uninformed consumers and consumers knowing that their loss is major. Moreover, under Assumption 2 the expert cannot make profits by attracting consumers knowing they have a minor loss.

Suppose that after a deviation $(p', \underline{p}', \bar{p}')$ a consumer visiting a non-deviating expert would still accept the recommendation $\underline{p} = \bar{c} - f$. Consider the following strategies: the non-deviating experts still accept consumers and are truthful, and the consumer visits a non-deviating expert at his first visit, and accepts both recommendations. Notice that because non-deviating experts are truthful it is indeed rational for the consumer to accept recommendation \bar{c} , since anyway the best alternative offer \bar{e} is at least equal to \bar{c} . And when the consumer accepts both recommendations an expert does not lie (indeed $\underline{p} = \bar{p} - f$). Therefore, these strategies form an equilibrium of the subgame in which the consumer chooses to first visit a non-deviating expert; in this subgame the consumer ends up paying C . For the consumer to choose a different visit pattern he must end up paying less than C ; but then the deviation cannot be strictly profitable unless some experts get negative profits, a contradiction with the fact that experts must play best responses in the continuation equilibrium following the deviation.

Now suppose that the deviating expert offers to treat minor losses for strictly less than $\bar{c} - f$. Because $f \geq f^*$, this implies that the deviating expert would get negative profits on customers with minor losses. A profitable deviation must therefore make profits on customers with major losses, so that $p + \bar{p} > \bar{c} + d$. Because $\bar{c} + d \geq \bar{c} - f$, the deviating expert cannot make the same offer to minor and major losses. He must therefore be truthful and propose \underline{p} to minor losses and \bar{p} to major losses, and both offers are accepted. But then

$$\bar{p} - \underline{p} = (p + \bar{p}) - (p + \underline{p}) > (\bar{c} + d) - (\bar{c} - f) = d + f \geq f,$$

in contradiction with the deviating expert being truthful.

It only remains to be shown that in all efficient equilibria profits must be zero. Consider such an efficient equilibrium, in which all players set the same efficient tariff $(p, \underline{p}, \bar{p})$. Suppose that experts make positive profits:

$$p - d + \mu(\underline{p} - \underline{c}) + (1 - \mu)(\bar{p} - \bar{c}) > 0.$$

Because $\bar{p} \leq \underline{p} + f$ and $\underline{p} \leq \underline{c} + d$ (under Assumption 2) and $f \leq \bar{c} - \underline{c}$, this implies that $p > 0$. Then any expert could deviate to $(p - \varepsilon, \underline{p}, \bar{p})$. With such a tariff the recommendation of a minor treatment is accepted, and the expert is truthful from our tie-breaking rules. Therefore the recommendation of a major treatment is also accepted. Clearly the consumer prefers to visit this expert first. The deviating expert's profits would increase since he would attract the whole market instead of only a share of it. ■

Next we state and prove a lemma which will be used in the proofs of Propositions 2-4.

Lemma 2 *Consider an inefficient equilibrium E in which all experts set the same tariff $(p, \underline{p}, \bar{p})$, with $\bar{p} \geq \bar{c}$ and $p + \bar{p} > \bar{c}$. Under Assumption 2, for any $p_1 \in [0, p + \bar{p} - \bar{c})$, we have:*

(i) *If an expert (say, expert 1) deviates towards the tariff $(p_1, \bar{c} - f, \bar{c})$, then in any continuation game following this deviation all major losses are repaired by expert 1.*

(ii) *Through this deviation expert 1 gets a profit $B_M \equiv p_1 - d$ on each customer with a major loss who visits him.*

(iii) *If through this deviation expert 1 does not repair the minor loss of a customer who visits him, then his profit on this customer is $B_m = B_M$. In any case $B_m \geq B_M$.*

(iv) *If he also repairs all minor losses, then total profits of expert 1 are $p_1 - d + \mu(\bar{c} - \underline{c} - f)$.*

Proof of Lemma 2. We proceed by contradiction and suppose that in a continuation equilibrium following the deviation major losses are repaired by non-deviating experts. Then it must be that the customer visits a non-deviating expert at his first visit: indeed, if the customer knew that his loss is major before such a visit, he would be strictly better off visiting expert 1 instead, since by assumption expert 1 offers a strictly lower price ($p_1 + \bar{c} < p + \bar{p}$) for fixing major losses.

If the non-deviating expert repairs only major losses, we are necessarily in case b of Lemma 1, so that $f = \bar{c} - \underline{c}$ and $\bar{p} = \bar{c}$. Therefore, the tariff of expert 1 is $(p_1, \underline{c},$

\bar{c}), and the assumption $p_1 + \bar{c} < p + \bar{p}$ reduces to $p_1 < p$. Also, from our tie-breaking rules expert 1 is truthful. By visiting a non-deviating expert first the customer pays $p + \mu\underline{e} + (1 - \mu)\bar{p}$; by visiting expert 1 first he would pay $p_1 + \mu\underline{c} + (1 - \mu)\bar{c}$, which is strictly less because $p_1 < p$ and $\underline{c} \leq \underline{e}$, and we reach a contradiction.

If the non-deviating expert repairs all losses through a major treatment, then the consumer pays $p + \bar{p}$ while by visiting expert 1 he could have paid $p_1 + \bar{c}$, which is strictly lower by assumption.

Finally, the non-deviating expert may be truthful and repair all losses, which would require $\bar{p} - f \leq \underline{p} \leq \underline{c} + d$. This implies that the non-deviating experts were also truthful when playing the equilibrium E in the absence of deviation, and all their recommendations were accepted. Because we have assumed that the equilibrium E was inefficient, then it must be that under E the consumer preferred to visit a minor specialist first, so that

$$d + \mu\underline{c} + (1 - \mu)(p + \bar{p}) \leq p + \mu\underline{p} + (1 - \mu)\bar{p}.$$

But since $p_1 + \bar{c} < p + \bar{p}$ this implies that

$$d + \mu\underline{c} + (1 - \mu)(p_1 + \bar{c}) < p + \mu\underline{p} + (1 - \mu)\bar{p}$$

so that the consumer strictly prefers to visit a minor specialist first, to visiting a non-deviating expert first. But then non-deviating experts cannot repair major losses, as explained at the beginning of this proof. We have thereby proved item (i) in the lemma.

The expression for B_M obtains directly. For minor losses, either the deviating expert rejects them and gets $B_m = B_M$; or, he repairs them and gets $B_M + \bar{c} - \underline{c} - f \geq B_M$. Items (iii) and (iv) follow. ■

Proof of Proposition 2. We first show that a specialization equilibrium exists if $f \in [f^{**}, f^*]$. Consider the following candidate: at least two experts (“minor specialists”) play the tariff $(d, \underline{c}, \underline{c})$; at least two experts (“major specialists”) play the tariff $(d, \underline{c},$

\bar{c}); all experts play one of these tariffs; consumers and experts behave as in Definition 2. One can easily check that in the continuation game associated to these tariffs no player can profitably deviate. Minor specialists reject major losses and are truthful. Major specialists lie because their offer of a major treatment is accepted, and because lying makes their profits increase ($\bar{c} - f > \underline{c}$ from $f < f^*$). Finally the consumer behaves as in Definition 2; indeed one can verify that $f^{**} < f^*$ implies that

$$d + \bar{c} > d + \mu \underline{c} + (1 - \mu)(\bar{c} + d),$$

so that visiting a major specialist first is dominated. It only remains to be checked that no deviation in tariffs can be profitable. Notice that after a deviation non-deviating minor specialists can still propose to fix minor losses for $\underline{c} + d$, while major specialists can still propose to fix major losses for $\bar{c} + d$, without having to fear negative profits. Hence, to be profitable the deviating expert must attract the customer when he is uninformed, i.e., at his first visit. We refer to Lemma 1 to establish the following. Case a of Lemma 1 cannot be strictly profitable since the deviating expert would have to charge $p > d$, and then the customer would prefer to visit a minor specialist first. Case b is irrelevant because $f < f^*$ implies $f < \bar{c} - \underline{c}$. Case c is associated with zero profits and cannot be profitable. Case d contradicts $f < f^*$. Finally case e is feasible only if $f = f^{**}$ because $f \geq f^{**}$. But then the deviating expert's profits are zero, since by definition f^{**} is such that $C + \mu f = \frac{d}{\mu} + \underline{c}$.

In the remainder of the proof we establish necessary conditions for the existence of specialization equilibria. Thus suppose that such an equilibrium exists, and denote by $(p_0, \underline{p}_0, \bar{p}_0)$ and $(p, \underline{p}, \bar{p})$ the tariffs played by minor and major specialists, respectively. We now show that we can apply Lemma 2 to this equilibrium, if we moreover assume $d > 0$.²⁸ Indeed at equilibrium major specialists repair major losses (so that $\bar{p} \geq \bar{c}$) and make non-negative profits (so that $p + \bar{p} \geq \bar{c} + d > \bar{c}$ from $d > 0$). Moreover minor

²⁸The case $d = 0$ is omitted for brevity. It is clear that if $d = 0$ information can be obtained at no cost, by visiting a minor specialist (as defined in Assumption 2) first and rejecting all recommendations. The game then degenerates into a simple Bertrand game.

specialists do not fix major losses: they either reject them ($\bar{p}_0 < \bar{c}$), or the consumer rejects their offer of a major treatment ($\bar{p}_0 > p + \bar{p}$). When a major specialist deviates as in Lemma 2, the best offer for fixing major losses is reduced to $p_1 + \bar{c} < p + \bar{p}$; as a result minor specialists still do not fix major losses. At best they fix minor losses, and make zero profits under Assumption 2; but this is exactly what the minor specialist postulated in Assumption 2 does. As a result one can ignore these experts, and we can therefore apply Lemma 2 as if all players were major specialists with the same tariff ($p, \underline{p}, \bar{p}$).

We now show that profits are zero in a specialization equilibrium. Notice that under Assumption 2 minor specialists make zero profits. Suppose that the major specialists' total profits $(1 - \mu)(p + \bar{p} - \bar{c} - d) \equiv (1 - \mu)B$ were positive. Consider a deviation $(p_1, \bar{c} - f, \bar{c})$ and choose $p_1 = p + \bar{p} - \bar{c} - \varepsilon$; this satisfies the assumptions in Lemma 2. Notice that $B_M = p_1 - d = p + \bar{p} - \bar{c} - \varepsilon - d = B - \varepsilon$, which is positive. Applying items (ii) and (iii) of Lemma 2, we get that the profits of the deviating expert (call him expert 1, as in the Lemma) are at least $(1 - \mu)B_M = (1 - \mu)B - (1 - \mu)\varepsilon$. If expert 1 did not deviate he would have to share $(1 - \mu)B$ with other major specialists, implying that the deviation is profitable.

Next, we show that a specialization equilibrium cannot exist if $f \geq f^*$. Suppose that it did exist. Let us apply Lemma 2 to the deviation $(p_1 = d - \mu(\bar{c} - \underline{c} - f) + \varepsilon, \bar{c} - f, \bar{c})$. Since $f \geq f^*$ the recommendation of a minor treatment is accepted. Therefore expert 1 is truthful, and the recommendation of a major treatment is also accepted. Finally it is easily verified that the consumer prefers to visit this efficient expert first. Item (iv) then implies that the deviation is profitable, thus yielding a contradiction.

We finally show that a specialization equilibrium cannot exist if $f < \min\{f^*, f^{**}\}$. Let us apply Lemma 2 to the deviation $(p_1 = d - \mu(\bar{c} - \underline{c} - f) + \varepsilon, \bar{c} - f, \bar{c})$. Because $f < f^*$, the recommendation of a minor treatment by expert 1 would be rejected; therefore, conditional on being chosen by a consumer, expert 1 would always recommend a major treatment. Moreover, Lemma 2 implies that expert 1 must repair major losses. From

the end of the proof of Lemma 1 we also know that $f < f^{**}$ is the condition ensuring that the consumer prefers to visit a fraudulent expert first, rather than visiting him after having had minor losses repaired by another expert. We can therefore conclude that expert 1 repairs all losses. Item (iv) in Lemma 2 then implies that the deviation is profitable . ■

Proof of Proposition 3. We first prove that there exists a fraud equilibrium if $f \leq f^{**}$ and $f < f^*$. Suppose that all experts announce the tariff $(d - \mu(\bar{c} - \underline{c} - f), \underline{c}, \bar{c})$, that they always claim that a major treatment is needed, and that the customer accepts the recommendation of a major treatment. Clearly these strategies form a PBE of the continuation subgame following the announcements of tariffs, because $f \leq f^{**}$ ensures that the customer does not gain by visiting a minor specialist first. It remains to be checked that there is no profitable deviation. Referring to Lemma 1: case a cannot be strictly profitable since the deviating expert would have to charge $p > d$, and then the customer would prefer to visit a minor specialist first; case b is irrelevant because $f < f^*$ implies $f < \bar{c} - \underline{c}$; case c is associated with zero profits and cannot be profitable; case d contradicts $f < f^*$; and case e cannot profitably attract customers.

We now check that profits must be zero in a fraud equilibrium when $f \leq f^{**}$ and $f < f^*$. So suppose to the contrary that there exists a fraud equilibrium in which all experts set a tariff $(p, \underline{p}, \bar{p})$ such that $p + \bar{p} > C + \mu f$. Recall from Lemma 1 that $p + \bar{p} \in [C + \mu f, \frac{d}{\mu} + \underline{c}]$. If $f^{**} = 0$, then $C + \mu f = d/\mu + \underline{c} = \bar{c}$, which yields a contradiction because then $p + \bar{p} = C + \mu f$. If $f^{**} > 0$, we have $C + \mu f > \bar{c}$; but then we can once more apply Lemma 2 to $p_1 = p + \bar{p} - \bar{c} - \varepsilon$. Because $f < f^*$, the consumer would reject the recommendation of a minor treatment; expert 1 therefore always recommends a major treatment and the consumer accepts. Moreover, from item (i) in Lemma 2 expert 1 must repair major losses. Hence, and given Assumption 2, all other experts can at best repair minor losses for a price equal to $d + \underline{c}$. Recall that $f < f^{**}$ is the condition ensuring that the consumer prefers to visit a fraudulent expert first, rather

than visiting him after having made minor losses repaired by another expert. Therefore the customer prefers visiting expert 1 first, and expert 1 repairs all losses. Item (iv) then applies; the profit of expert 1 is $p_1 - d + \mu(\bar{c} - \underline{c} - f) = [p + \bar{p} - (C + \mu f)] - \varepsilon$, which is to be compared to a share of the total profits $[p + \bar{p} - (C + \mu f)]$. As a result the deviation is profitable.

Finally, we prove that there exists no fraud equilibrium if $f > f^{**}$, or if $f \geq f^*$. First, if a fraud equilibrium exists then case e in Lemma 1 implies that $f \leq f^{**}$. Now consider the case $f \geq f^*$. Then one can apply Lemma 2 to $p_1 = d - \mu(\bar{c} - \underline{c} - f) + \varepsilon$. Our tie-breaking rules would imply that the deviating expert is truthful. The deviation would attract consumers at their first visit, and the deviation would therefore be profitable. ■

Proof of Proposition 4. From Lemma 1 it only remains to be shown that case b cannot be part of an equilibrium. In case b the consumer visits an expert with a tariff $(p, \underline{p}, \bar{c})$ (so that $p \geq d$ to avoid negative profits) who repairs only major losses. If the consumer is rejected he visits a minor specialist, who under Assumption 2 must make zero profits. Because $f = \bar{c} - \underline{c}$, we can apply Lemma 2 to $p_1 = d - \mu(\bar{c} - \underline{c} - f) + \varepsilon$. Once more expert 1 is truthful from our tie-breaking rules, he attracts consumers at their first visit because he is efficient, and thus increases his profits. ■

Proof of Proposition 6. Let us consider a candidate efficient equilibrium without any minor specialists. Suppose that all experts set the tariff $(d - \mu(\bar{c} - \underline{c} - f), \bar{c} - f, \bar{c})$ and behave truthfully, so that the consumer ends up paying the efficient amount C . In the case $f \geq f^*$, we can proceed exactly as in the proof of Proposition 1: a deviation can be profitable only if it proposes to fix minor losses for less than $\bar{c} - f$, but then the expert cannot be truthful and the deviation would not be profitable. Assume now that $f < f^*$. Then experts offer the minor treatment at a price $\underline{p} > \underline{c} + d$. Consider an expert who deviates towards a minor specialist tariff $(d + \varepsilon, \underline{c}, \underline{c})$. Then a non-

deviating expert would anticipate that his offer of a minor treatment would be rejected by the customer. So, if a non-deviating expert still accepts customers, then he must always recommend the major treatment. But then non-deviating experts are in effect using the tariff of the most efficient fraud equilibrium (see the beginning of the proof to Proposition 3). If $f > f^{**}$, the customer prefers to first visit the minor specialist, so that the non-deviating experts would attract only major losses. Whether or not the non-deviating experts accept consumers at all, the deviation would be profitable by fixing the customers' minor losses. By contrast, if $f \leq f^{**}$ the consumer would prefer to first visit a non-deviating expert, even though this would imply an additional expected cost of μf . In this case the deviation is not profitable. ■

References

- DARBY, M. and E. KARNI (1973). "Free Competition and the Optimal Amount of Fraud," *Journal of Law and Economics*, 14:67-88.
- DE JAEGHER, K. AND M. JEGERS (2001). "The Physician-Patient Relationship as a Game of Strategic Information Transmission," *Health Economics*, 10:651-668.
- DIAMOND, P.A. (1971). "A Model of Price Adjustment," *Journal of Economic Theory*, 3:156-68.
- DRANOVE, D. (1988). "Demand Inducement and the Physician/Patient Relationship," *Economic Inquiry*, 26:281-298.
- DULLECK, U. and R. KERSCHBAMER (2006). "On Doctors, Mechanics and Computer Specialists: The Economics of Credence Goods" *Journal of Economic Literature*, 44:5-42.
- EMONS, W. (1997). "Credence Goods and Fraudulent Experts," *RAND Journal of Economics*, 28:107-119.
- EMONS, W. (2001). "Credence Goods Monopolists," *International Journal of Industrial Organization*, 19:375-389.
- FARLEY, P.J. (1986). "Theories of the Price and Quantity of Physician Services: A Synthesis and Critique," *Journal of Health Economics*, 5:315-333.
- FONG, Y.-F. (2005). "When Do Experts Cheat and Whom Do They Target?" *RAND Journal of Economics*, 36:113-130.
- LAL, R. and C. MATUTES (1994). "Retail Pricing and Advertising Strategies" *Journal of Business*, 67:345-370.
- MCGUIRE, T.G. (2000). "Physician Agency," in Culyer, A.J. and J.P. Newhouse, eds., *Handbook of Health Economics*, Amsterdam, New York and Oxford: Elsevier Science, North-Holland.
- PESENDORFER, W. and A. WOLINSKY (2003). "Second Opinions and Price Competition: Inefficiency in the Market for Expert Advice," *Review of Economic Studies*,

70:417-37.

PITCHIK, C. and A. SCHOTTER (1987). "Honesty in a Model of Strategic Information," *American Economic Review*, 77:1032-1036.

RILEY, J. (1979). "Informational Equilibrium," *Econometrica*, 47:331-59.

ROTHSCHILD, M. and J. STIGLITZ (1976). "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, 90:629-649.

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

TIROLE, J. (1988). *A Theory of Industrial Organization*, Cambridge MA: MIT Press.

WOLINSKY, A. (1993). "Competition in a Market for Informed Experts' Services," *RAND Journal of Economics*, 24:380-398.




Figure 1.




Figure 2. Welfare as a function of the diagnosis cost d .




Figure 3. Welfare as a function of the fraud cost f .