Consumer Information in a Market for Expert Services*

Kyle Hyndman[†]
Southern Methodist University

Saltuk Ozerturk[‡]
Southern Methodist University

May 14, 2011

Abstract

We present a model of credence goods in which the consumers are heterogenous in terms of the valuation they place for getting a serious problem fixed. We introduce consumer information into this framework by assuming that, prior to visiting an expert, some consumers receive an information signal about whether they have a serious or a minor problem. We show that when the fraction of consumers with low willingness to pay is sufficiently high, the expert does not cheat any low valuation consumer regardless of their information status, but cheats the high valuation consumers: those high-valuation consumers with bad signals are the most frequent victims of cheating, whereas those with good signals are the least likely victims. When the fraction of consumers with low willingness to pay is below a certain threshold, however, the unique equilibrium involves no cheating.

Keywords: Credence Goods; Expert Cheating; Consumer Information

*We would like to thank seminar participants at CUNY-Hunter College, Drexel University, Fordham University, Koç University and SUNY-Stony Brook for their valuable comments.

†*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. This informational asymmetry between experts and consumers creates obvious incentive problems: experts may attempt to over-treat 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:[1]

> 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 expert to behave opportunistically and cheat. In the literature on expert services, however, there has been no formal analysis of the implications of consumer information. In this paper, we theoretically explore the implications of consumer information on the expert's cheating behavior. In particular, we pose the question of whether uninformed consumers are indeed the most likely victims of expert cheating.

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

Following most existing models, in our analysis a consumer's problem can either be serious (requiring an expensive treatment) or minor (requiring a cheap treatment). *Ex ante*, the consumer does not know the type of the problem, whether it is serious or minor. The consumer also cannot *ex post* verify the actual treatment provided by the expert, the expensive or the cheap treatment, to solve the problem. Accordingly, the potential fraud we consider in this paper is one of overcharging: the expert may recommend and charge for an expensive treatment when the problem is minor and can be solved with a cheap treatment.

In our model, consumers are heterogenous in terms of the valuation they place for getting a serious problem fixed with an expensive treatment. Into this framework, we introduce consumer information by assuming that, 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 optimistic that their problem is minor and some will be more pessimistic. The pessimistic consumers believe quite strongly that their problem is serious when, in reality, it is actually minor. Since we are interested in describing how observable consumer characteristics affect the expert's cheating behavior, we mainly consider the case where the expert observes the type of consumer's problem, the particular information consumer has and the consumer's willingness to pay for having the serious problem fixed.

Our analysis illustrates that, as well as the consumer's information on the type of the problem, perhaps an equally, if not more, important aspect that derives the expert's cheating behavior is the wedge between a consumer's willingness to pay for having the serious problem fixed and the treatment price the expert is charging. In a setting where the expert endogenously sets treatment prices before a consumer's visit, we show that a consumer's willingness to accept an expensive treatment recommendation is driven by two considerations, (i) the wedge between their willingness to pay and the price the expert is charging for the expensive treatment, (ii) their beliefs on the probability of having a serious treatment. As long as the expensive treatment price is below their full valuation, all consumers, regardless of their particular information, tolerate to some cheating. Not surprisingly, *for the same wedge between valuation and price*, those with bad signals and hence pessimistic beliefs about their problem are the ones most tolerant for expert cheating, and hence are the most likely victims of cheating. However, when the monopolist expert sets the expensive treatment price equal to consumer's full valuation, all consumers, regardless of their particular information, become completely intolerant to cheating. This consumer behavior implies that from the expert's point of view, charging a higher price and cheating are substitutes: as the price increases consumers become less and less tolerant towards cheating, the extent of their tolerance driven by their beliefs (hence information) on the nature of their problem.

In this framework, we show that whether the equilibrium involves cheating or not depends crucially not on the fraction of informed consumers, but on the fraction of high versus low valuation consumers.

- When the fraction of consumers with low willingness to pay is sufficiently high, the expert sets the expensive treatment price equal to the full valuation of low value consumers, and does not cheat any low value consumer regardless of the information status. However, with the treatment price set below their valuation, the high valuation consumers enjoy a wedge between their valuation and the treatment price, and hence are tolerant to the expert cheating with positive probability. Among the high valuation consumers, those with bad signals are the most frequent victims of cheating, whereas those with good signals and optimistic beliefs about their problem are the least frequent victims.

- When the fraction of consumers with low willingness to pay is below a certain threshold, however, the unique equilibrium involves no cheating: the expert sets the expensive treatment price equal to the full valuation of high value consumers. In this equilibrium, the expert is truthful to all types of consumers regardless of their information status, however, those with low valuations are priced out of the expensive treatment market.

These results make an interesting contribution to the folk wisdom that we mentioned above. As well as a consumer's information status, which determines the consumer's beliefs about the seriousness of the problem, an important and typically observable characteristic that affect an expert's cheating incentives is the consumer's willingness to pay for getting a serious problem fixed. Our analysis suggests that cheating is likely to emerge in situations where the expert observes that the consumer finds the expensive treatment price relatively low with respect to his/her valuation. In those situations, the expert is more likely to take advantage of this wedge between price and valuation, as the expert understands that the consumer has some tolerance for being cheated. Provided that there is room for cheating due to this willingness to pay and price, the extent of the expert's cheating then depends on whether the consumer strongly believes to have a serious problem or not, as those with pessimistic beliefs due to their information are more likely to be cheated than uninformed consumers. In that respect, a consumer may benefit more from pretending to have a low valuation for having a serious problem fixed than pretending to know more about the nature of his/her problem.

Related Literature. In many real life examples, the quality and amount of information that consumers possess differ substantially. However, the existing literature on credence goods 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. To the best of our knowledge, ours is the first paper that considers a credence goods market with heterogeneously informed consumers.

The theoretical literature on credence goods is small but growing.[2] 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 Salanié (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 overtreatment 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 model with fixed exogenous prices and homogenous consumers, Pitchik and Schotter (1987) demonstrate a mixed strategy equilibrium that involves cheating.[3]

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

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

[3] 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. Sülzle and Wambach (2005) study the impact of variations in the degree of insurance on the amount of fraud in a physician-patient relationship. Dranove (1988) analyzes how demand inducement by physicians relates to the treatment price and other exogenous variables. 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. Eső and Szentes (2007) consider a setting where the 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 advice for the client's welfare. None of these papers address the implications of heterogenous consumer information.

the focus of our paper.

The rest of the paper proceeds as follows. In the next section, we lay out our model. Section 3 provides the analysis and contains our main results. Section 4 concludes. All proofs not presented in the text can be found in various appendices. The case when the expert can observe the consumer's information but not the willingness to pay is presented in Appendix B.

2 The Model

In this section, we describe a basic model of the credence good market that we use in the paper. We also summarise the information held by consumers and the expert, and describe the sequence of events.

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.[4] 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).

Heterogenous Consumer Valuations for Treatment of Serious Problems. We assume that a minor problem, if left untreated, will lead to a loss of $\ell_m > c_m > 0$ for all consumers. Accordingly, we refer to ℓ_m as the consumers' willingness to pay for getting a minor problem treated. With respect to the untreated serious problems, we assume that the consumers are heterogenous in the losses they incur. When their problem is serious,

[4]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 Salanié (2006), the consumers can partially verify the actual inputs the expert uses during her treatment.

depending on the consumer's type, the expensive treatment avoids a loss of either ℓ_s^h or ℓ_s^l with $\ell_s^h > \ell_s^l > \ell_m$ and $\ell_s^l > c_s$.[5] We refer to ℓ_s^i with $i \in \{l, h\}$ as the consumer of type i's willingness to pay for getting a serious problem fixed, and assume that all consumers know their willingness to pay for getting a serious problem fixed, but the expert may or may not observe a consumer's willingness to pay.

Consumer Information. Before visiting the expert, a fraction λ of consumers observe an informative signal on whether their problem is serious or minor. The 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 $z = b$. To rule out a trivial fixed price solution, we also assume that

$$\alpha_b \ell_s^h + (1 - \alpha_b)\ell_m < c_s. \tag{A1}$$

Expert's Information. The expert can perfectly observe whether the consumer has a serious problem or a minor problem. Since our main objective is to investigate whether the expert cheats uninformed consumers, we assume that the expert can also distinguish perfectly whether the consumer has received a good signal, a bad signal or no signal (uninformed). With respect to the consumer's valuation for getting a serious problem fixed, in the main body of the paper we analyze the case in which the expert also observes the consumer's valuation. The case in which the expert does not observe a consumer's valuation is analyzed in Appendix B. We also assume that the expert cannot price discriminate across consumers.

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

[5] Since $\ell_m > c_m$ and $\ell_s^l > c_s$, both problems are efficient to fix.

also decides whether a consumer with a serious problem is willing to pay ℓ_s^h or ℓ_s^l for getting a serious problem fixed. The consumers do not know if their problem is minor or serious, but they know whether they are of type ℓ_s^h or ℓ_s^l. A fraction λ of consumers also observe a signal $\tilde{z}$ which is informative about the type of their problem.

- 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 informed with a good or bad signal, or if the consumer is uninformed. The expert also observes if the consumer is of type ℓ_s^h or ℓ_s^l. 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

In this section, we derive the equilibrium when the expert can observe and condition her recommendation strategy on the information status of the consumer, as well as the consumer's willingness to pay. From the point of view of the expert, a consumer may be one of six different types: those consumers with low valuation ℓ_s^l and information status $z \in \{g, b, n\}$, and consumers with high valuation ℓ_s^h and information status $z \in \{g, b, n\}$. In what follows, we describe a consumer's type with the index $t \in T \equiv \{g_l, b_l, n_l, g_h, b_h, n_h\}$. To clarify this notation, the type b_l, for example, refers to a consumer who has observed a bad signal and who has a low valuation ℓ_s^l, whereas type n_h refers to a consumer who has not observed a signal and has a high valuation ℓ_s^h.

A mixed strategy profile for the expert in a recommendation sub-game (p_m, p_s) is given by the probabilities $\{\rho_i^t, \beta_i^t, 1 - \beta_i^t - \rho_i^t\}$ for $i \in \{m, s\}$ and $t \in T$ where ρ_i^t refers to the probability that the expert rejects to treat a consumer of type t when the underlying problem is $i \in \{m, s\}$.[6] For example, $\beta_m^{g_l}$ is the probability that the expert recommends an expensive

[6]It is useful to note that one cannot dispense with ρ_i^t and remove the possibility that the expert refuses to treat consumers from the expert's action set without loss of generality. In particular, since we must consider all pricing subgames, including those with $p_i < c_i$, allowing the expert to refuse to treat customers guarantees that she will earn non-negative profits. In other words, if we are to assume that the expert has

treatment to a consumer who has a minor problem, when the consumer has received a good signal and has a low valuation ℓ_s^l for getting a serious problem fixed. A mixed strategy profile for a consumer of type $t \in T$ is given by the probability γ_i^t of accepting a recommendation $i \in \{m, s\}$ by the expert.

We begin our analysis by restricting the set of prices which are possible in equilibrium.

Lemma 1. *In any equilibrium the expert makes non-negative profits. Furthermore, in any equilibrium where the expert makes positive profits, the set of prices must belong to the range* $p_m \in [c_m, \ell_m]$ *and* $p_s \in [c_s, \ell_s^h]$.

Proof. See Appendix A. □

This follows because the expert always has the option of refusing to treat consumers (which she will exercise if the price is below her cost of providing the recommended treatment). Moreover, prices cannot be too high since otherwise consumers will automatically reject the recommended treatment.

The next lemma establishes that the expert will always recommend an expensive treatment when the problem is serious, and the consumers always accept a cheap treatment recommendation with probability one regardless of their information status and willingness to pay.

Lemma 2. *In any subgame in the relevant price range* $p_m \in [c_m, \ell_m]$ *and* $p_s \in [c_s, \ell_s^h]$*, we have* $\gamma_m^t = 1$ *and* $\beta_s^t = 1$ *for all* $t \in T$.

Proof. See Appendix A. □

This result follows because of the limited liability assumption. Namely, since an expert must fix the problem, she will always provide the expensive repair for the serious problem, and will consequently never mis-report that a serious problem is minor. Consequently, the consumer will always accept a cheap treatment recommendation, provided that $p_m \leq \ell_m$.

We now rule out pure strategy equilibrium for the relevant price range $p_s \in [c_s, \ell_s^h]$ and $p_m \in [c_m, \ell_m]$ with the following lemma.

Lemma 3. *(i) For* $p_m \in [c_m, \ell_m]$ *and* $p_s \in [c_s, \ell_s^l]$*, we have* $\gamma_s^t \in (0,1)$ *and* $\beta_m^t \in (0,1)$ *for all* $t \in T$*. (ii) For* $p_m \in [c_m, \ell_m]$ *and* $p_s \in [c_s, \ell_s^h]$*, we have* $\gamma_s^t \in (0,1)$ *and* $\beta_m^t \in (0,1)$ *for* $t \in \{g_h, b_h, n_h\}$*, however* $\gamma_s^{g_l} = \gamma_s^{b_l} = \gamma_s^{n_l} = 0$ *and* $\beta_m^{g_l} = \beta_m^{b_l} = \beta_m^{n_l} = 0$.

no ability to reject consumers ex post, then we need to exogenously assume that the prices are always chosen as $p_i > c_i$ so that the expert does not lose money by accepting to treat a consumer. Allowing for the expert to reject consumers ex post simply guarantees that the experts ex ante sets the prices in the region $p_i > c_i$. This is the content of Lemma 1, below.

Proof. See Appendix A. □

The intuition is relatively straight-forward. If the expert never cheated, then, following an expensive treatment recommendation, those consumers for which the price is below their valuation will accept with probability 1. However, if consumers accept an expensive treatment recommendation with probability 1, then the expert will strictly prefer to cheat consumers. However, if the expert cheats with probability 1, then consumers strictly prefer to reject. Finally, if consumers always reject, the expert prefers to be honest. Thus, we have come full circle, which means that there cannot be a purse strategy equilibrium

Having ruled out pure strategy equilibria, we proceed by deriving the equilibrium mixing probabilities first for the price range $p_s \in [c_s, \ell_s^l]$. For the expert to mix between recommending the expensive or cheap treatments to a type $t \in T$ consumer with a minor problem, we must have

$$\gamma_s^t(p_s - c_m) = p_m - c_m \Rightarrow \gamma_s^t = \frac{p_m - c_m}{p_s - c_m}. \tag{1}$$

Now consider the consumer's acceptance strategy that will leave the expert indifferent between cheating and not cheating when the problem is minor. For a consumer with valuation l_s^l or l_s^h and with signal $z \in \{g,\, b,\, n\}$ to be indifferent between accepting and rejecting the expensive treatment, the expert's respective cheating probabilities must be given by

$$\beta_m^{z_l} = \frac{\alpha_z(l_s^l - p_s)}{(1-\alpha_z)(p_s - \ell_m)} \text{ for } z \in \{g, b, n\}. \tag{2}$$

$$\beta_m^{z_h} = \frac{\alpha_z(l_s^h - p_s)}{(1-\alpha_z)(p_s - \ell_m)} \text{ for } z \in \{g, b, n\}. \tag{3}$$

A few observations on these cheating probabilities are in order. First, consider the expert's cheating probability $\beta_m^{z_l}$ above that makes a low valuation consumer with information $z \in \{g,\, b,\, n\}$ indifferent. Note that, if the expert cheats that consumer with a probability strictly higher than $\beta_m^{z_l}$, then consumer of type z_l will always reject an expensive treatment recommendation, whereas if the expert cheats with a probability strictly less than $\beta_m^{z_l}$, this consumer will always accept an expensive treatment recommendation. In other words, the probability $\beta_m^{z_l}$ describes the tolerance of consumer of type z_l for being cheated. It is straightforward to see that, since $\alpha_g < \alpha_n < \alpha_b$, among the consumers with low valuation, the consumer with the bad signal (consumer type b_l) is the most tolerant for being cheated for a given price $p_s < l_s^l$, since this is the consumer who is the most pessimistic about her problem. Formally, we have

$$\beta_m^{b_l} > \beta_m^{n_l} > \beta_m^{g_l} \text{ for } p_s < l_s^l$$

and the tolerance for being cheated is increasing in the difference between her willingness to pay l_s^l and the price p_s expert charges. Intuitively, a consumer with a given belief α_z where $z \in \{g, b, n\}$ becomes more tolerant to be cheated as the price p_s for expensive treatment becomes smaller.

Note, however, that if $p_s = \ell_s^l$ and $\beta_m^{z_l} > 0$, then low-valuation consumers strictly prefer to reject an expensive treatment recommendation, regardless of their type. In other words, at the full valuation price $p_s = l_s^l$, all low valuation consumers (regardless of their information) lose all their tolerance for being cheated. For the expert to charge $p_s = l_s^l$ and get an expensive treatment recommendation accepted with a positive probability, the expert can never cheat these consumers. Formally,

$$\beta_m^{z_l} = 0 \text{ for } z \in \{g, b, n\} \text{ at } p_s = l_s^l$$

Similar observations hold for consumers with high willingness to pay as well. The probability $\beta_m^{z_h}$ above describes the tolerance of consumer of type z_h for being cheated. Among the consumers with high valuation, the consumer with the bad signal (consumer type b_h) is the most tolerant for being cheated for a given price $p_s < l_s^h$. Formally, we have

$$\beta_m^{b_h} > \beta_m^{n_h} > \beta_m^{g_h} \text{ for } p_s < l_s^h$$

and the tolerance for being cheated is increasing in the difference between the willingness to pay l_s^h and the price p_s expert charges. At the full valuation price $p_s = l_s^h$, all high valuation consumers (regardless of their information) lose all their tolerance for being cheated, and hence

$$\beta_m^{z_h} = 0 \text{ for } z \in \{g, b, n\} \text{ at } p_s = l_s^h$$

Given these recommendation strategies and the acceptance probabilities, one can now write the ex ante expected profit to the expert. The next lemma derives the expert's ex ante expected profit function:

Lemma 4. *(i) For the price range $p_m \in [c_m, \ell_m]$ and $p_s \in [c_s, \ell_s^l]$, the expert's ex ante expected profits is given by*

$$\Pi_1(p_m, p_s) = \alpha \left[\frac{p_m - c_m}{p_s - c_m}\right] (p_s - c_s) + (1 - \alpha)(p_m - c_m). \tag{4}$$

(ii) For the price range $p_m \in [c_m, \ell_m]$ and $p_s \in (\ell_s^l, \ell_s^h]$, the expert's ex ante expected profits is given by

$$\Pi_2(p_m, p_s) = \alpha\theta \left[\frac{p_m - c_m}{p_s - c_m}\right] (p_s - c_s) + (1 - \alpha)(p_m - c_m). \tag{5}$$

Proof. See Appendix A. □

Consider the expected profit function $\Pi_1(\cdot)$ in part (i) of Lemma 4 for the expensive treatment price range $p_s \in [c_s, \ell_s^l]$. In this price range, both high and low valuation consumers, regardless of their information status, can afford to accept an expensive treatment recommendation. The first term in $\Pi_1(\cdot)$ refers to the expected profit the expert makes from consumers with serious problems: such a consumer arrives with probability α, and always receives an expensive treatment recommendation, which he/she accepts with a probability $(p_m - c_m) / (p_s - c_m)$. Every time a consumer with a serious problem accepts an expensive treatment recommendation, the expert makes $(p_s - c_s)$. The second term in $\Pi_1(\cdot)$ refers to the expected profits the expert makes from consumers with a minor problem. Recall that such a consumer arrives with probability $1-\alpha$, and regardless of the consumer's information status, the expert is indifferent between being truthful and making $(p_m - c_m)$ by performing a cheap treatment or lying. This indifference ensures that the expert makes a profit of $(p_m - c_m)$ whenever a consumer with a minor problem arrives.

Consider now the expected profit function $\Pi_2(\cdot)$ in part (ii) of Lemma 4 for the expensive treatment price range $p_s \in (\ell_s^l, \ell_s^h]$. In this range, only the high valuation consumers can afford to accept an expensive treatment recommendation with a positive probability. All low valuation consumers are priced out of the expensive treatment regardless of their information status. For this reason, for the price range $p_s \in (\ell_s^l, \ell_s^h]$, the expert can make a profit from only a fraction θ of the consumers with a serious problem (those with high valuations), which is captured by the first term in $\Pi_2(\cdot)$.

An important observation that emerges from Lemma 4 is that both expected profits $\Pi_1(\cdot)$ and $\Pi_2(\cdot)$ are monotone increasing in the expensive treatment price p_s and the cheap treatment price p_m. This observation implies that the expert maximizes $\Pi_1(\cdot)$ by setting $p_s^* = \ell_s^l$ and $p_m^* = \ell_m$, whereas $\Pi_2(\cdot)$ is maximized by setting $p_s^* = \ell_s^h$ and $p_m^* = \ell_m$. Therefore, while setting the optimal treatment prices, the expert has the following trade-off: If $\Pi_1^*(\ell_m, \ell_s^l) > \Pi_2^*(\ell_m, \ell_s^h)$, then the expert sets $p_s^* = \ell_s^l$ and hence allows the low valuation consumers to accept and expensive treatment recommendation with positive probability. On the other hand, if $\Pi_1^*(\ell_m, \ell_s^l) < \Pi_2^*(\ell_m, \ell_s^h)$, then the expert is better off from pricing out all the low valuation consumers by setting $p_s^* = \ell_s^h$. The pricing choice of the expert then depends on whether the fraction of high valuation consumers is high enough, so that pricing out all low valuation consumers is justified profit-wise. If the fraction of high valuation consumers is above a certain threshold, then the expert's optimization will yield $p_s^* = \ell_s^h$, as otherwise the expert cannot afford to price out all low valuation consumers. One can characterize this threshold value of θ by setting $\Pi_2^*(\ell_m, \ell_s^h) = \Pi_1^*(\ell_m, \ell_s^l)$ which yields the following threshold fraction of high valuation consumers

$$\theta^* \equiv \frac{\ell_s^l - c_s}{\ell_s^h - c_s} \cdot \frac{\ell_s^h - c_m}{\ell_s^l - c_m}.$$

For $\theta > \theta^*$, we have $\Pi_2^*(\ell_m, \ell_s^h) > \Pi_1^*(\ell_m, \ell_s^l)$ and hence the expert sets $p_s^* = \ell_s^h$. For $\theta < \theta^*$, we have $\Pi_2^*(\ell_m, \ell_s^h) < \Pi_1^*(\ell_m, \ell_s^l)$ and hence the expert sets $p_s^* = \ell_s^l$.

This optimal pricing behaviour suggests that, as far as cheating is concerned, there are two types of possible equilibria. For $\theta > \theta^*$, the expert will set $p_s^* = \ell_s^h$ and not cheat any of the consumers; whereas for $\theta < \theta^*$, the expert will set $p_s^* = \ell_s^l$ and cheat the high valuation consumers with a positive probability. Furthermore, the frequency that the high valuation consumers will be targeted for cheating will depend on their information status. The following proposition provides a full characterization of these two equilibrium outcomes.

Proposition 1. ***Case 1 (No Cheating):*** *For $\theta > \theta^*$, in the unique equilibrium the expert sets $p_s^* = \ell_s^h$ and $p_m^* = \ell_m$ and does not cheat any consumers regardless of their valuations and information status; that is,*

$$\beta_m^{z_l} = \beta_m^{z_h} = 0 \text{ for } z \in \{g, b, n\}.$$

In this no-cheating equilibrium, all low valuation consumers always reject an expensive treatment recommendation regardless of their information status, and all high valuation consumers accept an treatment recommendation with the common probability

$$\gamma_s^{z_h} = \frac{\ell_m - c_m}{\ell_s^h - c_m} \equiv \gamma \text{ for } z \in \{g, b, n\}$$

regardless of their information.

Case 2 (Cheating Equilibrium) *For $\theta < \theta^*$, in the unique equilibrium the expert sets $p_s^* = \ell_s^l$ and $p_m^* = \ell_m$, and cheats all the high valuation consumers, but not any of the low valuation consumers. Those high valuation consumers with bad signals are cheated with the highest frequency, whereas the high valuation consumers with good signals are cheated with the lowest frequency. In particular, we have*

$$\begin{aligned} \beta_m^{z_l} &= 0 \text{ for } z \in \{g, b, n\}, \\ \beta_m^{z_h} &= \frac{\alpha_z(\ell_s^h - p_s)}{(1-\alpha_z)(p_s - \ell_m)} > 0 \text{ for } z \in \{g, b, n\} \end{aligned}$$

and hence

$$\beta_m^{b_h} > \beta_m^{n_h} > \beta_m^{g_h},$$

In this cheating equilibrium, all consumers, regardless of their valuations and information status, accept an expensive treatment recommendation with the common probability

$$\gamma_s^{z_h} = \gamma_s^{z_l} \equiv \gamma = \frac{\ell_m - c_m}{\ell_s^l - c_m} \text{ for } z \in \{g, b, n\}.$$

It is interesting to note that, λ, the probability of a consumer receiving an informative signal does not appear to play a role in the equilibrium. This follows because, in the mixed strategy equilibrium, all types of consumers accept an expensive treatment recommendation with a common probability (thus the first term in equations (4) and (5) does not depend on λ). Furthermore, Lemma 3 shows us that, *for each type of consumer*, the expert must be indifferent between falsely making an expensive recommendation or honestly recommending a cheaper treatment (thus the second term in equations (4) and (5) does not depend on λ).

Remark 1 (Fixed Prices and Cheating)**.** *Observe that in the equilibrium described by Proposition 1, cheating occurs if there are sufficiently few high-value consumers (i.e., $\theta < \theta^*$); otherwise there is no cheating. Furthermore, when cheating occurs, it is not the case that the expert specifically targets the uninformed; rather, she cheats the high-value consumers. Within the group of high-value consumers, as was noted above,*

$$\beta_m^{b_h} > \beta_m^{n_h} > \beta_m^{g_h} > 0.$$

That is, the most likely victims of expert cheating are those consumers who received a bad signal, followed by the uninformed and, finally, by those who received a good signal.

This observation suggests that the ability to set prices is a powerful tool that partially mitigates the expert's incentive to cheat. In particular, a small increase in the price, combined with a small reduction in the frequency of cheating, leads to higher profits. Thus, pricing power and the ability to increase the expensive treatment price closer to the consumer's willingness to pay serves as a substitute for cheating. If instead, as in Pitchik and Schotter (1987), the expert faced an exogenous price vector $(p_m, p_s) \in [c_m, \ell_m) \times [c_s, \ell_s^h)$, then it is true that there will be cheating in the corresponding mixed strategy Nash equilibrium. However, our analysis shows that, for all $z \in \{g, b, n\}$, we have $\beta_m^{z_h} > \beta_m^{z_l}$. ***That is, a high-valuation consumer who received signal z is more likely to be cheated than a low-valuation consumer who received the same signal****. Furthermore, for both high- and low-value consumers, $\beta_m^{b_j} \geq \beta_m^{n_j} \geq \beta_m^{g_j}$ (with strict inequality if p_s is strictly less than the consumer's valuation). That is, the probability of being cheated is increasing in the consumer's pessimism about how likely it is that the problem is serious. Our analysis suggests that cheating is likely to emerge in situations where the expert observes that the consumer finds the treatment price* *relatively low with respect to his/her valuation. In those situations, the expert understands that the consumer has some tolerance for being cheated. Provided that there is room for cheating due to this willingness to pay and price, the extent of the expert's cheating then depends on whether the consumer strongly believes to have a serious problem or not. In that respect, an uninformed consumer is always less likely to be cheated than an informed consumer who has received a bad signal and hence is pessimistic about the nature of the problem.*

4 Conclusion

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.

We 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. While this is true, it is incomplete. Specifically, our results show that the ability to set prices acts as a substitute for opportunistic cheating. Moreover, the expert does strictly better by raising prices and an appropriate reduction in the probability of cheating. The end of this process concludes at one of two points. First, if there are relatively few high-value consumers, then the expert sets the price to repair a serious problem at the low-value consumers' reservation value. Regardless of their information status, because the price is at the reservation value, low-value consumers are not cheated by the expert. In contrast, the few high-value consumers are cheated, with the most likely victims of cheating being those who are most pessimistic about the likelihood that their problem is serious. Second, if there are sufficiently many high-value consumers, then the expert chooses a price at their reservation value, meaning that the low-value consumers are priced out of the market. Moreover, in this equilibrium, there is no cheating at all.

We have also analyzed the case (in Appendix B) in which the expert is unable to distinguish between high- and low-value consumers. Here we showed that in the unique equilibrium of the game, there is actually *no cheating*. The same two types of equilibrium outcomes emerge — that is, the price to repair a serious problem is either at the reservation value of the low-value or high-value consumers. In the latter case, that there can be no cheating is obvious. In the former case, since the price is at the reservation value of the low-value consumers, the expert cannot cheat these consumers. However, because the expert is unable to distinguish between high- and low-value consumers, it also means that the expert cannot cheat high-value consumers. This suggests that consumers with observable characteristics, which might tempt an expert to behave opportunistically, would do well to seek out ways to hide this characteristic, or to otherwise try to mimic the type who is less likely to be cheated.

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

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

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

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

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

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

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

[9] 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.

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

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

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

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

A Omitted Proofs

Proof of Lemma 1: Obviously, if $p_s > \ell_s^h$, then all consumers will reject an expensive treatment with probability one. Similarly, if $p_m > \ell_m$, all consumers will reject a minor treatment recommendation with probability one. Next, if $p_m < c_m$, the expert would refuse to treat the consumer (because doing so would generate a loss). Therefore, to generate positive profits we must have $p_m \in [c_m, \ell_m]$. Finally, if $p_s < c_s$, the expert would either provide the minor treatment (which will be rejected by consumers if $p_s \in (\ell_m, c_s)$) or refuse to treat the consumer. Therefore, in any equilibrium where the expert makes positive profits, the set of possible prices must be such that $p_m \in [c_m, \ell_m]$ and $p_s \in [c_s, \ell_s^h]$.

Proof of Lemma 2: First, observe that since $p_m \leq \ell_m < c_s$ and due to the limited liability assumption, regardless of the type of consumer, the expert would never recommend a minor treatment when the problem is serious, as otherwise she would be required to provide the expensive treatment at a loss. Therefore, we must have $\beta_s^t = 1$ for all $t \in T$. Second, because $\beta_s^t = 1$, when the consumer receives a minor treatment recommendation from the expert, he can be sure that the problem is minor. Since $p_m \leq \ell_m$, the consumer strictly prefers to accept. Hence, $\gamma_m^t = 1$ for all $t \in T$.

Proof of Lemma 3: (i) Suppose $p_m \in [c_m, \ell_m]$ and $p_s \in [c_s, \ell_s^l]$. Consider a consumer of type $t \in T$ and suppose to the contrary that $\gamma_s^t = 1$. In this case, since such consumers accept an expensive treatment with probability one, the expert will always cheat these consumers and set $\beta_m^t = 1$.

However, when $\beta_m^t = 1$, the expected benefit of accepting an expensive treatment recommendation for a consumer is at most $\alpha_b \ell_s^h + (1 - \alpha_b)\ell_m$, which is the case if the consumer has a high valuation and has received a bad signal. But note that by assumption A1, we have $\alpha_b \ell_s^h + (1-\alpha_b)\ell_m < c_s < p_s$, which implies that even for that consumer (and hence for all other types of consumers as well) we have $\gamma_s^t = 0$. Therefore, $\beta_m^t = 1$ implies $\gamma_s^t = 0$ for all $t \in T$.

Next note that, if $\gamma_s^t = 0$, we must have $\beta_m^t = 0$: the expert never cheats a consumer who rejects an expensive treatment recommendations with probability one. This follows, because when $\gamma_s^t = 0$, setting $\beta_m^t = 0$ yields $p_m - c_m > 0$ with probability one, whereas cheating yields zero with probability one. Therefore, $\gamma_s^t = 0$ implies $\beta_m^t = 0$ for all $t \in T$.

But if $\beta_m^t = 0$, then we must have $\gamma_s^t = 1$, that is, if the expert is never cheating, then all types of consumers must accept an expensive treatment with probability one. This follows since if $\beta_m^t = 0$, even for a low valuation consumer with a good signal (the consumer which expects to benefit the least from an expensive treatment), the expected benefit of accepting an expensive treatment recommendation is $\ell_s^l \geq p_s$. Hence, $\beta_m^t = 0$ for all t implies $\gamma_s^t = 1$ for all t. But as we already argued, $\gamma_s^t = 1$ implies $\beta_m^t = 1$. Therefore, the above arguments rule out a pure strategy equilibrium and hence we must have $\beta_m^t \in (0,1)$ and $\gamma_s^t \in (0,1)$ for all $t \in T$.

(ii) Note that for $p_s \in (\ell_s^l, \ell_s^h]$, regardless of the expert's cheating behaviour, the price of the xpensive treatment is too high for low valuation consumer, and hence all low valuation consumers regardless of their information status will reject an expensibe treatment recommendation with probability one,

$$\gamma_s^{g_l} = \gamma_s^{b_l} = \gamma_s^{n_l} = 0$$

Consequently, the expert will never cheat a low valutiation consumer, regardless of information status, and we have

$$\beta_m^{g_l} = \beta_m^{b_l} = \beta_m^{n_l} = 0$$

since cheating yields a payoff of zero, whereas truthfullly revealing that the consumer needs a cheap treatment yields $p_m - c_m$ with certainty.

The argument that for high valuation consumers we have $\gamma_s^t \in (0,1)$ and $\beta_m^t \in (0,1)$ for $t \in \{g_h, b_h, n_h\}$ is identical to the proof of part (i) of Lemma 3 and hence is skipped.

Proof of Lemma 4: (i) For the price range $p_m \in [c_m, \ell_m)$ and $p_s \in [c_s, \ell_s^l]$, the expert's ex ante expected profits is given by

$$\Pi(p_m, p_s) = \alpha \left[\lambda \pi_1 + (1-\lambda)\pi_2\right](p_s - c_s) + (1-\alpha)[\lambda \pi_3 + (1-\lambda)\pi_4]$$

where

$$\pi_1 = \phi\theta\gamma_s^{b_h} + \phi(1-\theta)\gamma_s^{b_l} + (1-\phi)\theta\gamma_s^{g_h} + (1-\phi)(1-\theta)\gamma_s^{g_l}$$

$$\pi_2 = \theta\gamma_s^{n_h} + (1-\theta)\gamma_s^{n_l}$$

$$\begin{aligned}\pi_3 &= \phi[(\theta\beta_m^{g_h}\gamma_s^{g_h} + (1-\theta)\beta_m^{g_l}\gamma_s^{g_l})(p_s - c_m) + (\theta(1-\beta_m^{g_h}) + (1-\theta)(1-\beta_m^{g_l}))(p_m - c_m)] \\ &\quad +(1-\phi)[(\theta\beta_m^{b_h}\gamma_s^{b_h} + (1-\theta)\beta_m^{b_l}\gamma_s^{b_l})(p_s - c_m) + (\theta(1-\beta_m^{b_h}) + (1-\theta)(1-\beta_m^{b_l}))(p_m - c_m)]\end{aligned}$$

$$\pi_4 = \theta[\beta_m^{n_h}\gamma_s^{n_h}(p_s - c_m) + (1-\beta_m^{n_h})(p_m - c_m)] + (1-\theta)[\beta_m^{n_l}\gamma_s^{n_l}(p_s - c_m) + (1-\beta_m^{n_l})(p_m - c_m)]$$

Now using the above expression for γ_s^t and β_m^t, one can simplify the above expected profit function and obtain

$$\pi_1 = \pi_2 = \frac{p_m - c_m}{p_s - c_m}$$

and

$$\pi_3 = \pi_4 = p_m - c_m.$$

Accordingly, we the ex ante expected profit function for $(p_m, p_s) \in [c_m, \ell_m) \times [c_s, \ell_s^l]$ takes the following form:

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

(ii) In this case, notice that since $p_s > \ell_s^l$, we must have $\gamma_s^{z_l} = 0$ for $z \in \{g, b, n\}$, so that all low valuation consumers, regardless of their information status, reject an expensive treatment with probability 1. Therefore, it must also be that $\beta_m^{z_l} = 0$. Using similar arguments as above, it can be shown that $\gamma_s^{z_h} \in (0,1)$ and $\beta_m^{z_h} \in (0,1)$ for $z \in \{g, b, n\}$. The acceptance probabilities for the uninformed and informed high valuation consumers can be written as

$$\gamma_s^{z_h} = \frac{p_m - c_m}{p_s - c_m} \equiv \gamma \text{ for } z \in \{g, b, n\},$$

while the cheating probabilities for the expert are given by

$$\beta_m^{z_h} = \frac{\alpha_z(\ell_s^h - p_s)}{(1-\alpha_z)(p_s - \ell_m)}$$

Given these recommendation strategies and the acceptance probabilities, for prices $p_m \in [c_m, \ell_m]$ and $p_s \in (\ell_s^l, \ell_s^h]$ one can write the ex ante expected profit of the expert as

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

B The Expert Cannot Distinguish High-Value From Low-Value Consumers

In this section, we focus on the case in which the expert is unable to distinguish between high- and low-value consumers. All of the primitives are the same as before. The main result here is that the unique equilibrium involves no cheating by the expert.

B.1 Consumers Do Not Receive an Informative Signal

In order to develop intuition for the more complicated case, we first focus on the case in which consumers do not receive an informative signal on whether or not the problem is serious. Since the expert is unable to distinguish between high-value and low-value consumers, she cannot condition her recommendation strategy on the consumer's type; only on whether the problem is serious or minor. To this end, let β_m denote the probability that the expert recommends an expensive repair to a consumer whose problem is actually minor.

We focus our attention on pricing subgames in which $p_m \in [c_m, \ell_m]$ and $p_s \in [c_s, \ell_s^h]$. Because of limited liability, we also immediately impose that $\beta_s = 1$, which means that the expert always recommends an expensive repair when the problem is serious. With these two restrictions in place, we can be sure that the expert will always agree to treat a consumer. Hence, we can ignore ρ (the probability that the expert refuses to treat a consumer).

We first show that if low-value consumers are indifferent between accepting and rejecting an expensive repair, then high-value consumers strictly prefer to accept. The expected loss from rejecting an expensive repair is given by:

$$\mathbb{E}^j(loss|s, reject) = \frac{\alpha\ell_s^j + (1-\alpha)\beta_m\ell_m}{\alpha + (1-\alpha)\beta_m},$$

where $j \in \{l, h\}$. It is easily apparent that, since $\ell_s^h > \ell_s^l$, $\mathbb{E}^h(loss|s, reject) > \mathbb{E}^l(loss|s, reject)$.

We next show that we cannot have a pure-strategy equilibrium. First suppose that $p_s < \ell_s^l$ and that $\beta_m = 0$. In this case, both low-value and high-value consumers will accept with probability 1. However, if all consumers accept with probability 1, then the expert strictly prefers to cheat, which means that $\beta_m = 1$. On the other hand, if $\beta_m = 1$, then consumers prefer to reject since now, by assumption:

$$\mathbb{E}^h(loss|s, reject) = \alpha\ell_s^h + (1-\alpha)\ell_m < c_s \leq p_s.$$

Therefore, the expert strictly prefers to be honest, which means that there can not be a pure-strategy equilibrium.

Taken together, these two points imply that in equilibrium, either the low-value consumers must be made indifferent between accepting and rejecting (leaving the high-value consumers to *accept* with probability 1), or the high-value consumers must be made indifferent between accepting and rejecting (leaving the low-value consumers to *reject* with probability 1).

B.1.1 Case 1: Low-value consumers are indifferent

We first look for an equilibrium in which low-value consumers are indifferent between accepting and rejecting an expensive recommendation. In this case, we must have that $\frac{\alpha \ell_s^j + (1-\alpha)\beta_m \ell_m}{\alpha + (1-\alpha)\beta_m} = p_s$, which implies that:

$$\beta_m = \frac{\alpha}{1-\alpha} \frac{\ell_s^l - p_s}{p_s - \ell_m}.$$

In order for the expert to be willing to mix, it must be that she is indifferent between recommending an expensive or cheap treatment when the problem is minor. That is:

$$[\theta + (1-\theta)\gamma^l](p_s - c_m) = p_m - c_m,$$

which implies that:

$$\gamma^l = \frac{1}{1-\theta} \frac{p_m - c_m}{p_s - c_m} - \frac{\theta}{1-\theta},$$

with the proviso that $\gamma^l \in [0, 1]$.

With this in hand, we are now ready to write down the expected profit function of an expert who sets prices p_m and p_s. In particular, we have:

$$\begin{aligned} \Pi(p_m, p_s) &= \alpha[\theta + (1-\theta)\gamma^l](p_s - c_s) + (1-\alpha)[\beta_m(\theta + (1-\theta)\gamma^l)(p_s - c_m) + (1-\beta_m)(p_m - c_m)] \\ &= \alpha \left[\frac{p_m - c_m}{p_s - c_m} \right] (p_s - c_s) + (1-\alpha)(p_m - c_m). \end{aligned}$$

Therefore, the decision of the expert is to choose $(p_m, p_s) \in [c_m, \ell_m] \times [c_s, \ell_s^l]$ to maximise $\Pi(p_m, p_s)$ subject to $\frac{1}{1-\theta} \frac{p_m - c_m}{p_s - c_m} - \frac{\theta}{1-\theta} \in [0, 1]$. Observe that the constraint can be re-written as:

$$\theta p_s + (1-\theta)c_m \leq p_m \leq p_s.$$

Observe also that the expected profit function is strictly increasing in both p_m and p_s. Therefore, the expert wishes to set p_m and p_s as high as possible, while respecting the constraints.

It is easy to see that in any equilibrium we must have $p_m < p_s$. To see this, suppose that $c_m < p_m = p_s \leq \ell_m$. If the expert increases p_s a little, while leaving p_m unchanged, then both constraints are still satisfied, but profits are higher. Therefore, in equilibrium, there are two possible outcomes:

E.1. $p_s = \ell_s^l$ and $p_m = \ell_m$, or

E.2. $p_s = \frac{1}{\theta}(\ell_m - (1-\theta)c_m) \leq \ell_s^l$ and $p_m = \ell_m$.

B.1.2 Case 2: High-value consumers are indifferent

Now suppose that the expert makes high-value consumers indifferent, meaning that low-value consumers strictly prefer to reject an expensive treatment recommendation. In this case, we have that

$$\beta_m = \frac{\alpha}{1-\alpha}\frac{\ell_s^h - p_s}{p_s - \ell_m}.$$

In order for the expert to be willing to mix, it must be that she is indifferent between recommending an expensive or cheap treatment when the problem is minor. That is:

$$\theta\gamma^h(p_s - c_m) = p_m - c_m,$$

which implies that:

$$\gamma^h = \frac{1}{\theta}\frac{p_m - c_m}{p_s - c_m},$$

with the proviso that $\gamma^h \in [0,1]$.

With this in hand, we are now ready to write down the expected profit function of an expert who sets prices p_m and p_s. In particular, we have:

$$\begin{aligned}\Pi(p_m, p_s) &= \alpha\theta\gamma^h(p_s - c_s) + (1-\alpha)[\beta_m(\theta\gamma^h)(p_s - c_m) + (1-\beta_m)(p_m - c_m)] \\ &= \alpha\left[\frac{p_m - c_m}{p_s - c_m}\right](p_s - c_s) + (1-\alpha)(p_m - c_m).\end{aligned}$$

Therefore, the decision of the expert is to choose $(p_m, p_s) \in [c_m, \ell_m] \times [c_s, \ell_s^h]$ to maximise $\Pi(p_m, p_s)$ subject to $\frac{1}{\theta}\frac{p_m - c_m}{p_s - c_m} \in [0,1]$. Observe that, because of our restriction on prices, $\gamma^h > 0$. Therefore, the only relevant constraint is:

$$p_m \leq \theta p_s + (1-\theta)c_m.$$

It is not too difficult to see that there can be two possible outcomes here:

E.3. $p_s = \ell_s^h$ and $p_m = \ell_m$, or

E.4. $p_s = \ell_s^h$ and $p_m = \theta \ell_s^h + (1-\theta)c_m \le \ell_m$.

Notice that in the event that E.2 determines the equilibrium prices, then the expert cheats with strictly positive probability in equilibrium. However, we now show that E.2 is dominated. Observe that since $p_s = \frac{1}{\theta}(\ell_m - (1-\theta)c_m)$, we are on the cusp between Case 1 and Case 2. Observe also the expected profits of the expert are continuous at $p_s = \frac{1}{\theta}(\ell_m - (1-\theta)c_m)$.[7] We also know that expected profits are strictly increasing in p_s, which means that the expert would do better to increase p_s and focus only on high-value consumers. Therefore, we have the following result:

Corollary 1. *In the unique equilibrium of the game, the expert never cheats with strictly positive probability.*

B.2 Some Consumers Receive an Informative Signal

We now assume that a fraction $\lambda \in (0,1]$ of consumers receive an informative signal about the severity of their problem. Let $\phi \in (0.5, 1)$ denote the precision of the signal. Therefore, 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. Since the signal is informative, it is apparent that $\alpha_b > \alpha > \alpha_g$. We also assume that the expert can observe the signal received by the consumer and, therefore, can condition her recommendation strategy on it. Denote by β_m^j the probability that the expert recommends an expensive treatment when the problem is minor and the consumer received signal $j \in \{b, g, n\}$ (where n denotes that no signal was received).

As was the case when consumers do not receive an informative signal about the severity of the problem, for all j, if β_m^j is such that a low-value, type j consumer is indifferent between accepting and rejecting, then the high-value, type j consumer strictly prefers to accept. Furthermore, since both the expert and the consumer know the signal received by the consumer, the same logic also shows that there cannot be a pure-strategy equilibrium. Table 1 summarises the cheating probability that will make a consumer with value $v \in \{l, h\}$ and signal $j \in \{g, b, n\}$ indifferent between accepting and rejecting an expensive recommendation.

[7] This follows because for $p_s < \frac{1}{\theta}(\ell_m - (1-\theta)c_m)$, the probability that an expensive recommendation is accepted is given by $\theta + (1-\theta)\gamma^l$, while for $p_s > \frac{1}{\theta}(\ell_m - (1-\theta)c_m)$, the probability that an expensive recommendation is accepted is given by $\theta\gamma^h$. However, when $p_s = \frac{1}{\theta}(\ell_m - (1-\theta)c_m)$, $\gamma^h = 1$ and $\gamma^l = 0$, which implies that expected profits are continuous.

Table 1: Critical Cheating Probabilities

Value		Signal Received: g	n	b
	ℓ_s^h	$\frac{\alpha_g}{1-\alpha_g}\frac{\ell_s^h - p_s}{p_s - \ell_m}$	$\frac{\alpha}{1-\alpha}\frac{\ell_s^h - p_s}{p_s - \ell_m}$	$\frac{\alpha_b}{1-\alpha_b}\frac{\ell_s^h - p_s}{p_s - \ell_m}$
	ℓ_s^l	$\frac{\alpha_g}{1-\alpha_g}\frac{\ell_s^l - p_s}{p_s - \ell_m}$	$\frac{\alpha}{1-\alpha}\frac{\ell_s^l - p_s}{p_s - \ell_m}$	$\frac{\alpha_b}{1-\alpha_b}\frac{\ell_s^l - p_s}{p_s - \ell_m}$

Observe that along each row of the table, as we move from left to right, the critical probability is increasing since consumers' pessimism about the severity of the problem is increasing. Also observe that high-value consumers are more tolerant of cheating than are low-value consumers.

In Table 2, we report the acceptance probability of a consumer with value $v \in \{h, l\}$ and signal $j \in \{g, b, n\}$ (assuming that the consumer is indifferent) upon receiving an expensive treatment recommendation. Notice that the acceptance probability is independent of the signal received by the consumer as the signal received by the consumer does not enter into the expert's payoff function.

Table 2: Consumer Acceptance Probabilities (assuming indifferent)

Value		Signal Received: g	n	b
	ℓ_s^h	$\frac{1}{\theta}\frac{p_m - c_m}{p_s - c_m}$	$\frac{1}{\theta}\frac{p_m - c_m}{p_s - c_m}$	$\frac{1}{\theta}\frac{p_m - c_m}{p_s - c_m}$
	ℓ_s^l	$\frac{1}{1-\theta}\frac{p_m - c_m}{p_s - c_m} - \frac{\theta}{1-\theta}$	$\frac{1}{1-\theta}\frac{p_m - c_m}{p_s - c_m} - \frac{\theta}{1-\theta}$	$\frac{1}{1-\theta}\frac{p_m - c_m}{p_s - c_m} - \frac{\theta}{1-\theta}$

Although tedious, it is possible to show that the expected profits function of the expert is given by:

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

for all $(p_m, p_s) \in [c_m, \ell_m] \times [c_s, \ell_s^h]$.[8]

Our goal is to show that there is no cheating in equilibrium. To this end, observe that if $\frac{1}{\theta}\frac{p_m - c_m}{p_s - c_m} < 1$, then $\frac{1}{1-\theta}\frac{p_m - c_m}{p_s - c_m} - \frac{\theta}{1-\theta} < 0$. This means that the expert cannot make indifferent some high-value consumer types and other low-value consumer types. Given this, we can conclude that the same four sets of prices determine potential optimal prices by the expert. That is,

[8] Of course, in different pricing regions, different consumers will be indifferent between accepting and rejecting. Obviously, for example, if $p_s > \ell_s^l$, then only high-value consumers are willing to accept an expensive treatment recommendation.

E.1. $p_s = \ell_s^l$ and $p_m = \ell_m$,

E.2. $p_s = \frac{1}{\theta}(\ell_m - (1-\theta)c_m) \leq \ell_s^l$ and $p_m = \ell_m$,

E.3. $p_s = \ell_s^h$ and $p_m = \ell_m$, or

E.4. $p_s = \ell_s^h$ and $p_m = \theta\ell_s^h + (1-\theta)c_m \leq \ell_m$.

Moreover, exactly the same reasoning as before allows us to conclude that E.2 (the only pricing vector with cheating) can never be optimal, since the expert would earn strictly higher profits by raising prices and focusing only on high-value consumers. Thus, we can conclude the following:

Proposition 2. *In the unique subgame perfect equilibrium of the game, the expert never cheats with strictly positive probability, regardless of whether or not some consumers receive an informative signal on the seriousness of their problem.*