

Economics Department
Working Papers in Economics

Boston College

Year 2004

A Theory of Fraud and
Over-Consumption in Experts Markets

Ingela Alger
Boston College,

Francois Salanie
INRA-LEERNA, Toulouse,

A Theory of Fraud and Over-Consumption in Experts Markets

Ingela Alger* and François Salanié**

July 2004

Abstract

Consumers often have to 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-consumption. Empirical evidence suggests that over-consumption is a pervasive phenomenon in experts markets. We offer and discuss conditions leading to equilibrium over-consumption in an otherwise purely competitive model. This market failure results from the freedom of consumers 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. Our model also yields predictions on the diagnosis price that are in line with stylized facts.

Keywords: experts, fraud, over-consumption.

JEL: D41, D82, L11.

*Boston College. ingela.alger@bc.edu

**LEERNA, INRA and Université de Toulouse. salanie@toulouse.inra.fr

Acknowledgement: We wish to thank Hideo Konishi, Levis Koshin, Ching-to Albert Ma, Chris Snyder, and several anonymous referees 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 also thankful to several anonymous referees who have provided 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. Although empirical research is relatively scarce for obvious reasons, there is evidence giving support to this concern. 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-consumption in experts markets. Is over-consumption 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-consumption has been particularly debated in health economics. In that literature,² the physician's market power over the consumer is a central factor behind the over-consumption 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-consumption.

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-consumption. The key problem in experts markets is that the consumer knows there is a loss, but only an expert can determine which treatment is needed.³ Once the diagnosis is made, the expert thus enjoys an informational advantage over the consumer. Moreover, there typically exist some economies of scope between diagnosis and treatment,⁴ making it costly for the consumer to get a second opinion. These characteristics of experts' markets do not create *per se* any incentive to induce over-consumption. They simply make overcharging more likely: once the diagnosis is made, the expert may increase its price without fearing that the customer rejects his offer.⁵

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 he claims to be needed for the repair, and computes the repair price accordingly. In this system, increasing the repair price now requires to justify the use of additional inputs. Still, this does not create over-consumption: if the customer cannot observe whether these inputs are actually used or not, the expert will not have to use these inputs anyway. Consequently over-consumption does not appear; simply the expert has to lie to the customer in order to increase his revenues.

For over-consumption to appear, the customer must be able to verify whether some of these unnecessary inputs were actually used. In that case, the expert has to incur

³Once the loss is fixed, evidence in favor of one treatment or another has disappeared. Taken together, these features define a credence good, a term coined by Darby and Karni (1973).

⁴These 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.

⁵In the literature this hold-up problem is associated with the work of Diamond (1971), who considers exogeneous switching costs from one seller to another; it may be so severe to make the monopoly price prevail as the unique equilibrium price.

an additional “fraud cost” when he lies, equal to the cost of unnecessary inputs used in the repair. The trade-off is now clear: lying allows to increase the price, but implies a simultaneous increase in the repair cost. Thus an interesting effect appears: avoiding fraud and over-consumption requires to make repair prices closer (in the limit, if the price does not depend on the type of repair, there is no fraud incentive); in other words, cross-subsidies between repairs are needed. However, competition makes these cross-subsidies less sustainable, through the threat of cream-skimming. It may thus well be that competition favors fraud.

In order to test these intuitions, we set up a model with the above key ingredients: informational advantage for the expert, economies of scope between diagnosis and repair, commitment to prices, verifiability of some inputs by the customer (fraud cost). The model incorporates an optimal visit pattern by the customer among experts, together with Bayesian updating of beliefs when an expert emits a recommendation. Experts compete in tariffs, which include repair prices and a price for the diagnosis. As will be discussed further below, a crucial assumption is that experts are not allowed to set repair prices below repair costs. Three different types of equilibria are characterized, depending on the parameters.

When the sum of the diagnosis cost and the fraud cost is sufficiently high, we show that there exists an efficient equilibrium in which the loss is fixed by the first visited expert, and this expert is truthful. Each repair price is set close enough to marginal cost to deter cream-skimming, without inducing fraud. Otherwise, inefficiencies arise. If the fraud cost is high compared to economies of scope, there exists an equilibrium without fraud in which the first visited expert is truthful, because the customer rejects his offer of an expensive repair to get the loss fixed by another expert. This “specialization” equilibrium, first exhibited by Wolinsky (1993), is inefficient because the diagnosis cost is sometimes incurred twice.

Finally, when the fraud cost is small relative to the economies of scope, all equilibria

involve fraud and over-consumption. Fraud means that the expert always claims that a costly repair is needed. The consumer does not learn anything, and therefore accepts to pay a high repair price, despite knowing that with some probability unnecessary inputs are used. Our theory thus relates fraud and over-consumption to information transmission and cream-skimming: the expert pools information in order to deter the consumer from seeking a better price elsewhere. Such a pooling requires to sometimes mimick an expensive repair, and thus creates over-consumption. This new type of equilibrium is our main contribution, and we now contrast this finding with other results in the literature.

Our model is inspired by that of Wolinsky (1993). This is a simple framework in which the loss may require either a minor or a major intervention. Our main innovation is the idea that the customer is able to observe some of the inputs used in the repair. As explained above, this creates a fraud cost, which measures the ability of the customer to monitor the expert.⁶ We allow this fraud cost to take arbitrary values, thus covering a variety of cases.

Previous papers have looked at two extreme cases. First, some have assumed that fraud is costless (Pitchik and Schotter, 1987, Wolinsky, 1993, Taylor, 1995, Dulleck and Kerschbamer, 2003, and Fong, 2003). Then lying is not inefficient *per se*, and the only source of inefficiency is the cost of getting a second opinion,⁷ incurred when the equilibrium involves multiple visits as in the specialization case; otherwise equilibria are efficient. In our model, the fraud cost is a parameter associated to the fact that the customer may verify the use of some inputs. This innovation allows us to properly define

⁶Notice that the fraud cost is also equal to the social cost of over-consumption. 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 at equilibrium. Caution may therefore be called for when estimating the welfare effects of a policy aiming at deterring fraud.

⁷Taylor (1995) adopts a framework where he can also analyze inefficiencies arising in the level of maintenance of the durable good. Emons (1997, 2001) and Fong (2003) also allow for undertreatment.

over-consumption, to confirm its empirical relevance since it appears as an equilibrium phenomenon, and to discuss its welfare effects in a meaningful way since now fraud has a welfare cost. In line with the intuition, it turns out that fraud occurs at equilibrium when the fraud cost is low enough, and the diagnosis cost is high enough to deter the customer from seeking a double advice as in the specialization equilibrium.

Conversely a high fraud cost means that the consumer has full control— it is as if 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, 2003, Emons, 1997 and 2001); they find no fraudulent consumption inducement in equilibrium.

Our model allows for an endogenous diagnosis price, often neglected in the literature.⁸ Because experts are allowed to reduce the diagnosis price below its cost, all equilibria are zero-profits equilibria, contrary to some equilibria in Wolinsky (1993) and Dulleck and Kerschbamer (2003). This is because repair profits may be redistributed to consumers, thanks to a lower diagnosis price. We find that whenever the consumer visits only one expert in equilibrium, the diagnosis price is set below the diagnosis cost, in line with the evidence. The diagnosis price is also shown to increase with the fraud cost. As an application, it seems quite plausible that the consumer typically observes more of the inputs used in the health care industry than for car or house repairs. Our model thus predicts a low diagnosis price for car or house repairs (and indeed diagnoses are often free in these industries), and a significant diagnosis price in health care or dentistry (patients do typically have to pay the physician's fee even if not treated by that physician).

Our equilibrium with over-consumption may also be usefully compared to those obtained in the health economics literature. This literature typically features a mo-

⁸Pitchik and Schotter (1987) take all prices as given. In Wolinsky (1993), Dulleck and Kerschbamer (2003), and Fong (2003), the diagnosis price is exogenously set at the diagnosis cost. In Taylor (1995), one of the intervention prices is exogenous; the precise level of the diagnosis price is indeterminate because only the sum of the repair price and the diagnosis price matters.

nopolistic competition setup. No strategic interaction occurs between physicians. Each physician selects both price and quantity of treatment, taken the reservation utility of the consumer (defined by other physicians' prices and quantities) as given. Overconsumption occurs in the sense that the patient would like to consume less treatment, given the unit price; this is reminiscent of the second-degree price discrimination literature (see Farley, 1986). Some papers have subsequently incorporated asymmetric 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 more easily a treatment offer, 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 exogeneous.

In the next section we present the setup. We then characterize pure-strategy equilibria, before turning to a welfare analysis. Before concluding, we devote a section to a discussion of the critical role played by several assumptions.

2 The Model

Our model features two classes of risk-neutral agents: consumers and experts. There is a continuum with mass one of consumers. Each consumer incurs a loss, which must be repaired; this last assumption will be relaxed in Section 4. The loss should be interpreted as a symptom: the car does not work properly, the house roof leaks, or a tooth aches. For each consumer a minor intervention is sufficient to get the loss fixed (state m) with probability $\mu \in (0, 1)$; otherwise, a major intervention is necessary (state M).⁹

⁹Using the vocabulary introduced by Dulleck and Kerschbamer (2003), consumers are homogenous. They also study the case where consumers differ in the probability distribution

The model is mathematically equivalent to one where there is only one consumer. For presentational simplicity, we will therefore often refer to *the* consumer. The consumer cannot distinguish between the two states; only experts can.

There are $n \geq 4$ identical experts, each of whom is randomly assigned an index $i = 1, \dots, n$. An expert may observe the consumer's state¹⁰ at a diagnosis cost $d \geq 0$, which must be understood as a diagnosis-and-switching cost. Including a switching cost incurred each time a consumer visits an expert, would in fact be equivalent to raising the diagnosis cost by that amount, and would not affect the results. The marginal cost of an intervention is assumed to be constant and equal to \underline{c} for a minor intervention and $\bar{c} > \underline{c}$ for a major one. For further use, we define C as the minimum expected cost of getting the loss fixed by an expert:

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

By assumption an expert is needed both at the diagnosis stage and at the repair stage. Moreover 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 .

Once the diagnosis is made, the expert recommends an intervention. We assume that whether the loss has been fixed or not is verifiable information; thus if an expert and a consumer agree that the expert should fix the loss, he does fix it.¹¹ However, since the consumer may not distinguish between the two states, recommendations may be false. This is the fraud issue we focus on: what does the consumer end up paying for getting the loss fixed, and what are the costs incurred by the expert?

over the two types of intervention. In Fong (2003), consumers may differ in the loss they incur if left untreated, or in the cost for the expert to provide the treatment; the monopolistic expert then uses fraud in the form of over-charging as a substitute for price discrimination.

¹⁰There is no moral hazard in the diagnosis itself. For an analysis of possible inefficiencies in the amount of effort provided at the diagnosis stage, see Pesendorfer and Wolinsky (2003) and Emons (2001). Wolinsky (1993) studies the possibility of diagnosis errors.

¹¹Using the vocabulary introduced by Dulleck and Kerschbamer (2003), we impose “liability”, as does Wolinsky (1993). By contrast Emons (1997, 2001) assumes that the consumer does not observe whether the loss has been fixed, which allows him to study fraud in the form of under-treatment.

Now we introduce the main innovation of our model. An expert usually provides a bill listing the inputs used, and we will assume that here. As a result, if the consumer observes that some of the inputs were effectively used, making a false recommendation is costly. For instance, suppose the roof leaks and the rooftop needs to be replaced (the minor intervention). However, the roofer (falsely) claims that the rafters are damaged and need to be changed as well (the major intervention). The consumer may easily verify *ex post* whether the rafters were replaced. This creates a positive fraud cost f for the expert, incurred when he recommends a major treatment when a minor one would have been sufficient. In this example, f is the cost of replacing rafters which were not damaged, and must be added to the cost of the minor intervention.¹² One may argue that $\underline{c} + f$ should typically lie below \bar{c} : this is the case for example if replacing the rafters is less effort-consuming than if they were really damaged, and effort is not observed by the consumer. It could also be higher, if for example the expert has to damage the rafters himself before replacing them. To keep some generality we simply assume $f \geq 0$; thus, we allow for $f = 0$ and $f = \bar{c} - \underline{c}$, which are the two special cases treated previously in the literature. *Ceteris paribus*, f is larger the more inputs the consumer observes. Finally it is typically harder to claim that the intervention was minor when in fact it was major, than the reverse. For instance, how could the roofer justify changing the rafters if he claimed they did not need to be changed? For simplicity, we assume that an expert may only make a false recommendation when the consumer needs a minor intervention; we will comment later on how allowing for false recommendations in the other direction would affect the results.

We avoid the hold-up problem referred to in the introduction by assuming that the tariffs posted by experts are public information, and that an expert can perfectly commit to his tariff.¹³ The tariff of expert i comprises a diagnosis price p_i , and prices \underline{p}_i and

¹²Alternatively one could justify the introduction of a positive fraud cost by the expected penalty associated with being caught by an external auditor, or by moral considerations.

¹³This is equivalent to assuming that experts may commit to input prices, and consumers know what inputs are needed for each type of treatment. Repair shops often post their input prices in a place visible by the customers. In some countries this is compulsory.

\bar{p}_i for the minor and the major treatment, respectively. Expert i 's recommendation if a consumer is in state m is denoted r_i , where $r_i \in \{m, M\}$. In what follows, we assume that “doing nothing” involves no cost.

The game has $n + 1$ stages. In Stage 0, all the experts simultaneously post their tariffs. In Stage j , $j = 1, \dots, n$, the consumer chooses between visiting some expert $i \in \{1, \dots, n\}$, and doing nothing. The latter action puts an end to Stage i . If the consumer chooses to visit expert i , he pays the diagnosis price p_i . The expert observes the state of nature at cost d , and then makes a recommendation: if the consumer's state is m , the expert recommends either a minor or a major intervention, $r_i \in \{m, M\}$, and if it is M he recommends a major intervention. The consumer either accepts or refuses the offer. If the consumer accepts, the expert fixes the loss and the consumer pays the corresponding price (\underline{p}_i if the expert recommended m , and \bar{p}_i if the expert recommended M). If the consumer refuses, the expert does nothing.

It is realistic to think that an expert may invoke some stratagem (for example unexpected delays or workload) to avoid serving a customer. To capture this idea, we impose that experts cannot commit to making a repair at a loss, as in Wolinsky (1993):

$$(2) \quad \underline{p}_i \geq \underline{c} \quad \forall i$$

$$(3) \quad \bar{p}_i \geq \bar{c} \quad \forall i.$$

In Section 4 we discuss how altering our assumptions would affect the results. It turns out that this assumption in particular is important.

By contrast, an expert is able to commit to a diagnosis price below marginal cost—that may simply be a tool to attract a consumer in the first place. We will of course impose that experts make non-negative expected profits. Because consumers may opportunistically visit an expert to collect a negative diagnosis price, we impose:

$$(4) \quad p_i \geq 0 \quad \forall i.$$

The consumer's overall 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 set of available tariffs and the recommendations that he may have received in the past. Expert i 's strategy includes a tariff $(p_i, \underline{p}_i, \bar{p}_i)$ and a recommendation strategy r_i for any consumer in state m . Given these strategies, one can compute the consumer's payoff as the sum of all prices paid to the experts he visits (only the diagnosis price if he turns down the expert's recommendation, but both the diagnosis price and the intervention price if he accepted the recommendation). This allows to compute an expected payoff, given the initial beliefs μ and the usual Bayesian updating rule, to be applied whenever the consumer receives a recommendation; updated beliefs given any recommendations received in stages up to and including Stage j are denoted $\hat{\mu}_j$. 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 some consumers may visit several experts.

We characterize Perfect Bayesian Equilibria (PBE) of this game. We restrict attention to pure strategy equilibria. As a tie-breaking rule, we assume that if an expert is indifferent between recommending m or M to a consumer in state m , he recommends m ; furthermore, if a consumer is indifferent between accepting and rejecting a recommendation, he accepts it.

3 Equilibria

3.1 Equilibrium recommendation and acceptance strategies

One of the difficulties with this model is that the number of stages where a consumer chooses to visit an expert is endogenous. Consider what may happen at the consumer's first visit to an expert. Then, either the expert recommends truthfully, in which case the consumer learns his true state, or he always recommends a major intervention, in which case the consumer learns nothing. However, without any further knowledge about the set of available tariffs, it is not obvious how the consumer should respond to the first

visited expert's recommendation in these cases. As mentioned previously, our model is fairly close to Wolinsky (1993). There are essentially two differences: we allow for a positive fraud cost and an endogenous diagnosis price. He finds that equilibrium may be of two types. Either it is efficient, i.e., the first visited expert makes a truthful recommendation that the consumer accepts. Or it involves inefficient specialization: the first visited expert makes a truthful recommendation, but the consumer accepts only if a minor intervention was recommended; if a major intervention was recommended, the consumer visits a second expert who fixes the loss. 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: the first visited expert always recommends a major intervention, and the consumer accepts. Finally, we show that these are the only three relevant types of equilibria.

As a preliminary step and benchmark, assume that each consumer knows which intervention he needs. A standard Bertrand competition argument shows that equilibrium then must be efficient. Assume to the contrary that in equilibrium either some consumers visited more than one expert, or costly over-consumption took place. To get the loss fixed, some consumers would therefore pay more than the sum of the diagnosis cost and the cost of the appropriate intervention. For instance, assume that consumers in state m paid more than $d + \underline{c}$. Then some expert could profit by offering a tariff $(d, \underline{c} + \varepsilon, +\infty)$, where $\varepsilon > 0$ would be chosen so as to attract all consumers in state m . The expert has no incentive to defraud the consumer with such a tariff.

Returning now to the case where consumers have initial beliefs μ , we derive sufficient and necessary conditions for efficient equilibria to exist. Formally, we define an "efficient equilibrium" as being one where the first visited expert makes a truthful recommendation, that the consumer accepts. Without loss of generality, we may assume that the expert first visited in equilibrium has index 1. For an efficient equilibrium to exist, the tariff $(p_1, \underline{p}_1, \bar{p}_1)$ must satisfy a number of constraints, which will turn out to imply a specific condition on the parameter values. This condition will turn out to also

be sufficient for efficient equilibria to exist.

When a consumer needs a minor intervention, expert 1 chooses between recommending a minor intervention, for a profit $\underline{p}_1 - \underline{c}$, and a major intervention, which yields a profit $\bar{p}_1 - \underline{c} - f$. For a truthful recommendation to be a best response to the consumer's acceptance of both recommendations, one must then have

$$(5) \quad \underline{p}_1 \geq \bar{p}_1 - f.$$

Next we check that accepting both recommendations is the consumer's best response to the expert's truthful recommendation. Given that expert 1 recommends truthfully, the consumer updates his beliefs to $\hat{\mu}_1 = 1$ upon being recommended a minor intervention. Now any expert would be willing to fix the loss of this consumer if he were paid at least $d + \underline{c}$. Therefore, unless

$$(6) \quad \underline{p}_1 \leq \underline{c} + d,$$

the consumer's acceptance of a minor recommendation by expert 1 cannot form part of an equilibrium. Indeed, some expert $i \geq 2$ may profit by offering a tariff, e.g., $(0, d + \underline{c} + \varepsilon, +\infty)$ for some $\varepsilon > 0$, that would attract a consumer who has been recommended a minor intervention by expert 1. Similarly, for the consumer to accept the first expert's recommendation to do a major intervention, we must have:

$$(7) \quad \bar{p}_1 \leq \bar{c} + d.$$

Taken together, constraints (5), (6), and $\bar{p}_1 \geq \bar{c}$ imply

$$(8) \quad f \geq \bar{c} - \underline{c} - d \equiv f^*.$$

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

Proposition 1 *There exists an efficient equilibrium, in which the first expert makes a truthful recommendation, that the consumer accepts, if and only if $f \geq f^*$. In any such equilibrium, experts make zero expected profits.*

This result highlights a fundamental tension between, on the one hand, the expert’s incentive to recommend a minor intervention truthfully, and on the other hand, the consumer’s incentive to accept the first expert’s recommendation to do a minor intervention. The former incentive requires the price of a minor intervention to be sufficiently large: in fact, this price cannot differ from the price of a major intervention by an amount larger than the fraud cost. The latter incentive puts an upper limit on the price of a minor intervention: it cannot exceed the cost of a minor intervention by more than the diagnosis cost, or else the consumer could obtain the minor intervention at a lower price elsewhere. Therefore, if the sum of the fraud and the diagnosis cost is too small, or, equivalently, if $f < f^*$, the two requirements are incompatible, and inefficiencies must arise.

We now investigate whether the equilibrium outcome may be that of full specialization, in which some experts perform only minor interventions, whereas others do only major ones. In a “specialization equilibrium,” the consumer finds out his true state by first visiting an expert specializing in minor interventions; any consumer needing a major intervention then visits a major specialist. The minor specialist removes any incentive to make a false recommendation by charging a price for the major intervention which is sufficiently high to ensure that the consumer would turn down a recommendation to do a major intervention.¹⁴

Proposition 2 *There exists a specialization equilibrium if and only if $f < f^*$ and $f \geq \frac{1-\mu}{\mu^2}[d - \mu(\bar{c} - \underline{c})] \equiv f^{**}$, or if $d = 0$. In such an equilibrium, the first expert makes a truthful recommendation; the consumer accepts only if a minor intervention was recommended. Otherwise, he visits a second expert, whose recommendation to do a major intervention he accepts. Experts make zero expected profits.*

As long as there are some economies of scope ($d > 0$), specialization involves a costly sec-

¹⁴A minor specialist and a consumer in state M would always have an incentive to renegotiate, so as to avoid a costly second visit to an expert. But of course, if such a renegotiation were allowed and anticipated, minor specialists would always recommend M . Prohibiting renegotiation, as we do, implies that the threat of cream-skimming is stronger, and thus makes the model more competitive.

ond expert visit for a consumer needing a major intervention. As a result, if the diagnosis cost and the probability of a major intervention are large, the consumer would prefer to get the loss fixed by one single expert, even if that would involve over-consumption: a relevant threat to specialization is an expert who always recommends the major intervention that the consumer accepts. There is therefore a fundamental trade-off between, on the one hand, costly double advice, and on the other hand, costly over-consumption. For this reason, over-consumption appears as an equilibrium phenomenon when the fraud cost is not too large, as will be shown next. In a “fraud equilibrium” the first visited expert always recommends a major intervention, which the consumer accepts.

For the purposes of the following proposition, we assume that at least one expert offers a minor specialist tariff in equilibrium, although he is not visited.¹⁵ We argue that this assumption is weak. If a fraction of customers were initially informed that their loss is minor, then active minor specialists would arise endogenously. This variation of the model would not alter our results in any respect. There are in fact many examples of minor interventions that consumers may demand on a regular basis for the sake of maintenance; for cars, such repairs include oil and brake changes. Also, since we know that at equilibrium some experts may not be visited, some of them may as well choose to be inactive minor specialists. Moreover, a minor specialist would not make losses if by accident a consumer decided to visit him.

Proposition 3 *There exists a fraud equilibrium involving fraudulent recommendations*

¹⁵This assumption ensures the existence of an equilibrium for $f < f^*$ and $f \leq f^{**}$. The fact that an equilibrium may not exist in a game with incomplete information is by itself not surprising. The archetypal example is the Rothschild and Stiglitz (1976) insurance model, in which, for some parameter values, pooling candidates are killed by an entry with separating contracts, and candidates with separating contracts are killed by a pooling entry. Here we have an even more complex, cyclical structure if there were no inactive minor specialists: efficient candidates are killed by a minor specialist tariff, specialization is killed by a tariff inducing fraud, and candidates with fraud are killed by the a tariff inducing a truthful recommendation which is accepted by the consumer. Interestingly it is because this cycle now counts three stages that an existence result may be obtained by introducing inactive agents. An alternative way to ensure equilibrium existence would be to allow for a second entrant, following Riley (1979).

and over-consumption, if and only if $f < f^*$ and $f \leq f^{**}$. Experts make zero expected profits.

This result shows that even in a purely competitive model, fraud and over-consumption may appear as an equilibrium phenomenon. The reason is that experts would lose the profitable customers if they offered both interventions at reasonable prices and recommended interventions honestly. Fraud is a simple way to avoid 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. When this expected cost is too large, another inefficiency arises in the form of specialization, whereby a costly second visit is made in state M .¹⁶

So far, we have only determined necessary and sufficient conditions for efficient, specialization, and fraud equilibria to exist. The next proposition states that, except in the degenerate case $d = 0$, these are the only type of equilibria that may exist, implying that fraud and over-consumption *must* arise as equilibrium phenomena for the relevant parameter values.

Proposition 4 *For any positive diagnosis cost $d > 0$, equilibrium is either an efficient, a specialization, or a fraud equilibrium.*

When the diagnosis cost is nil, finding out the true state costs nothing, and specialization in the diagnosis may therefore be an equilibrium. However, as soon as $d > 0$ so that there are some economies of scope between the diagnosis and the repair, such a specialization in the diagnosis is dominated by the specialization equilibrium where the minor specialist not only reveals the true state to the consumer but also sometimes fixes the loss.

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 $f = 0$; as in our model, he finds that there then exists two regimes. If the diagnosis cost d is small, there

¹⁶Recall that we have assumed that an expert may only falsely recommend a major intervention. Relaxing this assumption would simply be that equilibria with fraud would involve experts always recommending the minor treatment whenever low-cost repairs are more likely than high-cost repairs.

is specialization; if it is large, equilibria are efficient. Indeed, when $f = 0$, fraudulent recommendations are costless and therefore involve no inefficiency.¹⁷ At the other end of the spectrum, Dulleck and Kerschbamer (2003) consider the case where $f = \bar{c} - \underline{c}$; as in our model, they find that equilibria are then efficient.

[FIGURE 1 ABOUT HERE]

3.2 Equilibrium tariffs

In the following discussion, by equilibrium tariff we mean a tariff which is offered by some expert who is visited with positive probability in equilibrium.

Although there always exists a unique equilibrium visit and recommendation pattern (except for $d = 0$), equilibrium tariffs are typically not unique. For instance, if $f \geq \bar{c} - \underline{c}$, marginal cost pricing $(d, \underline{c}, \bar{c})$ is always an equilibrium tariff; but another equilibrium tariff would offer the diagnosis for free: $(0, \underline{c} + d, \bar{c} + d)$. In some experts markets, consumers may indeed get a diagnosis free of charge. This is typically the case for car repairs. However, it is not true in the health care industry. Here we try to shed some light on this issue, by analyzing whether, and if so, why, we should expect a low diagnosis price in a systematic manner. In the following proposition, we determine the highest possible diagnosis price for each type of equilibrium.

Proposition 5 *If the consumer visits only one expert in equilibrium, and if $f < \bar{c} - \underline{c}$, the largest possible equilibrium diagnosis price \hat{p} is strictly smaller than the diagnosis cost: $\hat{p} = d - \mu(\bar{c} - \underline{c} - f) < d$. Otherwise, marginal cost pricing may occur in equilibrium.*

Proof: From the proof of Proposition 2, we know that the diagnosis and the interventions may be prices at marginal cost in a specialization (in fact, the unique equilibrium minor specialist tariff sells both the diagnosis and the minor intervention at marginal

¹⁷The threshold value for d below which specialization occurs in our model is $(\bar{c} - \underline{c})\mu$; this is smaller than the one obtained by Wolinsky, namely, $(\bar{c} - \underline{c})\mu/(1 - \mu)$, because Wolinsky does not endogenize the diagnosis price. Then in the equilibrium where consumers visit one single expert, the experts make strictly positive profits, which in turn renders specialization more appealing.

cost). In a fraud equilibrium, the largest diagnosis price is obtained when the price of a major intervention is minimized: $\bar{p} = \bar{c}$. Then, for expected profits to be nil, we get $p = d - \mu(\bar{c} - \underline{c} - f)$. Finally, in an efficient equilibrium, marginal cost pricing is an equilibrium tariff if $f \geq \bar{c} - \underline{c}$. Otherwise, the largest diagnosis price is obtained by minimizing the prices of both interventions: $\bar{p} = \bar{c}$ and $\underline{p} = \bar{c} - f$. Again, for expected profits to be nil, we get $p = d - \mu(\bar{c} - \underline{c} - f)$. ■

We find that at any equilibrium where the consumer visits only one expert, and provided it is not more costly to falsely claim that a major intervention was performed than actually performing it ($f < \bar{c} - \underline{c}$), the diagnosis price must be below marginal cost. The reason is that consumers needing a minor intervention then pay a price above marginal cost for the repair. When equilibrium is efficient, this occurs because the price of the minor intervention must be sufficiently large to deter the expert from making a fraudulent recommendation. When equilibrium involves fraud, the consumer pays at least the cost of a major intervention to get a minor intervention. In both cases experts make profits on the minor intervention; competition will therefore lead experts to incur a loss at the diagnosis stage which reimburses the full amount of the profits made at the treatment stage.

To the best of our knowledge, our model is the first to predict that the diagnosis price *must* sometimes be strictly below marginal cost.¹⁸ The rationale is clear: either fraud or the threat of cream-skimming implies that an expert makes a profit at the treatment stage; competition causes experts to use the diagnosis price to reimburse those excess profits to consumers.

Our results also show that the largest possible diagnosis price is increasing in the fraud cost f . The smaller is f , the larger is the profit made on a minor intervention, and thus the smaller is the diagnosis price. One may argue that this may explain why car

¹⁸Taylor (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 set at marginal 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.

and home repair services typically go together with low diagnosis prices, whereas health care providers in 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 believe that the opposite may be true.

3.3 Welfare

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

First, it is obvious that for a given number of experts visited in equilibrium, welfare is decreasing in the diagnosis cost d . However, non-trivial changes may occur when switching from one type of equilibrium to another. It is straightforward to verify that welfare increases in a discrete manner when leaving the specialization regime to enter 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 discrete manner when reaching the efficient regime. Note however that it would be better to have a very small diagnosis cost and specialization, than efficiency together with a larger diagnosis cost. See Figure 2. In the figure, the values d^* and d^{**} correspond to the threshold values for d , for a given value of the fraud cost f .

[FIGURE 2 ABOUT HERE]

For a given diagnosis cost, the maximum welfare is attained for a value of the fraud cost f sufficiently large to enable experts to truthfully reveal the state without inviting cream-skimming. But this maximum welfare is also obtained if equilibrium involves costless fraudulent recommendations ($f = 0$). For values in between, equilibrium is inefficient, and welfare decreases in the fraud cost f within the fraud regime. Interestingly, increasing the fraud cost so as to eliminate fraud worsens welfare if it leads to

specialization instead.

[FIGURE 3 ABOUT HERE]

Casual observations suggests that fraud in experts markets is generally viewed as a problem. Our model indicates that cautious evaluation is called for: fraud may be an issue from a welfare point of view only if fraud is costly ($f > 0$). If the fraud cost is nil, equilibrium involving fraud is efficient. Receiving a truthful recommendation is not valuable *per se*; equilibrium recommendation and visit strategies matter only for the costs that they entail. This feature may convince the reader that models where fraud is not associated with over-consumption may miss important insights.

4 Discussion

Further insights may be gained by clarifying the role of some of the above assumptions. We discuss issues related to those insights, and also comment on some extensions.

4.1 The role of detailed bills and commitment

In our model experts supply a bill listing inputs used. This is a common practice, as is clear from a visit to the local mechanic or a home improvement subcontractor. It is therefore reasonable to include it as an assumption in our analysis. Nevertheless, future research may seek to endogenize this feature. An interesting explanation may be that part of the fraud cost represents a moral cost, and that experts are heterogeneous in this moral cost. Then intuition would suggest that the more honest ones may have an incentive to use a detailed bill to signal themselves. Or, if consumers differ in their capacity to verify the inputs used by experts, then these may have an incentive to use a detailed bill to attract consumers with a high ability.

Relaxing the assumption that expert provide a bill would make the model less realistic. However, it would shed light on a possible advantage of not requiring experts to provide bills detailing the inputs used. In such a setting the notion of fraud would be meaningless since the consumer would not be able to verify anything anymore. In

our model, this would correspond to the case where $f = 0$. In particular equilibria with costly over-consumption would disappear. Hence, if experts did not supply detailed bills, *ex post* cream-skimming could be deterred at no cost, so that the market would be able to perform more efficiently.

Our results rely on the fact that the relationship between the customer and the expert is short-term, which in turn implies a lack of commitment by both parties. In particular, we argued that an expert may not commit to making repairs at a loss, implying the restrictions $\underline{p} \geq \underline{c}$ and $\bar{p} \geq \bar{c}$. We argued that such restrictions are sensible, because experts may always use some stratagem to scare away consumers (for instance by pretending that ordering spare parts will be a lengthy process). In fact, an experts market with low intervention prices would also constitute an easy prey for crooks: a fake expert could pretend to make a diagnosis, cash in on the diagnosis price, and then somehow scare away the consumer. The risk of attracting such crooks would put a downward pressure on the diagnosis price and therefore an upward pressure on repair prices.

Yet, it is instructive to further explore what would happen if the conditions $\underline{p} \geq \underline{c}$ and $\bar{p} \geq \bar{c}$ were no longer imposed. Then an expert could choose a tariff such that truthful recommendations do not invite *ex post* cream-skimming, e.g., $(p, \underline{p}, \bar{p}) = (C, 0, 0)$. As a result, equilibria involving costly fraud would clearly not exist. However, *ex ante* cream-skimming would still prevent the existence of efficient equilibria. Indeed, if $d < \mu(\bar{c} - \underline{c})$, a consumer would first visit a minor specialist, so that experts offering the tariff $(p, \underline{p}, \bar{p}) = (C, 0, 0)$ would only attract consumers in state M and thereby make a loss. This discussion reveals the value of long-term contracts: for instance, assume that the consumer pays C to a given expert before the loss arises, and that the expert agrees to fix the consumer's subsequent loss at no charge.¹⁹ Because the consumer pays the price before the loss arises, *ex ante* cream-skimming is not an issue. This is reminiscent

¹⁹Dulleck and Kerschbamer (2003) study the case where consumers may (for exogenous reasons) commit to visiting only one expert; then equilibria are efficient.

of contracts such as maintenance contracts, (extended) warranties, or in the health care case an insurance contract together with an exclusivity clause (managed health care). Our model thus provides an explanation for why there may be a market for such long-term contracts, even if consumers are risk neutral and even if administrative costs are significant.²⁰ This result is not new in the literature,²¹ although the commitment rationale is new as far as we know.

Despite their appeal, long-term contracts have obvious drawbacks. First, since the consumer is linked to a particular expert, the expected transportation cost is higher; expert networks may somewhat mitigate that problem. Second, it makes the market even more vulnerable to crooks than a market based on spot contracts, since it is possible to cash in on the fixed fee from a large set of consumers before any losses arise. Based on these and the above observations, we argue that our setup is a plausible one. Further research would be useful to yield sharper predictions as to the exact conditions under which the alternative setups discussed here may be considered.

4.2 Repair or replace?

The above results were derived under the assumption that the loss was fixed with certainty in equilibrium; this amounts to assuming that the consumer's willingness to pay is at least $d + \bar{c}$. While this may be true for expensive durables, for cheaper durables it is sometimes more costly to repair than to replace a defective good. Consider now the case where replacing the good costs some $r < d + \bar{c}$.

First, assume that $r > \bar{c}$; then *ex post* efficiency requires that a defective good be repaired, even if a major intervention is required (since the diagnosis cost then is sunk). The insights of the above analysis may then be directly applied, with a minor variation: since $r < d + \bar{c}$, if in equilibrium the consumer first visits a minor specialist, he gets the good replaced upon being recommended a major intervention. However, it is easy to

²⁰According to Lutz and Padmanabhan (1998), 27% of new car buyers purchase an extended warranty; the proportion is even higher for home electronics products.

²¹See the analysis of Taylor (1995), who also points out the inefficiencies linked to the moral hazard at the maintenance stage arising from such warranties.

verify that all other results remain valid (with some changes in the threshold values).

Things change drastically if r falls below \bar{c} ; then the consumer would turn down any recommendation to get a major intervention (since $\bar{p} \geq \bar{c}$). As a result, the only potential equilibrium is that of specialization: the consumer first visits a minor specialist, who recommends truthfully; the consumer accepts only a minor intervention. The expected cost for the consumer would then be:

$$d + \mu \underline{c} + (1 - \mu)r.$$

But of course, the consumer would effectively visit the minor specialist only if the replacement cost were sufficiently large to warrant the diagnosis cost:

$$d + \mu \underline{c} + (1 - \mu)r < r, \text{ or } r > \underline{c} + \frac{d}{\mu}.$$

By contrast to the case where $r \geq \bar{c}$, here the market performs efficiently.

5 Concluding remarks

In this paper, we have argued that the question of fraud in experts' markets cannot be analyzed without reference to the question of over-consumption. Our theory of fraud 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 the consumer from seeking a better price elsewhere. This argument is general enough to be applied to various industries, such as car repairs or health care, in which information is scarce and each visit to an expert is costly. The model we have solved provides several additional insights.

First, fraud in the form of inefficient over-consumption may arise in equilibrium, even if the market is competitive, and in the absence of insurance. This finding is new in the literature, and it may offer a sound and more 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-consumption. In our theory, the consumer's

inability to commit to visit only one expert is indeed at the heart of the matter: if consumers could commit, a competitive experts market would perform efficiently, since then cross-subsidies which are necessary to remove incentives to defraud are not threatened by cream-skimming. A credible commitment device exists, however, in the form of warranties (or extended warranties), or managed health care; this suggests that even risk-neutral consumers may demand such warranties.

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 under specialization. As we have seen, welfare may be lower in this regime; eradicating equilibrium fraud may thus be a misleading objective.

Third, our model explains 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 allows experts to transfer to customers the profits originating in treatment prices exceeding 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 obviously is a highly relevant aspect of many experts markets, it would be desirable to extend our model to allow for insurance.

Appendix

Proof of Proposition 1:

We first prove that expert make zero profits in an efficient equilibrium. Assume to the contrary that experts offer a set of tariffs such that the consumer first picks a tariff $(p_1, \underline{p}_1, \bar{p}_1)$ satisfying $\underline{p}_1 \in [\underline{c}, \underline{c} + d]$, $\bar{p}_1 \in [\bar{c}, \bar{c} + d]$, and $\underline{p}_1 \geq \bar{p}_1 - f$, and yielding positive expected profits: $p_1 + \mu \underline{p}_1 + (1 - \mu) \bar{p}_1 > C$. Combining these inequalities implies that any such tariff has $p_1 > 0$. But then an expert $i \neq 1$ can profitably attract the consumer on his first visit by offering the tariff $(p_i, \underline{p}_i, \bar{p}_i) = (p_1 - \varepsilon, \underline{p}_1, \bar{p}_1)$ for some $\varepsilon > 0$.

Next we show that $f \geq f^*$ is a sufficient condition for an efficient equilibrium to exist. Assume that all experts offer some tariff $(p, \underline{p}, \bar{p})$ such that the first visited expert makes a truthful recommendation, that the consumer accepts, and such that experts make zero expected profits. Such a tariff exists if and only if $f \geq f^*$. We now prove that if $f \geq f^*$, there exists no profitable deviating tariff.

The arguments preceding the proposition show that as long as $f \geq f^*$, any deviating tariff that would induce the consumer with beliefs $\hat{\mu}_1 = 0$ or $\hat{\mu}_1 = 1$ to turn down the first visited expert's recommendation must be loss-making. Therefore, it only remains to be checked that there exists no profitable deviating tariff that would attract the consumer on his first visit.

Let us assume that expert i deviates by offering a tariff $(p_i, \underline{p}_i, \bar{p}_i)$ that attracts the consumer with beliefs μ ; we show that this tariff must be loss-making. We need to verify this for all the possible recommendation and acceptance decisions that this tariff may give rise to.

(i) Suppose that $r_i = m$ and that the consumer accepts both recommendations. For the consumer to pick expert i for his first visit, it must be that $p_i + \mu \underline{p}_i + (1 - \mu) \bar{p}_i < C$. But then expert i would make a loss.

(ii) Suppose that $r_i = m$ and that the consumer accepts only recommendation m . Upon receiving recommendation M the consumer would visit an expert with the tariff $(p, \underline{p}, \bar{p})$, implying that his expected cost of getting the loss fixed would be $p_i + \mu \underline{p}_i +$

$(1 - \mu)(p + \bar{p})$. Since $p + \bar{p} \geq \bar{c}$, this implies that $p_i + \mu \underline{p}_i$ must be smaller than $d + \mu \bar{c}$ for the consumer to choose expert i for his first visit. But then expert i would make a loss. A very similar argument may be used to show that there cannot exist a profitable deviating tariff such that $r_i = m$ and the consumer accepts only recommendation M .

(iii) Suppose that $r_i = m$ and that the consumer accepts no recommendation. If the consumer chooses expert i for his first visit, his expected cost of getting the loss fixed would be $p_i + C$, which implies $p_i < 0$ so that expert i would make a loss.

(iv) Suppose that $r_i = M$. Whether or not the consumer accepts the recommendation M , expert i would make a loss. If the consumer does not accept the recommendation, $p_i < 0$ is necessary for the tariff $(p_i, \underline{p}_i, \bar{p}_i)$ to attract the consumer on his first visit. If the consumer accepts the recommendation, the expected revenues that expert i may collect must be smaller than C for the consumer to be attracted in the first place, whereas his expected costs would be $C + \mu f$ since he incurs the fraud cost f when the consumer is in state m . ■

Proof of Proposition 2:

Using the index m to denote a minor specialist tariff, and the index M for a major specialist tariff, let us assume that at least one and at most $n - 1$ experts offer the minor specialist tariff, while the rest offer the major specialist tariff. Assume that a consumer first visits a minor specialist, who makes a truthful recommendation, which the consumer accepts only if it is a minor one. If the minor specialist recommends a major intervention, the consumer turns him down, and visits a major specialist whose recommendation to do a major intervention he accepts. First, since a consumer in state M has beliefs $\hat{\mu} = 0$ upon being recommended a major intervention by the first visited expert, a standard Bertrand competition argument shows that major specialists make zero expected profits. Thus a major specialist tariff satisfies $p_M + \bar{p}_M = d + \bar{c}$, $p_M \geq 0$, and $\bar{p}_M \geq \bar{c}$. Turning now to the tariff of a minor specialist, a consumer who has been recommended a major intervention rejects that recommendation if and only if

$\bar{p}_m > d + \bar{c}$; for any $\underline{p}_m \geq \underline{c}$, the minor specialist's best response is to make a truthful recommendation. It is straightforward to show that minor specialists must make zero expected profits. Assume that they made a profit, either on the diagnosis or on the minor intervention or on both. Then an expert could profit by offering a slightly lower price on the profitable service, while offering the same price for a major intervention.

We can further pin down the minor specialist tariff by showing that $p_m + \underline{p}_m \leq d + \underline{c}$. Assume the contrary. Then an expert could make a profit by offering a tariff with a slightly lower price for the diagnosis and the minor intervention: this would attract a consumer who was recommended a minor intervention by a minor specialist.

Thus a minor specialist tariff satisfies $p_m + \mu \underline{p}_m = d + \mu \underline{c}$, $p_m \geq 0$, $\underline{p}_m \geq \underline{c}$, and $p_m + \underline{p}_m \leq d + \underline{c}$. Taken together these conditions imply $\underline{p}_m = \underline{c}$ and $p_m = d$. The consumer's expected cost of getting the loss fixed is $C + (1 - \mu)d$, where the second term reflects the costly second diagnosis for a consumer in state M .

It remains to be checked whether an expert could profitably deviate by offering a tariff that would attract the consumer on his first visit. As in the proof to Proposition 1, let us assume that expert i deviates by offering a tariff $(p_i, \underline{p}_i, \bar{p}_i)$ designed to attract the consumer with beliefs μ . Again we need to rule out a profitable deviation for every possible combination of recommendation and acceptance decisions that this tariff may give rise to.

(i) Suppose that $r_i = m$ and that the consumer accepts both recommendations. Such a tariff exists if and only if $f \geq f^*$, in which case it would represent a profitable deviation as long as $d > 0$. If $d = 0$, the deviation would make zero profits.

(ii) Suppose the consumer accepts no recommendation by expert i . Whether or not expert i makes a truthful recommendation, if the consumer chooses expert i for his first visit, his expected cost of getting the loss fixed would be $p_i + C + (1 - \mu)d$, which implies $p_i < 0$ for the tariff to attract the consumer. But then expert i would make a loss.

(iii) Suppose that $r_i = m$ and that the consumer accepts only recommendation M . For the expert to truthfully recommend a minor intervention (that will be turned down),

it must be that $0 \geq \bar{p}_i - f - \underline{c}$, which together with $\bar{p}_i \geq \bar{c}$ implies $f \geq \bar{c} - \underline{c}$. But then $f \geq f^*$, and we know from (i) that specialization is not an equilibrium.

(iv) Suppose that $r_i = M$ and that the consumer accepts this recommendation. For such a tariff not to be loss-making, it must have $p_i + \bar{p}_i \geq C + \mu f$ since the fraud cost f is incurred in state m . If $p_i + \bar{p}_i \geq \bar{c} + d$, the consumer would not visit expert i . If $p_i + \bar{p}_i < \bar{c} + d$, a consumer who first visited a minor specialist and was recommended a major intervention would then visit expert i . But in that case expert i would make a loss, since his revenue $p_i + \bar{p}_i$ would fall short of his cost $d + \bar{c}$. Thus expert i could profit if and only if the consumer prefers to immediately visit him, i.e., if

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

which is equivalent to

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

■

Proof of Proposition 3:

Assume that there exists an equilibrium where the first visited expert always recommends a major intervention, and the consumer accepts this recommendation. Further assume that there exists at least one expert offering a minor specialist tariff $(d, \underline{c}, +\infty)$. A standard Bertrand competition argument may be used to show that profits are nil in equilibrium. Thus the consumer's expected cost for getting the loss fixed is $C + \mu f$. Since the consumer learns nothing along the equilibrium path, we only need to determine whether or not an expert offering a tariff leading to some other recommendation and/or acceptance decisions may attract the consumer on his first visit. Thus, we check whether there exists a profitable tariff $(p_i, \underline{p}_i, \bar{p}_i)$; again we need to take into account all the possible recommendation and acceptance decisions that this tariff could give rise to.

(i) Suppose that the consumer accepts no recommendation by expert i . Whether or not expert i makes a truthful recommendation, if the consumer chooses expert i for his

first visit, his expected cost of getting the loss fixed would be $p_i + C + \mu f$, which implies $p_i < 0$ for the tariff to attract the consumer. But then expert i would make a loss.

(ii) Suppose that $r_i = m$ and that the consumer accepts only recommendation m . From the arguments developed in (iv) in the proof of Proposition 2, it is clear that there exists such a tariff that would attract the consumer on his first visit and that would be profitable if and only if $f > f^{**}$.

(iii) Suppose that $r_i = m$ and that the consumer accepts both recommendations. Since there is one expert offering a minor specialist tariff, such a tariff would attract the consumer on his first visit and would be profitable if and only if $f \geq f^*$.

(iv) Suppose that $r_i = m$ and that the consumer accepts only recommendation M . For the expert to truthfully recommend a minor intervention (that will be turned down), it must be that $0 \geq \bar{p}_i - f - \underline{c}$, which together with $\bar{p}_i \geq \bar{c}$ implies $f \geq \bar{c} - \underline{c}$. But then $f \geq f^*$, and we know from (iii) that fraud is not an equilibrium. ■

Proof of Proposition 4:

At the first visited expert, there are six possible sequences of events. First, if the expert makes a truthful recommendation, there are four possibilities: the consumer may accept both, only one, or none of the recommendations. Second, if the first visited expert always recommends a major intervention, the consumer may either accept or turn down the recommendation. Propositions 1-3 characterize conditions under which three of those sequences of events may be part of an equilibrium. Here we investigate whether the three that remain may be part of an equilibrium.

(i) Assume that the first expert recommends truthfully but the consumer accepts only a recommendation to do a major intervention. Using arguments that by now should be familiar to the reader, one can show that in such an equilibrium experts would make zero profits. The first visited expert would therefore offer a tariff where $p_1 + \bar{p}_1 = d + (1 - \mu)\bar{c}$. Upon being recommended a minor intervention, the consumer would visit a second expert who would fix the loss for a total price of $d + \underline{c}$. The

consumer's expected cost of getting the loss fixed would therefore be $C + \mu d$. Now, for the first expert to truthfully recommend a minor intervention (that will be turned down), it must be that $0 \geq \bar{p}_1 - f - \underline{c}$, which together with $\bar{p}_1 \geq \bar{c}$ implies $f \geq \bar{c} - \underline{c}$. But then $f \geq f^*$, and an expert could offer a tariff inducing a truthful recommendation and consumer acceptance. This tariff would attract the consumer and make a profit as long as $d > 0$. However, if $d = 0$, such a deviation would make zero profits.

(ii) Assume that the first expert recommends truthfully but the consumer accepts no recommendation. The consumer knows his true state after having visited the first expert, and his loss would be fixed by a minor specialist with a tariff satisfying $p_m + \underline{p}_m = d + \underline{c}$ if he is in state m , and by a major specialist with a tariff satisfying $p_M + \underline{p}_M = d + \bar{c}$ if he is in state M . The expected cost of getting the loss fixed would be $d + C$. As long as $d > 0$, the consumer would be better off by immediately visiting a minor specialist: the expected cost of getting the loss fixed would then drop to $C + (1 - \mu)d - (1 - \mu)(\underline{p} - \underline{c})$. However, if $d = 0$, this could be an equilibrium.

(iii) Assume that the first expert always recommends a major intervention and the consumer turns down the recommendation. The consumer would learn nothing at cost d . However, if $d = 0$, this could be an equilibrium. ■

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.
- DULLECK, U. and R. KERSCHBAMER (2003). "On Doctors, Mechanics and Computer Specialists Or Where are the Problems with Credence Goods?" WP 0101, University of Vienna.
- 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.
- 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. (2003). "When Do Experts Cheat and Whom Do They Target?" forthcoming, *RAND Journal of Economics*.
- LUTZ, N.A. and V. PADMANABHAN (1998). "Warranties, Extended Warranties, and Product Quality," *International Journal of Industrial Organization*, 16:463-493.
- 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.

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 .