Consumer Information in a Market for Expert Services

Kyle Hyndman[*] Saltuk Ozerturk[†]

Southern Methodist University Southern Methodist University

April 28, 2008

Abstract

This paper analyzes the implications of heterogeneously informed consumers in a market for expert services. We consider a credence good model where an expert can cheat consumers by recommending an expensive treatment while only a cheap one is needed. Our main question is to investigate whether uninformed consumers are the most likely victims of expert cheating. We show that when consumers are heterogeneously informed on their true benefit from an expensive treatment, there is no equilibrium where the expert only cheats uninformed consumers. In fact, informed high-value consumers are the most frequent victims of cheating. Surprisingly, more information on the consumer side increases the inefficiency of the market outcome in terms of the foregone, but required, treatments. When some consumers receive noisy information signals on whether their problem is serious or minor, while others remain uninformed, in the unique equilibrium the expert is truthful to all types of consumers, regardless of their information status.

[*] *Department of Economics, Southern Methodist University, 3300 Dyer Street, Suite 301R, Dallas, TX 75275. E-mail: hyndman@smu.edu, url: http://faculty.smu.edu/hyndman.*

[†] *Department of Economics, Southern Methodist University, 3300 Dyer Street, Suite 301X, Dallas, TX 75275. E-mail: ozerturk@smu.edu, url: http://faculty.smu.edu/ozerturk.*

1 Introduction

One of the most frequent consumer complaints involve so-called *credence goods*. These are products and services purchased from informed 'experts' such as auto mechanics, home improvement contractors, appliance service-persons, physicians and lawyers. An important feature of these services is that the provider of the service also assumes the role of an expert and determines how much or what type of service the consumer needs. Even when the success of the service is observable to the consumer *ex post*, consumers typically can never determine the type of the service they needed in the first place. In certain instances, the consumers may never know what type of service was *actually* performed by the expert. Furthermore, most consumers are unable to evaluate their true benefit from receiving a certain type of treatment (*e.g.,* how much changing a car part actually adds to the well-being of a car). This informational asymmetry between experts and consumers creates obvious incentive problems: a mechanic may easily claim that a car needs a major and expensive repair, while only a minor and inexpensive repair is necessary.[1] Experts may attempt to overtreat consumers by providing unnecessary and expensive services, or overcharge them by claiming to provide an expensive treatment, although they actually solve the problem with an inexpensive treatment.

The concern in everyday life that experts may behave fraudulently is so common that consumer groups regularly provide tips to protect consumers from expert cheating. One common piece of advice given to consumers is that they should gather information about their problem and possible remedies before visiting an expert. It is argued that by appearing to be more informed, consumers can prevent the expert from cheating. The following excerpt from a consumer advice website captures this folk wisdom:[2]

> Often you can get a good idea of what's wrong with a vehicle by entering the keywords of the symptoms at your favorite internet search engine. There are message boards and helpful websites designed to help diagnose car problems. Although this won't aide in the repair of your vehicle, you will be more informed when you contact a car repair shop. *If you sound as if you know something about cars you are more likely to obtain a fair estimate. Uneducated individuals are more likely to be taken advantage of.*

The argument behind this folk wisdom is straightforward: the more substantial the informational asymmetry between an expert and a consumer, the easier it becomes for the

[1] A recent field study by Schneider (2006) reports that at only 27 of the 40 garages he visited, mechanics correctly diagnosed that the car had a disconnected battery cable (which was the real problem), while 10 of them recommended costly repairs that were plainly unnecessary, like replacing the starter or the battery.

[2] See www.essortment.com/hobbies/overpricingrip_sfsa.com. Italics added in quote above.

expert to behave opportunistically and cheat. Perhaps it is because of this straightforward intuition that, to date, there has been no formal analysis of the implications of consumer information in a market for expert services. In this paper, we question this conventional thinking by introducing heterogenously informed consumers in a credence good model. In particular, we ask whether uninformed consumers are indeed the most likely victims of expert cheating, and how the efficiency of the market outcome is affected with more information on the consumer side. Our analysis illustrates that the folk wisdom summarized above is somewhat misguided in the sense that uninformed consumers may not always be the most likely victims of expert cheating. Furthermore and perhaps more surprisingly, we show that more information on the consumer side may actually increase the inefficiency of the market outcome in the form of required but foregone treatments.

Following most existing models, a consumer's problem in our framework can either be serious (requiring an expensive treatment) or minor (requiring a cheap treatment). *Ex ante*, a consumer does not know whether his problem is serious or minor. Consumers can visit a monopolist expert who can perfectly diagnose and treat their problem. Upon the consumer's visit and the subsequent diagnosis, the expert can provide an expensive treatment that solves both serious and minor problems, or a cheap treatment which only solves the minor problem. The consumers cannot ex post verify the actual treatment they receive from the expert, but are protected with limited liability.[3] Due to the unverifiability of treatments and the limited liability protection, the potential fraud we consider is one of overcharging. The expert may recommend and charge for an expensive treatment when the problem is minor and the cheap treatment is provided. We also assume that the consumers with a serious problem are heterogenous in the true benefit they receive from having their problem fixed: the expert's expensive treatment can either yield a high or a low benefit for a consumer with a serious problem.[4] In this setting, we focus on the implications of two potentially different pieces of consumer information: (i) information on the true benefit of an expensive treatment, (ii) information on the seriousness of the problem.

We first introduce consumer heterogeneity in information by assuming that some consumers are informed of their true benefit from receiving an expensive treatment, while some are not. All consumers, however, remain uninformed about the type of their problem. The expert, on the other hand, can perfectly identify not only the type of the consumer's problem, but also the true benefit of a required expensive treatment for that specific consumer. This assumption is meant to capture a second dimension of the expert's informational superiority over the consumer: typically an expert can tell not only the type of the treatment required,

[3]As standard in the literature on credence goods, limited liability protection implies that the expert cannot provide the cheap treatment if the problem is serious.

[4]As we further explain when we lay out the model, this assumption captures the notion that usually sophisticated and expensive treatments work differently across different consumers.

but also the extent that a sophisticated treatment actually suits a specific consumer. For example, a marketing expert can identify that a corporate client needs an expensive marketing campaign to penetrate into a new market. From earlier experience, the expert may also know how much this expensive campaign would increase demand for the client's specific product.

Since our main objective is to investigate if the expert selectively cheats uninformed consumers, we primarily analyze the game under the assumption that the expert can perfectly distinguish between informed and uninformed consumers and hence can condition her recommendation strategy on the information status of a consumer.[5] The analysis of this model yields the following results. First, when all consumers are uninformed, the unique equilibrium involves no cheating. Second, when some consumers are informed about their true benefit from an expensive treatment, there is no equilibrium outcome in which the expert only cheats uninformed consumers. Depending on parameter values, the unique equilibrium involves one of the following three outcomes: (i) the expert only cheats informed high-value consumers, (ii) the expert cheats informed high-value and uninformed consumers, but is truthful to informed low-value consumers, and (iii) the expert is truthful to all types of consumers. Accordingly, it is the informed high-value consumers, and not uninformed consumers, who are the most frequent victims of expert cheating. Finally, and perhaps most surprisingly, more information on the consumer side increases the inefficiency of the market outcome in terms of the foregone but required treatments. All types of equilibrium outcomes that emerge when some consumers are informed about their true expensive treatment benefit involve more efficiency loss than the truthful equilibrium that arises when all consumers are uninformed.

As an extension of our basic model, we next study the case in which, prior to visiting the expert, some consumers receive signals about whether they have a serious or a minor problem. The signal, though noisy, is informative and depending upon its realization some consumers will be more (less) optimistic that their problem is minor. Given this information structure, some informed consumers may be pessimistic and believe quite strongly that their problem is serious when, in reality, it is actually minor. With such pessimistic beliefs, it is these consumers who are most likely to accept an expensive treatment recommendation. As such, when the expert can identify whether a consumer is informed or not, and the particular signal he has observed, one might expect the equilibrium to involve some cheating, with the pessimistic consumers the victims of expert cheating. Instead, however, we show that the unique equilibrium involves no cheating, independently of whether the expert is able to

[5] In Appendix B, we also provide a detailed analysis of the game when the expert cannot perfectly identify informed consumers from uninformed ones. We show that the equilibrium outcomes in this case are qualitatively similar to the ones that emerge when the expert can identify perfectly the information status of all consumers.

distinguish informed from uninformed consumers.

The intuition for the above result is as follows. In a recommendation subgame for a given list of treatment prices, a consumer's incentive to accept an expensive treatment recommendation depends on his beliefs that the problem is serious, the difference between the price of an expensive treatment and his benefit of having a serious problem fixed, and the expert's cheating behaviour at the posted prices. In particular, as long as the price of the expensive treatment is strictly less than the consumer's valuation, the more strongly the consumer believes his problem to be serious, the more tolerant of expert cheating he is. When the expert chooses the treatment prices *ex ante*, she faces the following trade-off: by increasing the price of the serious treatment, consumers reject such recommendations more frequently, but the profit margin increases for those consumers who still accept. It turns out that the increase in profit margin dominates the lower acceptance rate, which causes the expert to set the price of the expensive treatment at the consumers' (in this case common) valuation for having a serious problem repaired. However, at this price, regardless of their information and initial beliefs that the problem is serious, all consumers would reject with certainty if the expert cheats with strictly positive probability. This is what induces the expert to be truth-telling.

Related Literature. In reality, the quality and amount of information that consumers possess differ substantially. In the specific case of car repairs, some consumers may have a good deal of prior knowledge about car parts, the nature of their problem and their true benefit from certain types of treatments. On the other hand, some consumers may be completely uninformed about car repairs that they cannot even tell whether the mechanic actually performed the service recommended. Despite this casual observation, the existing literature consider models where all the consumers are equally informed about the nature of their problem and their potential benefits from receiving treatment by the expert. This point has also been raised by Dulleck and Kerschbamer (2006) in their critical survey article on the economics of credence goods where they write: "Thus, technical expertise, or expert's expectation of its existence on the consumer side, may affect market outcomes. The existing literature has ignored consumers' heterogeneity in expertise so far" (p. 31).

The closest to our paper is Fong (2005) who formally introduces the notion that an expert's recommendation strategy is typically selective and can be best understood to be conditional on observable and heterogenous consumer characteristics. Our main focus is to investigate how the information status of a consumer as an observable characteristic determines an expert's recommendation strategy and affects the market outcome. To sharpen this focus, we abstract away from price discrimination considerations and build upon Fong's framework, which shows that selective cheating may arise as a substitute for price discrimination. While Fong introduces an elegant framework to illustrate how an expert can selectively cheat high valuation and high cost consumers, his analysis does not investigate the implications of heterogenous consumer information on the expert's cheating behaviour, which is the focus of our paper.

To the best of our knowledge, ours is the first paper that considers a credence goods market with heterogenously informed consumers. The theoretical literature on credence goods is small but growing.[6] One set of papers examine the implications of a consumer's ability to search for second opinions. Wolinksy (1993) considers a competitive setting with many experts, and show that cheating can be eliminated when consumers search for second opinions and experts have reputational concerns. Pesendorfer and Wolinsky (2003) show that consumers' search for second opinions motivates experts to exert costly effort that improves the accuracy of their diagnosis. Alger and Salanie (2006) introduce a fraud cost by allowing the consumers to partially verify the actual inputs the expert uses during her treatment: they show that fraudulent over-treatment may appear as an equilibrium even in a competitive model. Emons (1997, 2001) examine how the market price mechanism can eliminate fraudulent behaviour when experts have capacity constraints and the actual treatment received is verifiable by consumers.

In a durable goods model, Taylor (1995) illustrates how *ex post* pricing and extended service plans provide incentives to customers to properly take care of their durable goods. In a model with exogenous prices and homogenous consumers, Pitchik and Schotter (1987) demonstrate a mixed strategy equilibrium that involves cheating. In another model with exogenous treatment prices, Sülzle and Wambach (2005) study the impact of variations in the degree of insurance on the amount of fraud in a physician-patient relationship. Again, in the context of medical services, Dranove (1988) analyzes how demand inducement by physicians relates to the treatment price and other exogenous variables.[7] Biglaiser (1993) shows that the presence of experts with the ability to identify sellers' qualities can serve to mitigate the informational problems between sellers and buyers. None of these papers address the implications of heterogenous consumer information.

The rest of the paper proceeds as follows. In the next section, we lay out our model. Section 3 analyzes the case when some consumers are informed about their true expensive treatment benefits. Section 4 focuses on the case when, some consumers receive information signals about the type of their problem. Section 5 concludes. All proofs not presented in the text can be found in various appendices.

[6]The seminal work in this literature is by Darby and Karni (1973) who coined the term "credence good."

[7]In another contribution in the health economics literature, De Jaegher and Jegers (2001) describe how the credence good framework can be applied to the analysis of supplier induced demand hypothesis in medical services.

2 The Model

In this section, we describe a basic model of a credence good market.

The consumers and the expert. There is a continuum of consumers with measure one. Each consumer (he) either has a serious problem (denoted by state $\omega = s$) that requires an expensive treatment; or a minor ($\omega = m$) problem that requires a cheap treatment. A consumer does not know whether his problem is serious or minor. The *ex ante* probability of having a serious problem is given by $\Pr(\omega = s) = \alpha \in (0, 1)$. As in Emons (2001) and Fong (2005), the consumers can visit a *monopolist* expert (she) who can perfectly diagnose and treat their problem. Based on the diagnosis, the expert can reject the consumer, or recommend an expensive treatment at a price p_s or a cheap treatment at a price p_m. Providing a cheap treatment costs the expert $c_m > 0$, whereas an expensive treatment costs $c_s > c_m$.

Verifiability and Liability. The consumers cannot observe or verify the actual treatment they receive.[8] They can only tell whether their problem is fixed or not. An expensive treatment fixes both types of problems, whereas a cheap treatment only fixes the minor problem. Furthermore, the consumers are protected by limited liability: the expert cannot recommend and perform a cheap treatment if an expensive treatment is required (*i.e.,* if the expert agrees to treat the consumer, she must fix the problem).

Treatment Benefits. If a minor problem is treated, all consumers receive a benefit $v_m > c_m > 0$. Following Fong (2005), we assume that the consumers are heterogeneous in the benefit they receive from having an expensive treatment. When their problem is serious, depending on the consumer's type the expensive treatment provides a benefit of either v_s^h or v_s^l with $v_s^h > v_s^l > v_m$ and $v_s^l > c_s$.[9] The *ex ante* probability that the consumer is of type v_s^h is given by $\Pr\left(v_s^h\right) = \theta$. Different than Fong (2005), however, we assume that *ex ante* the consumers do not know whether their benefit from the expensive treatment is v_s^h or v_s^l. This assumption enables us to introduce the notion that a consumer may have an informed or uninformed valuation for an expensive treatment. We explain how a consumer can be informed about the true benefit of an expensive treatment shortly.[10]

As part of her diagnosis, the expert can perfectly determine the true benefit of her

[8] Previous work by Pitchick and Schotter (1987), Wolinsky (1993) and Fong (2005) also assume that the actual treatment the expert provides is not verifiable. In Alger and Salanie (2006), the consumers can partially verify the actual inputs the expert uses during her treatment.

[9] Since $v_m > c_m$ and $v_s^l > c_s$, both problems are efficient to fix.

[10] We only consider heterogeneity of treatment benefits from an expensive treatment. Introducing heterogenous benefits from a cheap treatment only complicates the analysis without changing the cheating behaviour of the expert. We explain why this is the case in detail in Remark 2 at the end of Section 3.

expensive treatment for a consumer with a serious problem. This assumption captures a second dimension of the expert's informational superiority over the consumer. Typically, an expert is able to identify not only the type of treatment required, but also the true benefit of a more sophisticated and expensive treatment for a specific consumer. For example consider a doctor-patient relationship where the issue is possible side-effects from medical treatments. An expensive procedure may have different side-effects depending on the patient's medical history and type. The physician, due to some experience with former patients, may be better informed about such side-effects than the patient. Our assumption seems to be consistent with many relevant settings where an expert's expensive treatment works differently across consumers depending on a consumer's type, and by former experience it is the expert who knows the true treatment benefit better than the consumer.[11] We also maintain the following assumptions:

Assumption 1. *The treatment benefits and costs satisfy:*

$$\alpha v_s^t + (1-\alpha)v_m < c_s \text{ for } t \in \{h, l\} \tag{A1}$$

Assumption 2. *The expert cannot price discriminate across consumers, but can follow recommendation strategies contingent on consumer type.*

Assumption 1 rules out a fixed price equilibrium in which the expert sets a single price and agrees to treat both minor and serious problems at that fixed price.[12] Assumption 2 states that while the expert may only submit a single price vector (p_m, p_s), she is able to condition her recommendation strategy on the observable characteristics of the consumer she faces

Consumer Information. Our key innovation is to introduce heterogenous consumer information in a credence good market. Many real life examples indicate that not all consumers are equally informed in their relationships with experts. Consider the case of car repairs as a motivating example. Suppose, if the problem with the car is serious, then a complete change

[11] An alternative approach is undertaken in a recent paper by Eső and Bzentes (2007). They consider a setting where the consultant/expert does not know the true impact of the information she provides for the client. They show that as long as the expert can tie her compensation to the decision undertaken by the client, she can still extract all the surplus despite not knowing perfectly the usefulness of her advise for the client's welfare.

[12] If $\alpha v_s^l + (1-\alpha)v_m > c_s$, then even if the expert cheats consumers with probability 1, consumers would still accept a serious treatment at a price of $p_s = \alpha v_s^l + (1-\alpha)v_m$, which could be shown to be the optimal pricing strategy for the expert. On the other hand, when Assumption 1 holds, both experts and consumers must employ strictly mixed strategies: Since the price of a serious treatment (p_s) must be higher than the cost (c_s), if the expert always cheated, consumers would reject because the price would be higher than the expected valuation of the consumer ($\alpha v_s^t + (1-\alpha)v_m$). More details can be found in Appendix A.1, where our main result is proven.

of the transmission is required, whereas for a minor problem replacing the clutch would be sufficient. Some consumers, simply as part of a lifetime hobby, might have developed a much better familiarity with car parts. An informed consumer may have a better idea when and why changing the transmission is essential for fixing the car. Furthermore, this consumer is also likely to possess some information about the general well-being of the car, and hence may know how long the new transmission would have the car running before creating another problem. As such, some consumers may have better information than others about the type of problem they face, and their true benefit from receiving a more sophisticated and expensive solution from an expert.

In the current framework, these ideas can be captured by introducing two potentially different pieces of information on the consumer side. Before visiting the expert, a consumer can receive an information signal indicating:

(i) his true benefit from receiving an expensive treatment if his problem is serious, or

(ii) the true nature of his problem (whether his problem is a minor or a serious one).

Although we separately analyze both scenarios, in the main body of the paper we first consider the case when a fraction $\lambda \in [0, 1]$ of consumers observe an information signal $\tilde{z}$ which indicates whether they benefit v_s^h or v_s^l from an expensive treatment.[13] The information signal $\tilde{z}$ can take two values: A high signal ($z = h$) perfectly indicates that the consumer is of type v_s^h, whereas a low signal ($z = l$) perfectly indicates that the benefit is only v_s^l.

Expert's Information. If a consumer visits the expert, the expert perfectly identifies whether the consumer has a serious or a minor problem, and whether he benefits v_s^h or v_s^l from a required expensive treatment. The expert, however, may or may not be able to identify whether a consumer is informed or not about his true expensive treatment benefit. Depending on the context, the expert and a consumer may have a close interaction that enables the expert to easily observe the consumer's background, experience and expertise level about the specific service in question. In some other situations, however, the expert may not even meet the client and hence may have no idea about the consumer's information status. Our primary focus is to investigate whether the expert will selectively and more frequently cheat uninformed consumers. This question calls for a setting where the expert can identify consumers as informed and uninformed so that she can base her recommendation strategy on this consumer characteristic. At the same time, it is also interesting to address the implications when the expert anticipates facing some informed consumers without being able to tell uninformed and informed ones from each other. In what follows, we shall analyze

[13] We analyze the case when the consumers can receive information on the type of their problem (serious or minor) in Section 4.

both settings and specifically distinguish between the cases when the expert can identify a consumer as uninformed or informed and when she cannot.

Sequence of Events. The timing of the game is as follows:

STAGE 1: Nature decides whether a consumer has a serious or a minor problem. Nature also decides whether a consumer with a serious problem benefits v_s^l or v_s^h from an expensive treatment. The consumers do not know if their problem is minor or serious and if they are of type v_s^l or v_s^h. A fraction λ of consumers learn perfectly their true benefit from receiving an expensive treatment.

STAGE 2: The expert optimally chooses and announces a price vector (p_m, p_s) where p_m and p_s are the prices for cheap and expensive treatments.

STAGE 3: The consumer visits the expert who perfectly identifies if the problem is serious or minor, and whether the consumer is of type v_s^h or v_s^l. The expert may or may not observe if the consumer is informed or not. Based on the diagnosis, the expert either rejects to treat the consumer or recommends an expensive or a cheap treatment.

STAGE 4: The consumer can accept or reject the expert's recommendation. If he accepts, the expert provides a treatment unobservable to the consumer and charges a fee according to the prices posted in Stage 2. If the consumer rejects, the problem remains untreated.

3 Analysis and Results

We first introduce some notation to analyze the game. Consider a recommendation subgame that starts upon the expert posting a price vector (p_m, p_s). In any such subgame, the expert observes the consumer's problem (serious or minor) and consumer's type (v_s^h or v_s^l). Conditioning on the problem being $i \in \{m, s\}$ and the consumer being of type v_s^t with $t \in \{h, l\}$, a pure strategy for the expert in the subgame (p_m, p_s) specifies whether she refuses to provide treatment, recommends a serious treatment or recommends a minor treatment. A mixed strategy assigns probabilities of taking these actions with $\rho_i^{t,k}$ denoting the probability of rejecting a type $(t, k) \in \{h, l\} \times \{I, N\}$ consumer with a problem $i \in \{m, s\}$, and $\beta_i^{t,k}$ denoting the probability of recommending a serious treatment to such a consumer. For clarity, note that the index $t \in \{h, l\}$ indicates whether the consumer is of high or low type, while the index $k \in \{I, N\}$ indicates whether the consumer is informed or uninformed about his true expensive treatment benefit.[14] Of course, it may or may not be possible for the expert to

[14]The probability of recommending a minor treatment to this consumer is then given by $1 - \beta_i^{t,k} - \rho_i^{t,k}$.

identify a consumer as informed or uninformed. If the expert cannot identify the consumer's information status, she cannot condition her recommendation on this additional consumer characteristic: while analyzing this case, we will drop the index $k \in \{I, N\}$.

A pure strategy for a consumer specifies whether he rejects or accepts the recommended treatment $i \in \{m, s\}$ at the posted prices (p_m, p_s). In terms of the information they might have, there are three possible consumer profiles: those with a high signal $z = h$, those with a low signal $z = l$, and uninformed consumers (we denote them as n). Accordingly, a mixed strategy for a consumer of type $z \in \{h, l, n\}$ assigns probabilities of accepting (γ_i^z) and rejecting $(1 - \gamma_i^z)$ a recommendation $i \in \{m, s\}$.

3.1 Benchmark: all consumers are uninformed

As a benchmark, we first analyze the case when *all* consumers are *uninformed* $(\lambda = 0)$ about their true benefit from an expensive treatment. In this case, all consumers have an *ex ante* valuation $\bar{v}_s$ from an expensive treatment where

$$\bar{v}_s = \theta v_s^h + (1 - \theta) v_s^l.$$

The expert can condition her recommendation strategy on the type of the problem (m or s) and the consumer's type (v_s^h or v_s^l).[15]

We now establish that when all consumers are uninformed, there is a unique equilibrium with $p_s^* = \bar{v}_s$ and $p_m^* = v_m$ which involves no cheating. Since this benchmark result is a variation of Proposition 1 in Fong (2005), here we only describe the key features of the argument and omit a formal proof. First, it can be shown that in any equilibrium we must have $(p_m, p_s) \in [c_m, v_m] \times [c_s, \bar{v}_s]$.[16] Second, observe that in recommendation subgames with $(p_m, p_s) \in [c_m, v_m] \times [c_s, \bar{v}_s]$, we must have $\beta_s^t = 1$ for $t \in \{h, l\}$. This follows, because due to limited liability the expert always recommends expensive treatment when the problem is serious. Third, for $p_m \in [c_m, v_m]$, the consumers always accept a cheap treatment recommendation. Finally, for expensive treatment prices $p_s \in [c_s, \bar{v}_s)$ we must have $\gamma_s^n \in (0, 1)$ and $\beta_m^t \in (0, 1)$ for $t \in \{h, l\}$.[17] For $p_s \in [c_s, v_s^l)$, an uninformed consumer mixes

[15]Since all consumers are uninformed, we omit the superscript $k \in \{I, N\}$. A mixed strategy profile for the expert in a subgame (p_m, p_s) is then given by the probabilities $\{\rho_i^t, \beta_i^t, 1 - \beta_i^t - \rho_i^t\}$ for $t \in \{h, l\}$ and $i \in \{m, s\}$. A mixed strategy profile for the uninformed (type n) consumer is given by the probability γ_i^n of accepting a recommendation $i \in \{m, s\}$

[16]To see this, note that if $p_s > \bar{v}_s$, then all consumers reject an expensive treatment with probability 1. Similarly, if $p_m > v_m$, all consumers will reject a minor treatment with probability 1. Next, if $p_m < c_m$, the expert would refuse to treat the consumer (because doing so would generate a loss). Therefore, in equilibrium we must have $p_m \in [c_m, v_m]$. Finally, if $p_s < c_s$, the expert would either always recommend the cheap treatment (which will be rejected by consumers if $p_s \in (v_m, c_s)$) or refuse to treat the consumer.

[17]To see this, consider an uninformed consumer and suppose to the contrary that $\gamma_s^n = 1$. In this case, the expert always cheats and sets $\beta_m^h = \beta_m^l = 1$. But then the expected benefit of accepting an expensive treat-

between accepting and rejecting an expensive treatment only when the expert cheats with a probability

$$\beta_m^h = \beta_m^l = \beta_m^N = \frac{\alpha(\bar{v}_s - p_s)}{(1-\alpha)(p_s - v_m)}.$$

For the expert to mix between recommending the expensive and cheap treatments when the problem is minor, the uninformed consumers must be accepting expensive treatments with probability:

$$\gamma_s^n = \frac{p_m - c_m}{p_s - c_m}.$$

Using these expressions, for $(p_m, p_s) \in [c_m, v_m] \times [c_s, \bar{v}_s]$ the expert's expected profit function can be computed as:

$$\Pi(p_m, p_s) = \alpha\,(p_s - c_s)\,(\frac{p_m - c_m}{p_s - c_m}) + (1-\alpha)(p_m - c_m).$$

Maximizing $\Pi(p_m, p_s)$ for $(p_m, p_s) \in [c_m, v_m] \times [c_s, \bar{v}_s]$ yields the equilibrium prices $p_m^* = v_m$ and $p_s^* = \bar{v}_s$. We report the full equilibrium in the Proposition below.

Proposition 1. *Suppose all consumers are uninformed about the true benefit of an expensive treatment. The expert's equilibrium price vector is* $(p_m^* = v_m, p_s^* = \bar{v}_s)$. *The expert does not cheat the consumers and sets* $\beta_m^h = \beta_m^l = 0$ *and* $\beta_s^h = \beta_s^l = 1$. *All consumers accept an expensive treatment recommendation with probability*

$$\gamma_s^n = \frac{v_m - c_m}{\bar{v}_s - c_m}$$

and a cheap treatment recommendation with probability $\gamma_m^n = 1$. *The expert never refuses to treat consumers.*

Proof. Omitted. □

The intuition for the no-cheating result is as follows. A consumer's best response behaviour when responding to an expensive treatment recommendation depends on his inference of the expert's cheating probability, β_m^t, and the difference between his valuation $\bar{v}_s$ and the price p_s. In particular, at a price $p_s < \bar{v}_s$ consumers still tolerate some cheating and accept an expensive treatment with a positive probability. As the price p_s approaches $\bar{v}_s$, which is the maximum that an uninformed consumer is willing to pay for an expensive treatment, the expert must reduce her cheating probability to get an expensive treatment recommendation accepted. At $p_s = \bar{v}_s$, for the consumer to accept at all, the expert must

ment to an uninformed consumer is $\alpha\bar{v}_s + (1-\alpha)v_m < c_s < p_s$, (by Assumption 1) which is a contradiction. On the other hand, if $\gamma_s^n = 0$, we must have $\beta_m^h = \beta_m^l = 0$ and the expert never cheats, which implies that the benefit from accepting an expensive treatment is $\bar{v}_s > p_s$, a contradiction.

always be truthful. For the expert, charging $p_s = \bar{v}_s$ is optimal because while increasing p_s reduces the acceptance rate γ_s^n, it increases the profit margin even more. As a result, the expert's expected profit $\Pi(p_m, p_s)$ is increasing in p_s. Therefore, it is optimal to set the highest possible price $\bar{v}_s$ and tell the truth, rather than set a price $p_s < \bar{v}_s$ and cheat. In this truthful equilibrium, the consumers are indifferent between accepting and rejecting because they pay their full valuation.

Despite the truthful revelation of expert's information, a feature of the above equilibrium is the efficiency loss in the form of foregone but required expensive treatments. The consumers must reject an expensive treatment recommendation with some probability to induce truthfulness. This positive equilibrium rejection rate, given by $1 - \gamma_s^n$, leads to an underprovision of required expensive treatments. Every required but foregone expensive treatment leads to an efficiency loss of $(\bar{v}_s - c_s)$. The following corollary reports the size of this equilibrium efficiency loss.

Corollary 1. *In the unique equilibrium when all consumers are uninformed about the true benefit of an expensive treatment, the efficiency loss is given by*

$$EL_{\lambda=0} = \alpha \left(\frac{\bar{v}_s - v_m}{\bar{v}_s - c_m} \right) (\bar{v}_s - c_s).$$

3.2 Some consumers are informed about true treatment benefits

Suppose now that a fraction $\lambda > 0$ of consumers are perfectly informed about their true benefit from an expensive treatment. As before, the expert can perfectly diagnose whether a consumer has a serious or a minor problem, and how much a consumer with a serious problem benefits from an expensive treatment. Our primary purpose is to investigate if the expert will selectively cheat uninformed consumers more than the informed ones. Therefore, we first consider the case where the expert can identify whether a consumer is informed or not. We describe the possible equilibrium outcomes in the proposition below.

Proposition 2. *Suppose a fraction $\lambda > 0$ of consumers are informed about their true benefit from an expensive treatment and the expert can identify consumers as informed and uninformed. Then there is a unique equilibrium in which, depending on the parameter values, there are three possible outcomes:*

> ***Type I Outcome****: $p_m = v_m$, $p_s = v_s^l$, the expert cheats informed high types and uninformed consumers, but is truthful to informed low types. All consumers accept an expensive treatment recommendation with a common positive probability.*
>
> ***Type II Outcome****: $p_m = v_m$, $p_s = \bar{v}_s$, the expert cheats informed high types, but is truthful to uninformed and informed low types. The informed low types always*

reject an expensive treatment, whereas uninformed consumers and informed high types accept with a common positive probability.

Type III Outcome: $p_m = v_m$, $p_s = v_s^h$, *the expert is truthful to all consumers. The informed low types and uninformed always reject an expensive treatment, whereas informed high types accept with a positive probability.*

In all outcomes, the expert is always truthful when the problem is serious and a cheap treatment recommendation is accepted with probability one by all types of consumers.

Proof. See Appendix A.1. □

An interesting feature of the above equilibrium characterization is that unlike the case when all consumers are uninformed, cheating may now emerge when some consumers are better informed about their true treatment benefit from an expensive treatment. Perhaps surprisingly, however, there is no equilibrium outcome in which the expert only cheats the uninformed consumers. In fact, the most frequent victims of expert cheating are informed high types, whereas the expert never cheats the informed low types. We illustrate the properties of the three possible equilibrium outcomes in Table 1, and discuss further their main features below.[18] Also, in Figure 1, for a specific set of parameters $(\alpha, v_s^h, v_l^h, v_m, c_s, c_m)$ we identify the ranges of (λ, θ) under which each of the three equilibrium outcome arises.[19]

TABLE 1: Properties of the Unique Subgame Perfect Equilibrium

Type	p_s^*	p_m^*	$\beta_m^{h,I}$	$\beta_m^{l,I}$	β_m^N	γ_s^h	γ_s^l	γ_s^n
I	v_s^l	v_m	+	0	+	+	+	+
II	$\bar{v}_s$	v_m	+	0	0	+	0	+
III	v_s^h	v_m	0	0	0	+	0	0

A + indicates that the variable is positive in equilibrium.

Type I Outcome: expert cheats uninformed consumers and informed high types.

In this equilibrium outcome, the expert sets $p_s^* = v_s^l$, and $p_m^* = v_m$ and cheats informed high types and uninformed consumers with a positive probability, while being always truthful to informed low types. All consumers accept a cheap treatment recommendation with probability one. All consumers accept an expensive treatment recommendation with a common probability:

$$\gamma_s^n = \gamma_s^h = \gamma_s^l = \frac{v_m - c_m}{v_s^l - c_m}.$$

[18] For both the cheating probabilities ($\beta_m^{t,k}$) and the acceptance rates (γ_s^z), the table also indicates whether they are positive (+) or zero (0) in equilibrium. The exact expressions are derived in the proof in the appendix.

[19] Specifically, $\alpha = 0.25$, $v_s^h = 5$, $v_s^l = 3$, $v_m = 1.5$, $c_s = 2.75$ and $c_m = 1$.

The Type I equilibrium outcome arises when λ is relatively large (most consumers are informed), and θ is relatively small (most consumers are of low type). In others words, this equilibrium outcome emerges when informed low type consumers form the majority of the market: in this outcome, the expert charges the highest possible price for an expensive treatment without losing these consumers all together. The expert is worse off from increasing the price beyond v_s^l, because by doing so, she can only profit from uninformed consumers and informed high-valuation consumers which form only a small fraction of the market when λ is large and θ is small. At the price $p_s^* = v_s^l$, the expert also makes some profit from informed high types and uninformed consumers, since despite being cheated these consumers still accept an expensive recommendation with a positive probability as they face a price low enough with respect to their benefit.

In this outcome, the uninformed consumers are cheated less often than informed high types ($\beta_m^N < \beta_i^{h,I}$) and informed low type consumers are not cheated at all. Despite this fact, all consumers accept an expensive treatment recommendation with the same common probability: informed low types are not cheated, but they pay their full valuation, whereas for example, uninformed consumers are sometimes cheated, but at $p_s^* = v_s^l$ the expensive treatment is relatively a bargain for them (and even more so for informed high types).

Type II Outcome: expert cheats only informed high types.

In the Type II outcome, the experts sets $p_s^* = \bar{v}_s$, and $p_m = v_m$. She cheats only informed high types, while she is always truthful to uninformed consumers and informed low types. Although the expert is always truthful to them, the informed low types always reject an expensive recommendation since the price $p_s^* = \bar{v}_s$ is too high relative to their valuation. The uninformed and informed high types accept an expensive treatment with a common probability

$$\gamma_n(s) = \gamma_h(s) = \frac{v_m - c_m}{\bar{v}_s - c_m}.$$

Although they are sometimes being cheated, informed high-valution types accept an expensive treatment as often as uninformed consumers. This observation again follows, because at a price $p_s^* = \bar{v}_s$, the expensive treatment is still a bargain for informed high types. Since at $p_s^* = \bar{v}_s$, the expert loses the business of all informed low type consumers, for this outcome to arise the fraction λ of informed consumers can not be too high, and the fraction θ of high-valuation consumers cannot be too low.

Type III Outcome: expert is truthful to all consumers.

In this outcome, the expert sets $p_s^* = v_s^h$, and $p_m = v_m$. She is always truthful to all types of consumers. Despite the expert's truthfulness, only the informed high types can afford to accept an expensive treatment recommendation at this high price. They do accept an

FIGURE 1: Equilibrium Outcomes Over Possible Values of θ and λ

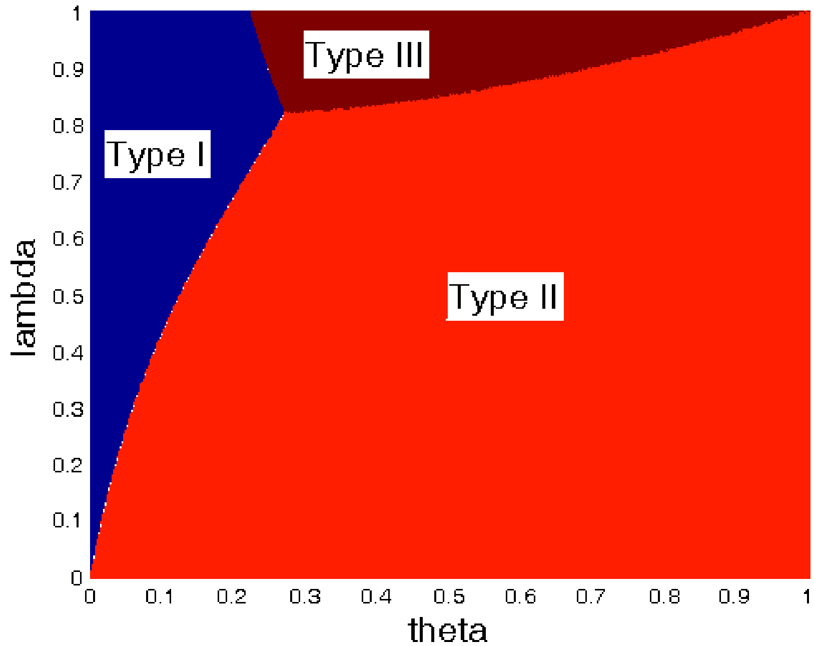

expensive treatment with a probability

$$\gamma_s^h = \frac{v_m - c_m}{v_s^h - c_m}.$$

At $p_s^* = v_s^h$, the expert makes a profit from an expensive treatment recommendation only from informed high types. She gives up all uninformed consumers and informed low types in the market by setting a too high price for them. Accordingly, for this outcome to arise the majority of the market must be of informed high type, i.e., both λ and θ must be sufficiently high.

Comparison of the Efficiency Loss. How does introducing heterogeneous and identifiable consumer information on treatment benefits affect the efficiency loss in the form of foregone but required treatments? We have shown in the previous section that the inefficiency stems from the fact that the consumers have to reject expensive treatment recommendations with positive probability to induce truthfulness to the expert. One may be tempted to suggest that with more informed consumers, efficiency is likely to improve as there will be less need to reject expensive treatments to discipline the expert. Perhaps surprisingly, we now show that this is not the case at all. In fact, all types of equilibrium outcomes that may arise when some consumers are informed about their true benefit from an expensive treatment involve more efficiency loss than the one in which all consumers are uninformed.

Comparison with Type I outcome.

In this equilibrium outcome, uninformed consumers and informed high types are cheated with strictly positive probability. The efficiency loss, which we denote with $EL^I_{\lambda>0}$, can be written as

$$EL^I_{\lambda>0} = \alpha\left(\frac{v^l_s - v_m}{v^l_s - c_m}\right)(\bar{v}_s - c_s) + \alpha\theta\left(\frac{v_m - c_m}{v^l_s - c_m}\right)(v^h_s - v^l_s)$$

The first term in the above expression is the efficiency loss due to required expensive treatments not being provided. Notice that this first term of $EL^I_{\lambda>0}$ is smaller than the efficiency loss in Corollary 1. Therefore, there is an efficieny gain as far as the required expensive treatments are concerned. However, due to expert cheating, now there is an additional efficiency loss with cheap treatments not provided. This loss is captured by the second term in $EL^I_{\lambda>0}$. Calculating the difference $EL_{\lambda>0}(I) - EL_{\lambda=0}$ yields

$$EL^I_{\lambda>0} - EL_{\lambda=0} = \frac{\alpha\theta(c_s - c_m)(v_m - c_m)(v^h_s - v^l_s)}{(v^l_s - c_m)(\theta(v^h_s - v^l_s) + v_l - c_m)} > 0$$

so that the efficiency loss is greater in the Type I equilibrium outcome. In this outcome, the uninformed consumers and the informed high types accept a serious treatment recommendations more often despite the fact that they are sometimes being cheated. This higher acceptance rate by uninformed consumers and informed high types reduces the efficiency loss due to foregone but required expensive treatments. However, precisely because of expert cheating at this low price $p^*_s = v^l_s$, there is now an additional source of inefficiency. Recall that the expert cheats by reporting a minor problem as a serious one and recommending an expensive treatment. When a consumer rejects this recommendation, a required minor treatment is foregone. It turns out that this additional inefficiency more than offsets the gains from required expensive treatments and as a result the Type I equilibrium outcome is more inefficient than the truthful equilibrium of Proposition 1.

Comparison with Type II outcome.

In this case only informed high types will be cheated, while informed low types will reject the expensive treatment with probability one. After some algebra, the efficiency loss in the Type II outcome can be written as:

$$EL^{II}_{\lambda>0} = \alpha\left(\frac{\bar{v}_s - v_m}{\bar{v}_s - c_m}\right)(\bar{v}_s - c_s) + \alpha\lambda\gamma(1-\theta)(v^l_s - c_s) + (1-\alpha)(1-\gamma)\lambda\theta\beta^{h,I}_m(v_m - c_m)$$

The first two terms in the above expression represent the efficiency loss from serious problems not being treated, while the third term represents the loss due to minor problems not being treated when the expert cheats informed high types. Comparing the first term of $EL^{II}_{\lambda>0}$ with the expression for $EL_{\lambda=0}$, we immediately see that $EL^{II}_{\lambda>0} > EL_{\lambda=0}$. In this outcome, the informed high types accept an expensive treatment recommendation more often due to the low price $p^*_s = \bar{v}_s$, despite being cheated with positive probability. However, now all informed low types reject with probability one. As a result, the inefficiency in the form of foregone expensive treatments is now worse. Even without taking into account the additional loss due to some minor problems being untreated, the exclusion of all informed low types makes Type II outcome more inefficient compared to the truthful equilibrium of Proposition 1.

Comparison with Type III outcome.

The reason that the Type III equilibrium is more inefficient than the truthful equilibrium of Proposition 1 is again an exclusion argument. Recall that in the Type III outcome the equilibrium price is $p^*_s = v^h_s$. At this high price, only the informed high types can afford to have their serious problems treated. With all uninformed consumers and informed high types excluded, the inefficiency gets worse in terms of foregone but required expensive treatments.[20] We include the formal expressions for this case in Appendix A.2.

This efficiency comparison establishes the interesting result that introducing identifiable consumer information on treatment benefits not only may give rise to cheating, but also increases efficiency loss in the form of required but foregone treatments. We state this result below.

Proposition 3. *All types of equilibrium outcomes when some consumers are informed about true treatment benefits exhibit more efficiency loss than the truthful equilibrium that arises when all consumers are uninformed.*

Proof. See Appendix A.2 □

Remark 1 (Expert cannot identify informed and uninformed)**.** *A natural question is the extent that the equilibrium characterised in this section depends on the expert being able to identify if a consumer is informed or not. When the expert cannot tell if a consumer is informed or not, then her recommendation strategy may only be a function of the consumer's type $t \in \{h, m\}$ and the true problem $i \in \{m, s\}$. Because of this the analysis is considerably more involved. For example, in pricing sub-games with $p_s > v^l_s$, depending on the measure of informed consumers (i.e., the size of λ), it may not be possible to make both uninformed and informed high types indifferent. In fact, in some sub-games, it may be that informed high types must accept with probability 1, even though they are being cheated with positive probability.*[21] *Despite these extra complications, it turns out that the results when the expert*

[20] Note that in the Type III outcome there is no cheating and hence there is no efficiency loss due to minor pronlems not being solved.

[21] Unlike when the expert can distinguish informed from uninformed, the fact that informed high types accept with probability 1, does not imply that the expert's best response is to cheat with probability 1; if

cannot distinguish informed from uninformed are qualitatively similar to those summarized by Proposition 2: in particular, there is a unique subgame perfect equilibrium with, depending on the underlying parameters, three possible outcomes. In two of these equilibrium outcomes the price is relatively low, and high types (possibly including uninformed consumers) are being cheated with strictly positive probability by the expert, while in the other equilibrium pricing outcome the expert is truthful. We relegate a complete analysis of this case to Appendix B.

Remark 2 (Heterogenous benefits from a cheap treatment)**.** *We have considered heterogenous benefits only from an expensive treatment. We now briefly argue that introducing heterogeneity of benefits from a cheap treatment would not add much insight to our analysis as far as the expert's cheating behaviour is concerned. Suppose that a cheap treatment can yield a benefit of with* v_m^h *or* v_m^l *with* $c_m < v_m^l < v_m^h < c_s$*. Furthermore, suppose that all consumers are perfectly informed whether they are of type* v_m^h *or* v_m^l*. In this case, if the expert sets* $p_m > v_m^l$*, a consumer of type* v_m^l *will reject a minor cheap treatment recommendation with certainty. This consumer with type* v_m^l *would reject any expensive treatment recommendation with certainty as well, since any strictly positive acceptance probability triggers the expert to always misreport minor problems as serious. Given that any price* $p_m > v_m^l$ *drives away completely all consumers who benefit* v_m^l *from a cheap treatment, the expert will either set* $p_m = v_m^h$ *and only serve consumers of type* v_m^h *or set* $p_m = v_m^l$ *and serve both types. This decision clearly depends on the relative magnitude of each consumer group. However, in either case the price* p_s *for an expensive treatment and hence the expert's cheating behaviour will not be affected. As a result, the insights from Proposition 2 will remain valid when we consider hetereogeity of benefits from a cheap treatment.*

4 Consumer information on the type of the problem

In the previous section, we have considered the implications of introducing consumer information on the true benefit from an expensive treatment. Perhaps an equally interesting task is to investigate a model where some consumers receive information on the type of their problem. Accordingly, in this section, we consider a variant of our basic model and analyze the possibility that before visiting the expert, a fraction λ of consumers observe an informative signal on whether their problem is serious or minor. As before, the *ex ante* probability of a serious problem is $\alpha \in (0, 1)$. For simplicity, we now assume that all consumers benefit v_s when a serious problem is treated where $v_s > v_m > 0$.

she did, since her strategy cannot be conditioned on the consumer's information status, the expert would lose all uninformed consumers.

The information signal $\tilde{z}$ can take two values: A good signal $(z = g)$ indicates that the problem is more likely to be minor, whereas a bad signal $(z = b)$ indicates that the problem is more likely to be serious. In particular, the precision of the signal, denoted by ϕ is defined as

$$\phi \equiv \Pr(z = b|\omega = s) = \Pr(z = g|\omega = m) \in (\frac{1}{2}, 1).$$

For those customers who receive a signal prior to visiting the expert, the posterior beliefs are given by

$$\alpha_g \equiv \Pr(s|g) = \tfrac{\alpha(1-\phi)}{(1-\alpha)\phi+\alpha(1-\phi)} \quad \text{and} \quad \alpha_b \equiv \Pr(s|b) = \tfrac{\alpha\phi}{\alpha\phi+(1-\alpha)(1-\phi)},$$

whereas a customer with no signal still believes that his problem is serious with probability α. It is useful to emphasize that the signals are noisy. In particular, a consumer with a minor problem might arrive in the expert's office believing that his problem is serious if he had observed a signal $z = b$. Eveything else equal, such a pessimistic consumer seems more willing to accept an expensive treatment recommendation than a consumer who has received a good signal. This construction allows us to address whether the expert will exploit and cheat those consumers who already believe that their problem is serious with high probability.

As before, there are two potentially different scenarios to consider: the expert may or may not be able to distinguish informed consumers from uninformed ones. We analyze both cases below. In both cases, to rule out a trivial fixed price solution, we again assume that

$$\alpha_b v_s + (1 - \alpha_b)v_m < c_s.$$

Expert can distinguish informed and uninformed consumers. Suppose that the expert can perfectly distinguish whether a consumer has an informative signal or not, and the particular signal he has observed. This scenario is relevant in situations where the information signal is public for the expert to observe as well, such as some unpleasant noise coming from the engine of a car.[22] In this modified game, the expert can condition her recommendation strategy on the type of the problem and also on whether the consumer has a good or a bad signal, or he is uninformed. A mixed strategy profile for the expert in a recommendation sub-game (p_m, p_s) is now given by the probabilities $\{\rho_i^t, \beta_i^t, 1 - \beta_i^t - \rho_i^t\}$ for $i \in \{m, s\}$ and $t \in \{g, b, n\}$ where n stands for uninformed. A mixed strategy profile for a consumer of type $t \in \{g, b, n\}$ is given by the probability γ_i^t of accepting a recommendation $i \in \{m, s\}$. For notational convenience, let us define $\alpha_n \equiv \alpha$. If a consumer of type $z \in$

[22]Of course, in this scenario we are assuming that the mechanic/expert is also able to tell whether the consumer can or cannot recognize the unpleasant noise from the engine as an "informative" signal.

$\{g, b, n\}$ accepts an expensive treatment, his expected payoff is

$$V_z^s = \frac{\alpha_z \beta_s^z v_s + (1-\alpha_z)\beta_m^z v_m}{\alpha_z \beta_s^z + (1-\alpha_z)\beta_m^z} - p_s.$$

It can again be shown that in any equilibrium we must have $(p_m, p_s) \in [c_m, v_m] \times [c_s, v_s]$, and in these recommendation sub-games we have $\beta_s^t = 1$, $\gamma_i^z \in (0,1)$ and $\beta_m^z \in (0,1)$ for $z \in \{g, b, n\}$. The mixing probabilities of the expert can then be computed as

$$\beta_m^z = \frac{\alpha_z (v_s - p_s)}{(1-\alpha_z)(p_s - v_m)} \text{ for } z \in \{g, b, n\}.$$

It is useful to note that for any given price $p_s < v_s$, we have $\beta_m^b > \beta_m^n > \beta_m^g$. This observation implies that a consumer with a bad signal would tolerate a higher cheating probability to accept an expensive treatment than an uninformed consumer or a consumer with a good signal. However, when the expert charges the full valuation by setting $p_s = v_s$, regardless of their information status, all consumers only accept with positive probability if the expert is always truthful; *i.e.*, only when $\beta_m^b = \beta_m^n = \beta_m^g = 0$.

On the other hand, for the expert to always mix between recommending the expensive and cheap treatments when the problem is minor, all consumers must be accepting an expensive treatment with probability:

$$\gamma_s^t = \gamma_s = \frac{p_m - c_m}{p_s - c_m} \text{ for } t \in \{g, b, n\}.$$

This common acceptance probability by all consumers regardless of their information status implies that by increasing the price p_s, the expert is reducing her acceptance rate in a uniform manner across all types of consumers, but this reduction in demand is more than compensated for by the higher profit margin $p_s - c_s$. Indeed, using the above indifference conditions, one can show that the expert's expected profit function is given by

$$\Pi(p_m, p_s) = \alpha \left(\frac{p_m - c_m}{p_s - c_m} \right) (p_s - c_s) + (1-\alpha)(p_m - c_m)$$

which is increasing in p_s. This observation suggests that the expert will set $p_s^* = v_s$ and be truthful to everyone regardless of their information status.

Expert cannot distinguish informed and uninformed consumers. Consider now the possibility that the expert cannot distinguish whether a consumer is informed or not, and the particular signal he might have observed. Accordingly, the expert can only condition her recommendation strategy on the type of the problem. A mixed strategy profile for the expert in a recommendation subgame (p_m, p_s) is now given by the probabilities $\{\rho_i, \beta_i, 1-\beta_i-\rho_i\}$ for $i \in \{m, s\}$. A mixed strategy profile for a consumer of type $z \in \{g, b, n\}$ is described by the probability γ_i^z of accepting a recommendation $i \in \{m, s\}$. If a consumer of type $z \in \{g, b, n\}$ accepts an expensive treatment, his expected payoff will be

$$V_z^s = \frac{\alpha_z \beta_s v_s + (1-\alpha_z)\beta_m v_m}{\alpha_z \beta_s + (1-\alpha_z)\beta_m} - p_s$$

In Appendix A.3, we show that in any equilibrium we again must have $(p_m, p_s) \in [c_m, v_m] \times [c_s, v_s]$, $\beta_s = 1$ and $\gamma_m^z = 1$ for $z \in \{g, b, n\}$. Furthermore, for a given $p_m \in [c_m, v_m]$ and $p_s \in [c_s, v_s]$, a consumer of type $z \in \{g, b, n\}$ sets

$$\gamma_s^z > 0 \text{ if } \beta_m < A_z \equiv \frac{\alpha_z}{1-\alpha_z} \frac{v_s - p_s}{p_s - v_m}, \text{ and } \gamma_s^z = 0 \text{ if } \beta_m \geq A_z$$

It can also be shown that in any equilibrium, the expert will always follow a recommendation strategy with $\beta_m \in [A_g, A_b]$. Accordingly, the expert must be indifferent between recommending the expensive and cheap treatments when the problem is minor, and hence we must have

$$\gamma_s^T \equiv \lambda[\phi \gamma_s^g + (1-\phi)\gamma_s^b] + (1-\lambda)\gamma_s^n = \frac{p_m - c_m}{p_s - c_m}$$

But note that the expert's total acceptance rate for an expensive treatment recommendation is again given by the ratio $(p_m - c_m) / (p_s - c_m)$. Indeed, the *ex ante* profit function for the expert is identical to the one above. Hence, the unique equilibrium price is again given by $p_s^* = v_s$ and the expert will be truthful to all consumers.

The following proposition formally establishes that the equilibrium outcomes are the same when the expert can and cannot identify informed and uninformed consumers; and they both involve no cheating.

Proposition 4. *Suppose a fraction $\lambda > 0$ of consumers observe an informative signal on whether their problem is serious or minor. The equilibrium outcome is unique and is the same when the expert can or cannot identify informed and uniformed consumers. In the unique equilibrium outcome, the expert sets $p_m^* = v_m$ and $p_s^* = v_s$ and is always truthful to all types of consumers. All consumers accept a cheap treatment with probability one. All consumers accept an expensive treatment with a positive probability. The expert never refuses to treat any consumer.*

Proof. See Appendix A.3. □

The intuition for the above result is as follows. Again, the expert's expected profit is increasing in the price of the expensive treatment, and hence she sets the price equal to the consumers' benefit from having a serious treatment fixed. However, at that maximum possible price, regardless of their information and initial beliefs that they have a serious problem, all consumers reject with certainty if the expert cheats with a positive probability, which induces the expert to always tell the truth to all types of consumers. To sustain truthtelling, the consumers reject expensive treatment recommendations with a common positive probability.

5 Conclusions

In this paper, we contribute to the literature on credence goods by analyzing the implications of consumer information in a market for expert services. An important yet somewhat overlooked feature of these markets is the consumer heterogeneity in expertise and information regarding the service that the expert provides. Perhaps, this lack of attention to the implications of consumer information in expert services stems from the implicit and widely unquestioned assumption that marks most conventional thinking: fraudulent experts are likely to target ignorant and uninformed consumers as their victims. By identifying what drives expert cheating, our analysis questions this folk wisdom and shows that it is somewhat misguided.

In our basic model, we consider the case when prior to visiting an expert, a certain fraction of consumers learn their true benefit from an expensive treatment. We show that when all consumers are uninformed about their true benefit, the unique equilibrium involves no cheating. With some consumers informed about their true benefit from an expensive treatment, there is no equilibrium outcome in which the expert only cheats uninformed consumers. Depending on parameter values, the unique equilibrium involves one of the following three outcomes: (i) the expert only cheats informed high-valuation consumers; (ii) the expert cheats informed high-valuation and uninformed consumers, but is truthful to informed low-valuation consumers; and (iii) the expert is truthful to all types of consumers. Accordingly, it is the informed high-valuation consumers, and not the uninformed consumers, that are the most frequent victims of expert cheating.

Another widely accepted assumption regarding the inefficiencies created by fraudelent expert behaviour is that more information on the consumer side decreases the inefficiency of the market outcome. A surprising insight of our analysis is that this is not necessarily the case. In terms of the efficiency loss due to foregone but required treatments, all types of equilibrium outcomes that emerge when some consumers are informed about their true expensive treatment benefit exhibit *more* efficiency loss than the truthful equilibrium that arises when all consumers are uninformed.

We also analyze the case when some consumers receive noisy information signals about the type of their problem. This information structure implies that a consumer may believe that his problem is likely to be serious, whereas it is only minor. As such, one would expect the expert to target such pessimistic consumers who are more likely to accept a fraudulent expensive treatment recommendation. However, we show that, regardless of whether the expert can or cannot identify informed and uninformed consumers, the unique equilibrium in this case involves no cheating. The intuition for this truthfulness result is that the monopolist expert finds it optimal to charge the highest possible price for an expensive treatment. At this price, regardless of their information and beliefs about their problem, all types of consumers accept the expert's expensive treatment recommendation with some positive probability only when the expert is always truthful.

References

[1] Alger, I. and Salanié, F. 2006. "A Theory of Fraud and Over-Consumption in Experts Markets." Journal of Economics and Management Strategy, **15**(4), 853-81.

[2] Biglaiser, G. 1993. "Middlemen as Experts." *RAND Journal of Economics*, **24**(1), 212–23.

[3] Darby, M.R. and Karni, E. 1973. "Free Competition and Optimal Amount of Fraud." *Journal of Law and Economics*, **16**, 67-88.

[4] De Jaegher, K. and Jegers, M. 2001. "The Physician–Patient Relationship as a Game of Strategic Information Transmission." *Health Economics*, **10**(7), 651–68.

[5] Dranove, D. 1988. "Demand Inducement and the Physician/Patient Relationship." *Economic Inquiry*, **26**, 281–98.

[6] Dulleck, U. and Kerschbamer, R. 2006. "On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods." *Journal of Economic Literature*, **44**, 5–42.

[7] Emons, W. 2001. "Credence Goods Monopolists." *International Journal of Industrial Organization*, **19**, 375–389.

[8] Emons, W. 1997. "Credence Goods and Fraudulent Experts." *RAND Journal of Economics*, **28**(1), 107–19.

[9] Eső, P. and Szentes, B. 2007. "The Price of Advice." *RAND Journal of Economics*, **38**, 863-80.

[10] Fong, Y-F. 2005. "When Do Experts Cheat and Whom Do They Target." *RAND Journal of Economics*, **36**, 113–130.

[11] Pesendorfer, W. and Wolinsky, A. 2003. "Second Opinions and Price Competition: Inefficiency in the Market for Expert Advice." *Review of Economic Studies*, **70**(2), 417–37.

[12] Pitchik, C. and Schotter, A. 1987. "Honesty in a Model of Strategic Information." *American Economic Review*, **77**(2), 815–829.

[13] Schneider, H. 2006. "Agency Problems and Reputation in Expert Services: Evidence from Auto Repair." Johnson School of Management, Cornell University.

[14] Sülzle, K. and Wambach, A. 2005. "Insurance in a Market for Credence Goods." *Journal of Risk and Insurance*, **72**(1), 159–76.

[15] Taylor, C. R. 1995. "The Economics of Breakdowns, Checkups, and Cures." *Journal of Political Economy*, **103**(1), 53–74.

[16] Wolinsky, A. 1993. "Competition in a Market for Informed Experts' Services." *RAND Journal of Economics*, **24**, 380–398.

A Proofs of Results Not Given in the Main Text

A.1 Proof of Proposition 2

The proof of this result proceeds in a number of steps. We begin by restricting the set of prices which are possible in equilibrium, then on this restricted set, we rule out pure strategy equilibria. Next, for each set of prices, we solve for the unique mixed strategy equilibrium of the subsequent subgame and derive an expression for expected profits of the expert. Finally, we optimise over the set of feasible prices and show that each of the three pricing outcomes discussed in the main body of the text can arise depending on the parameters of the model.

Step 1: Restricting the set of possible prices. Obviously if $p_s > v_s^h$, then all consumers will reject an expensive treatment with probability 1. Similarly, if $p_m > v_m$, all consumers will reject a minor treatment with probability 1. Next, if $p_m < c_m$, the expert would refuse to treat the consumer (because doing so would generate a loss). Therefore, $p_m \in [c_m, v_m]$. Finally, if $p_s < c_s$, the expert would either provide the minor treatment (which will be rejected by consumers if $p_s \in (v_m, c_s)$) or refuse to treat the consumer. Therefore, in any equilibrium $(p_m, p_s) \in [c_m, v_m] \times [c_s, v_s^h]$.

Step 2: Deriving best responses in any pricing subgame such that $(p_m, p_s) \in [c_m, v_m) \times [c_s, v_s^h)$**.** First observe that $\gamma_m^z = 1$ for all $z \in \{h, l, n\}$, so that all consumers will accept with probability 1 the minor treatment. Second, because $p_m \leq v_m < c_s$ and the limited liability assumption, regardless of the type of consumer, the expert would never recommend a minor treatment when the problem is serious, since otherwise, she would be required to provide the serious treatment as a loss. Third, even though the expert knows whether an uninformed consumer is a high or a low type, since the uninformed consumers only condition their response on the proposed treatment, we have $\beta_m^{h,N} = \beta_m^{l,N} = \beta_m^N$.

We will break our subsequent analysis of the game into three regions, depending on p_s.

Case (a): $p_s \in [c_s, v_s^l)$.

We claim that $\gamma_s^z \in (0,1)$ for $z \in \{h,l,n\}$ and $\beta_m^{t,k} \in (0,1)$ for $t \in \{h,l\}$ and $k \in \{I,N\}$. Consider an uninformed consumer and suppose to the contrary that $\gamma_s^n = 1$. In this case, since such consumers accept with probability 1 a serious treatment, the expert will always cheat these consumers; *i.e.*, $\beta_m^N = 1$. However, in this case, the expected benefit to an uninformed consumer is $\alpha\bar{v}_s + (1-\alpha)v_m < c_s < p_s$, which is a contradiction. On the other hand, if $\gamma_s^n = 0$, it must be that $\beta_m^N = 0$, so that the expert never cheats an uninformed consumer, but then in this case the benefit from accepting a serious treatment is $\bar{v}_s > p_s$, a contradiction. Therefore, $\beta_m^N \in (0,1)$ and $\gamma_s^n \in (0,1)$. The argument for informed low and high type consumers is identical and is, therefore, omitted.

Having ruled out pure strategy equilibria, we now derive the equilibrium mixing probabilities when $p_s \in [c_s, v_s^l)$. In order for an uninformed consumer to be indifferent between accepting and rejecting the expensive treatment, it must be that:

$$\beta_m^N = \frac{\alpha(\bar{v}_s - p_s)}{(1-\alpha)(p_s - v_m)}. \tag{1}$$

Similarly, the indifference conditions for informed high and low types are:

$$\beta_m^{l,I} = \tfrac{\alpha(v_s^l - p_s)}{(1-\alpha)(p_s - v_m)} \quad \text{and} \quad \beta_m^{h,I} = \tfrac{\alpha(v_s^h - p_s)}{(1-\alpha)(p_s - v_m)}. \tag{2}$$

In order for the expert to mix between recommending the expensive or cheap treatments to a type $z \in \{h,l,n\}$ consumer with a minor problem, we must have, for all $z \in \{h,l,n\}$:

$$\gamma_s^z = \frac{p_m - c_m}{p_s - c_m}. \tag{3}$$

Given the recommendation strategies and the acceptance probabilities, we can write the expected profit to the expert for prices $p_m \in [c_m, v_m)$ and $p_s \in [c_s, v_s^l)$ as:

$$\Pi(p_m, p_s) = \alpha\pi_1 + (1-\alpha)[\lambda\pi_2 + (1-\lambda)\pi_3]$$

where

$$\begin{aligned}
\pi_1 &= \left[\lambda(\theta\gamma_s^h + (1-\theta)\gamma_s^l) + (1-\lambda)\gamma_s^n\right](p_s - c_s) \\
\pi_2 &= \theta\left[\beta_m^{h,I}\gamma_s^h(p_s - c_m) + (1-\beta_m^{h,I})(p_m - c_m)\right] \\
&+ (1-\theta)\left[\beta_m^{l,I}\gamma_s^l(p_s - c_m) + (1-\beta_m^{l,I})(p_m - c_m)\right] \\
\pi_3 &= \beta_m^N\gamma_s^n(p_s - c_m) + (1-\beta_m^N)(p_m - c_m).
\end{aligned}$$

Making use of (1), (2) and (3), we can simplify the above equations to:

$$\pi_1 = \tfrac{p_m - c_m}{p_s - c_m}(p_s - c_s) \quad \pi_2 = p_m - c_m \quad \pi_3 = p_m - c_m.$$

Consequently, we have that:

$$\Pi(p_m, p_s) = \left[\alpha\frac{p_s - c_s}{p_s - c_m} + (1-\alpha)\right](p_m - c_m). \tag{4}$$

Case (b): $p_s \in (v_s^l, \bar{v}_s)$.

In this case, notice that since $p_s > v_s^l$, it must be that $\gamma_s^l = 0$, so that all informed low types reject with probability 1. Therefore, it must also be that $\beta_m^l(I) = 0$. In the same way as in Case (a), it can be shown that $\gamma_s^z \in (0,1)$ for $z \in \{n, h\}$ and $\beta_m^h(I)$, $\beta_m(N) \in (0,1)$.

The acceptance probabilities for the uninformed and informed high type consumers can be written as $\gamma_s^z = \frac{p_m - c_m}{p_s - c_m}$ for $z \in \{n, h\}$, while the cheating probabilities for the expert are:

$$\beta_m(N) = \tfrac{\alpha(\bar{v}_s - p_s)}{(1-\alpha)(p_s - v_m)} \quad \text{and} \quad \beta_m^h(I) = \tfrac{\alpha(v_s^h - p_s)}{(1-\alpha)(p_s - v_m)}.$$

From this, after some simplifications, we can get an expression for the expected profits of the expert in the price range under consideration:

$$\Pi(p_m, p_s) = \left[\alpha(1 - \lambda + \theta\lambda)\frac{p_s - c_s}{p_s - c_m} + (1-\alpha)\right](p_m - c_m). \tag{5}$$

Case (c): $p_s \in (\bar{v}_s, v_s^h)$.

The techniques are by now familiar and so we only sketch the details. In this case, since $p > \bar{v}_s$, both informed low type and uninformed consumers will reject the serious treatment with probability 1. Therefore, $\gamma_s^l = \gamma_s^n = \beta_m^{l,I} = \beta_m^N = 0$. In this case, the informed high type will accept a serious offer with probability $\gamma_s^h = \frac{p_m - c_m}{p_s - c_m}$, while the export will cheat the informed high type with probability $\beta_m^{h,I} = \frac{\alpha(v_s^h - p_s)}{(1-\alpha)(p_s - v_m)}$. After some simplification, for prices in the relevant region, we are able to write the expected profit function of the expert as:

$$\Pi(p_m, p_s) = \left[\alpha\theta\lambda\frac{p_s - c_s}{p_s - c_m} + (1-\alpha)\right](p_m - c_m). \tag{6}$$

Step 3: Solving for the optimal price. Examining (4) - (6), it is easily seen that the expressions for expected profits are increasing in both p_m and p_s. Therefore, in all cases, $p_m = v_m$, while for the serious treatment we have:

Case	Optimal Price p_s	Maximized Profit
(A) $p_s \in [c_s, v_s^l)$	$p_s = v_s^l$	$\Pi_1^* = \left[\alpha\frac{v_s^l - c_s}{v_s^l - c_m} + (1-\alpha)\right](v_m - c_m)$
(B) $p_s \in (v_s^l, \bar{v}_s)$	$p_s = \bar{v}_s$	$\Pi_2^* = \left[\alpha(1-\lambda+\lambda\theta)\frac{\bar{v}_s - c_s}{\bar{v}_s - c_m} + (1-\alpha)\right](v_m - c_m)$
(C) $p_s \in [\bar{v}_s, v_s^h)$	$p_s = v_s^h$	$\Pi_3^* = \left[\alpha\lambda\theta\frac{v_s^h - c_s}{v_s^h - c_m} + (1-\alpha)\right](v_m - c_m)$

Step 4: Each of the three possible cases may be optimal. One can show that:

(i) The Type I equilibrium outcome will arise when:

$$\begin{aligned} 1-\lambda+\lambda\theta &< \frac{v_s^l - c_s}{v_s^l - c_m}\cdot\frac{\bar{v}_s - c_m}{\bar{v}_s - c_s} \\ \lambda\theta &< \frac{v_s^l - c_s}{v_s^l - c_m}\cdot\frac{v_s^h - c_m}{v_s^h - c_s}. \end{aligned}$$

(ii) The Type II equilibrium will arise when:

$$\begin{aligned} 1-\lambda+\lambda\theta &> \frac{v_s^l - c_s}{v_s^l - c_m}\cdot\frac{\bar{v}_s - c_m}{\bar{v}_s - c_s} \\ \frac{1-\lambda+\lambda\theta}{\lambda\theta} &> \frac{v_s^h - c_s}{v_s^h - c_m}\cdot\frac{\bar{v}_s - c_m}{\bar{v}_s - c_s}. \end{aligned}$$

(iii) the Type III equilibrium outcome will arise when:

$$\begin{aligned} \lambda\theta &> \frac{v_s^l - c_s}{v_s^l - c_m}\cdot\frac{v_s^h - c_m}{v_s^h - c_s} \\ \frac{1-\lambda+\lambda\theta}{\lambda\theta} &< \frac{v_s^h - c_s}{v_s^h - c_m}\cdot\frac{\bar{v}_s - c_m}{\bar{v}_s - c_s}. \end{aligned}$$

The example reported in Figure 1 shows that all three equilibrium outcomes are possible for a specific set of parameter values. Of course, by continuity, all three outcomes will arise for an open set of parameter values. □

A.2 Proof of Proposition 3

By Corollary 1 we have that the efficiency loss in the case in which $\lambda = 0$ so that all consumers are uninformed is given by:

$$EL_{\lambda=0} = \alpha \left(\frac{\bar{v}_s - v_m}{\bar{v}_s - c_m} \right) (\bar{v}_s - c_s).$$

We only provide the comparison for Type III outcome since the comparison for Type I and Type II outcomes has been provided in the main text. Recall that for the Type III equilibrium outcome, the expert is truthful so the only source of inefficiency arises due to serious treatments which are rejected by consumers. In this outcome, all informed low types and uninformed consumers reject the expensive treatment with probability 1 (due to the high price), while informed high types probabilistically reject (to discipline the expert). In this case, the efficiency loss may be written as:

$$EL^{III}_{\lambda>0} = \alpha \left[\lambda \left(\theta(1-\gamma)(v^h_s - c_s) + (1-\theta)(v^l_s - c_s) \right) + (1-\lambda)(\bar{v}_s - c_s) \right]$$

where $1-\gamma = \frac{v^h_s - v_m}{v^h_s - c_m}$ is the rejection rate for the expensive treatment by informed high types. Again, after some effort, we may re-write the efficiency loss as:

$$EL^{III}_{\lambda>0} = \alpha \left(\frac{v^h_s - v_m}{v^h_s - c_m} \right) (\bar{v}_s - c_s) + \alpha \left(\lambda\theta\gamma(v^l_s - c_s) + (1-\lambda)\gamma(\bar{v}_s - c_s) \right).$$

It is apparent that $\alpha \left(\frac{v^h_s - v_m}{v^h_s - c_m} \right) (\bar{v}_s - c_s) > EL_{\lambda=0}$, which means that $EL^{III}_{\lambda>0} > EL_{\lambda=0}$, which completes the proof. $\square$

A.3 Proof of Proposition 4

Expert cannot identify uninformed and informed consumers. Similar to the proof of Proposition 2, it is possible to show that the prices, (p_m, p_s) must lie in the set $[c_m, v_m] \times [c_s, v_s]$. Moreover, one can also rule out the existence of pure strategy equilibria in which the expert is either always truthful or always dishonest. Therefore, given the expressions in the text, one can write the expert's expected profit function for $(p_m, p_s) \in [c_m, v_m) \times [c_s, v_s)$ as

$$\Pi(p_m, p_s) = \alpha\pi_1(p_s - c_s) + (1-\alpha)[\lambda\pi_2 + (1-\lambda)\pi_3]$$

where

$$\begin{aligned}
\pi_1 &= \left[\lambda(\phi\gamma_s^b + (1-\theta)\gamma_s^g) + (1-\lambda)\gamma_s^n\right] \\
\pi_2 &= \phi\left[\beta_m^g\gamma_s^g(p_s - c_m) + (1-\beta_m^g)(p_m - c_m)\right] \\
&+ (1-\phi)\left[\beta_m^b\gamma_s^b(p_s - c_m) + (1-\beta_m^b)(p_m - c_m)\right] \\
\pi_3 &= \beta_m^n\gamma_s^n(p_s - c_m) + (1-\beta_m^n)(p_m - c_m)
\end{aligned}$$

and it can be shown that $\pi_1 = \frac{p_m - c_m}{p_s - c_m}$ and $\pi_2 = \pi_3 = p_m - c_m$. Consequently, we have

$$\Pi(p_m, p_s) = \alpha\left(\frac{p_s - c_s}{p_s - c_m}\right)(p_m - c_m) + (1-\alpha)(p_m - c_m).$$

which is maximized at $p_m^* = v_m$ and $p_s^* = v_s$. At these prices, we have:

$$\beta_m^z = 0 \text{ and } \gamma_s^z = \frac{v_m - c_m}{v_s - c_m} \text{ for } z \in \{g, b, n\}.$$

Expert cannot identify uninformed and informed consumers. Consider now the case in which the expert cannot distinguish whether a consumer has an informative signal or not, and the particular signal he might have observed. For notational convenience, let us again define $\alpha_n \equiv \alpha$. If a consumer of type $z \in \{g, b, n\}$ accepts an expensive treatment, his expected payoff will be

$$V_z^s = \frac{\alpha_z\beta_s v_s + (1-\alpha_z)\beta_m v_m}{\alpha_z\beta_s + (1-\alpha_z)\beta_m} - p_s.$$

One can again show that in any equilibrium we must have $(p_m, p_s) \in [c_m, v_m] \times [c_s, v_s]$, $\beta_s = 1$ and $\gamma_m^z = 1$ for $z \in \{g, b, n\}$. For a given $p_m \in [c_m, v_m]$ and $p_s \in [c_s, v_s]$, a consumer of type $z \in \{g, b, n\}$ sets

$$\gamma_s^z > 0 \text{ if } \beta_m < A_z \equiv \frac{\alpha_z}{1-\alpha_z}\frac{v_s - p_s}{p_s - v_m}, \text{ and } \gamma_s^z = 0 \text{ if } \beta_m \geq A_z.$$

Furthermore, we have $A_g < A_n < A_b$ for $p_s \in [c_s, v_s)$ and $A_g = A_n = A_b = 0$ for $p_s = v_s$. In any equilibrium, the expert will always follow a recommendation strategy with $\beta_m \in [A_g, A_b]$. To see this, note that any $\beta_m > A_b$ yields $\gamma_s^z = 0$ for $z \in \{g, b, n\}$ which implies we must have $\beta_m = 0$, a contradiction. Also if $\beta_m < A_g < 1$, we have $\gamma_s^z = 1$ for $z \in \{g, b, n\}$ which implies a best response $\beta_m = 1$, a contradiction. Accordingly, the expert must be indifferent between recommending the expensive and cheap treatments when the problem is minor, and hence we must have

$$p_m - c_m = \gamma_s^T (p_s - c_m)$$

where the total acceptance probability γ_s^T of a serious treatment is given by

$$\gamma_s^T = \lambda[\phi\gamma_s^g + (1-\phi)\gamma_s^b] + (1-\lambda)\gamma_s^n.$$

The expert's expected profit function for $(p_m, p_s) \in [c_m, v_m] \times [c_s, v_s]$ then takes the form

$$\Pi(p_m, p_s) = \alpha \left(\frac{p_m - c_m}{p_s - c_m} \right) (p_s - c_s) + (1-\alpha)(p_m - c_m)$$

which has a maximum at $p_m^* = v_m$ and $p_s^* = v_s$. But then we have $A_g = A_n = A_b = 0$ and $\beta_m = 0$; *i.e.,* the expert is always truthful. □

B Some consumers are informed but the expert cannot distinguish informed and uninformed

In this section, we consider a variation on our model of Section 3.2. In particular, we study the case in which the expert is not able to distinguish between whether the consumer is informed or uninformed about his true valuation for having a serious problem repaired. Therefore, her recommendation strategy may only be a function of the consumer's type $t \in \{h, m\}$, which the expert still observes. Exactly as in our earlier analysis, the limited liability assumption ensures that the expert will always recommend the serious treatment when the problem is serious: that is $\beta_s^t = 1$ for all $t \in \{h, l\}$. Moreover, an identical argument to that in the proof of Proposition 2 allows us to restrict attention to price intervals $(p_m, p_s) \in [c_m, v_m] \times [c_s, v_s^h]$.

In this setting, the fact that the expert cannot distinguish uninformed from informed leads to an extra complication. Suppose that $p_s \in (\bar{v}_s, v_s^h)$ so that both informed low types and uninformed consumers reject the serious treatment with probability 1, and suppose that the expert observes that the consumer's problem is minor. If she recommends the expensive treatment, the her expected profits are $\lambda\gamma_s^h(p_s - c_m)$, while if she recommends the minor treatment her profits are $p_m - c_m$. If λ is small enough, then it is possible that the expert *strictly* prefers to be truthful, even if the informed high type accepted the serious recommendation with certainty. Therefore, in some pricing subgames, it is not possible to rule out pure strategy equilibria.

As before, we break our analysis into three price regions and begin with $(p_m, p_s) \in [c_m, v_m] \times [c_s, v_s^l)$. In this region, the prices are low enough so that all types of consumers receive strictly positive surplus from having a serious problem repaired. It is easily shown that for all prices in this range, both the expert and all types of consumers must employ mixed strategies. For the expert, who observes a low type, to mix it must be that:

$$(1-\lambda)\gamma_s^n + \lambda\gamma_s^l = \frac{p_m - c_m}{p_s - c_m} \tag{7}$$

and similarly, for the expert, who observes a high type, to mix it must be that:

$$(1-\lambda)\gamma_s^n + \lambda\gamma_s^h = \frac{p_m - c_m}{p_s - c_m} \tag{8}$$

Consider now the consumers. For an informed consumer of type $t \in \{h, m\}$, to make the consumer indifferent between accepting and rejecting an expensive treatment, it must be that:

$$\beta_m^t = \frac{\alpha(v_s^t - p_s)}{(1-\alpha)(p_s - v_m)}$$

while for an uninformed consumer to be indifferent it must be that:

$$\theta\beta_m^h + (1-\theta)\beta_m^l = \frac{\alpha(\bar{v}_s - p_s)}{(1-\alpha)(p_s - v_m)}$$

That is, if informed high and low types are indifferent, then an uninformed consumer is automatically indifferent.

We may now write the expected profits of the expert when the prices are in the range $(p_m, p_s) \in [c_m, v_m] \times [c_s, v_s^l)$.

$$\Pi(p_m, p_s) = \alpha\pi_1 + (1-\alpha)[\theta\pi_2 + (1-\theta)\pi_3]$$

where

$$\begin{aligned}
\pi_1 &= (p_s - c_s)\left[\theta\left(\lambda\gamma_s^h + (1-\lambda)\gamma_s^n\right) + (1-\theta)\left(\lambda\gamma_s^l + (1-\lambda)\gamma_s^n\right)\right] \\
\pi_2 &= \beta_m^h\left(\lambda\gamma_s^h + (1-\lambda)\gamma_s^n\right)(p_s - c_m) + (1-\beta_m^h)(p_m - c_m) \\
\pi_3 &= \beta_m^h\left(\lambda\gamma_s^l + (1-\lambda)\gamma_s^n\right)(p_s - c_m) + (1-\beta_m^l)(p_m - c_m)
\end{aligned}$$

Making use of the indifference conditions, we can simplify the expression for expected profits to:

$$\Pi(p_m, p_s) = \left[\alpha\frac{p_s - c_s}{p_s - c_m} + 1 - \alpha\right](p_m - c_m)$$

and note that $\Pi(p_m, p_s)$ obtains a maximum over the relevant range of prices at $p_m^* = v_m$ and $p_s^* = v_s^l$. Therefore, when the expert optimally chooses prices in this range, she is honest to low types and cheats high types.

Next consider the price range $(p_m, p_s) \in [c_m, v_m) \times (v_s^l, \bar{v}_s)$. Here, we know that $\gamma_s^l = 0$. It is necessary to break our analysis into two sub-regions, depending upon the prices and the value of λ. To see this, recall (7). Substituting in for $\gamma_s^l = 0$, we see that if λ is sufficiently high it may be that $(1-\lambda)(p_s - c_m) < p_m - c_m$ so that even if an uninformed, though low valuation, consumer accepts the expensive treatment with probability 1, the expert still prefers to be truthful to low valuation consumers.

Therefore, first suppose that $(1-\lambda)(p_s - c_m) < p_m - c_m$. This immediately implies that $\beta_m^l = 0$. In this case, we will have $\beta_m^h = \frac{\alpha(v_s^h - p_s)}{(1-\alpha)(p_m - v_m}$, while $\gamma_s^n = 0$ and $\gamma_s^h = \frac{p_s - c_m}{\lambda(p_m - c_m)}$. One can show that the profit function of the expert in this price range can be written as:

$$\Pi(p_m, c_m) = \left[\alpha\theta \frac{p_s - c_s}{p_s - c_m} + (1-\alpha)\right](p_m - c_m)$$

which implies that $p_m^* = v_m$ and $p_s^* = \min\{\bar{v}_s, \max\{c_s, \frac{v_m - c_m}{\lambda} + c_m\}\}$.

Next, consider the case in which $(1-\lambda)(p_s - c_m) > p_m - c_m$. In this case, we cannot immediately conclude that $\beta_m^l = 0$; however, we claim that this must be so. If $\beta_m^l > 0$, it must be that $\gamma_s^n > 0$. Since $\gamma_s^n > 0$, if, additionally, $\gamma_s^h > 0$, then (8) shows that the expert strictly prefers to cheat the high types; *i.e.*, $\beta_m^h = 1$. However, we know this to be impossible. On the other hand, if $\gamma_s^h = 0$, then it must be that $\beta_m^h \geq \frac{\alpha(v_s^h - p_s)}{(1-\alpha)(p_s - v_m)}$. However, in order for the uninformed agents to accept a serious treatment, it must be that $\theta\beta_m^h + (1-\theta)\beta_m^l = \frac{\alpha(\bar{v}_s - p_s)}{(1-\alpha)(p_s - v_m)}$, and one can see that this is incompatible with the aforementioned restriction on β_m^h. That is, the uninformed consumer would strictly prefer to accept. Therefore, we have proven that $\beta_m^l = 0$.

Now for the expert to be indifferent between cheating and not cheating the high types, it must be that:

$$\lambda\gamma_s^h + (1-\lambda)\gamma_s^n = \frac{p_m - c_m}{p_s - c_m}$$

and given that $\beta_m^l = 0$, it must be that $\gamma_s^h = 1$ and $\gamma_s^n = \frac{1}{1-\lambda}\left[\frac{p_m - c_m}{p_s - c_m} - \lambda\right]$ and also that $\beta_m^h = \frac{\alpha(\bar{v}_s - p_s)}{\theta(1-\alpha)(p_s - v_m)}$.[23]

We can calculate the profits of the expert in this price range as:

$$\Pi(p_m, p_s) = \alpha\left[\frac{p_m - c_m}{p_s - c_m} - (1-\theta)\lambda\right](p_s - c_s) + (1-\alpha)(p_m - c_m)$$

The profit function, unlike other cases, is not strictly increasing in prices; however, its continuity assures us that a maximum exists. In general, for this case, p_s^* must be strictly less than $\bar{v}_s$. This follows because, in order to cheat the informed high types and also to insure that the uninformed consumers accept a serious recommendation with positive probability, the expert must leave some surplus to the uninformed consumers who, since the expert

[23]To be sure, one must check that when the expert cheats high types with probability β_m^h that the informed high types strictly prefer to accept the expensive treatment. One can easily verify that this is so.

cannot distinguish informed from uninformed, are also victims of expert cheating.

Finally, consider the case in which $(p_m, p_s) \in [c_m, v_m) \times (\bar{v}_s, v_s^h)$. In this case, we have that $\gamma_s^l = \gamma_s^n = 0$, which implies that $\beta_m^l = 0$. As with the previous pricing range, it is possible, if λ is sufficiently small for the equilibrium outcome to be in pure strategies. To see this, look at 8, recall that $\gamma_s^n = 0$ and observe that if λ is sufficiently small, then $p_m - c_m \geq \lambda(p_s - c_m)\}$; hence the expert prefers to be honest. Consider then the (possibly empty) set of prices $(p_m, p_s) \in \{[c_m, v_m) \times (\bar{v}_s, v_s^h)\} \cap \{(p_m, p_s) : p_m - c_m \geq \lambda(p_s - c_m)\}$. We can express the expected profits of the expert as:

$$\Pi(p_m, p_s) = \alpha\lambda\theta(p_s - c_s) + (1 - \alpha)(p_m - c_m)$$

It is then easily seen that $p_m^* = v_m$ and $p_s^* = \min\{v_s^h, \max\{c_s, \frac{v_m - c_m}{\lambda} + c_m\}\}$. Even though p_s^* may be strictly below v_s^h, the expert is still truthful because there are sufficiently few informed high types, which means that the expert does not want to risk having them reject due to dishonesty on her part.

Next consider the opposite case in which $(p_m, p_s) \in \{[c_m, v_m) \times (\bar{v}_s, v_s^h)\} \cap \{(p_m, p_s) : p_m - c_m < \lambda(p_s - c_m)\}$. Here, we will have the usual mixed strategy equilibrium in which $\gamma_s^h \in (0, 1)$ and $\beta_m^h \in (0, 1)$. The expected profits of the expert in this range of prices are:

$$\Pi(p_m, p_s) = \alpha\lambda\theta\gamma_s^h(p_s - c_s) + (1 - \alpha)((1 - \theta)(p_m - c_m) + \theta(\beta_m^h\lambda\gamma_s^h(p_s - c_m) + (1 - \beta_m^h)(p_m - c_m)))$$

which can be simplified to:

$$\Pi(p_m, c_m) = \left[\alpha\theta\frac{p_s - c_s}{p_s - c_m} + (1 - \alpha)\right](p_m - c_m)$$

where here we have, over the relevant range of prices, $p_s^* = v_s^h$ and $p_m^* = \min\{v_m, (1 - \lambda)c_m + \lambda v_s^h\}$. Observe that since $p_s^* = v_s^h$, there is no cheating in equilibrium.

Of course, since the two pricing regions are mutually exclusive, one of them must be empty and the other non-empty in any given pricing subgame. The first pricing region is likely to be non-empty when λ is relatively low, while when λ is relatively high, the second region is likely to be non-empty. In the former case, the expert strictly prefers to be honest, while in the latter, both the expert and the informed high type employ strictly mixed strategies.

To summarise, the unique subgame perfect equilibrium resembles the case discussed at length in the main body of the text in which the expert was able to distinguish informed from uninformed. In particular, there are three possible types of equilibrium outcomes (which, of course, depend upon the underlying parameter values) – two of which involve cheating, while one involves truth telling. The main difference between the two models is that here, since the expert cannot distinguish informed from uninformed, if she cheats she will cheat all high

Figure 2: Equilibrium Outcomes Over Possible Values of θ and λ

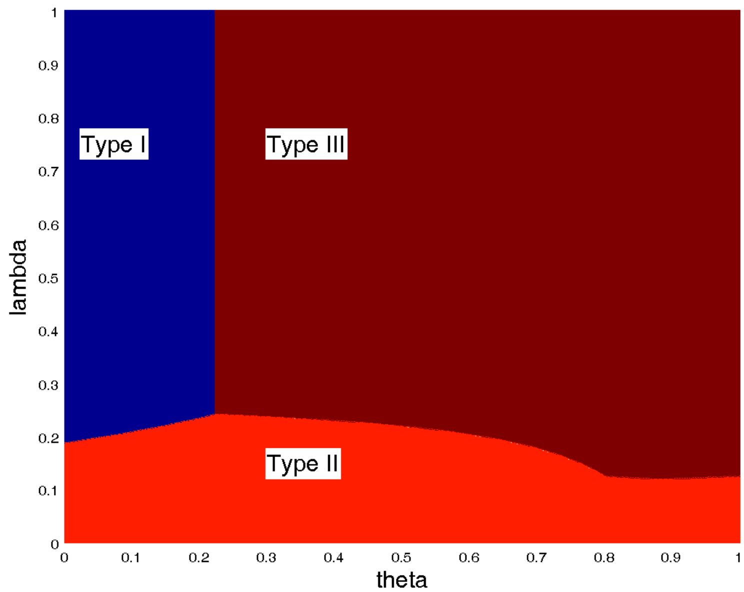

types, regardless of their information status (which is unknown to her). Similar to Figure 1, we can show in Figure 2 the sets of (θ, λ) values under which each type of equilibrium arises.[24] Interestingly, though not surprising, we see that the Type II equilibrium outcome appears to shrink at the expense of both the Type I and Type III outcomes. Unless the proportion of uninformed consumers is sufficiently high (*i.e.*, λ is low), the expert will find it more often optimal to set either a price at $p_s = v_s^l$ (Type I equilibrium outcome which keeps all consumers buying), or set a price high enough that both informed low types and uninformed consumers will always reject and the expert will have business with only informed high types (Type III equilibrium outcome).

[24]This figure is drawn using the same set of exogenous parameters as in Figure 1.