

On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods

UWE DULLECK AND RUDOLF KERSCHBAMER*

Most of us need the services of an expert when our apartment's heating or our washing machine breaks down, or when our car starts to make strange noises. And for most of us, commissioning an expert to solve the problem causes concern. This concern does not disappear even after repair and payment of the bill. On the contrary, one worries about paying for a service that was not provided or receiving some unnecessary treatment. This article studies the economics underlying these worries. Under which conditions do experts have an incentive to exploit the informational problems associated with markets for diagnosis and treatment? What types of fraud exist? What are the methods and institutions for dealing with these informational problems? Under which conditions does the market provide incentives to deter fraudulent behavior? And what happens if all or some of those conditions are violated?

1. Introduction

“If a mechanic tells you that he has to replace a part in your car, don't forget to ask him to put the replaced part into the trunk of your car.” This is an example of day-to-day advice regarding the services of car mechanics. One hopes to avoid two types of fraud by following this advice. On the one hand, the mechanic has to change a part and

cannot only claim (and charge for) the replacement of a part he or she did not actually change. On the other hand, the customer might be able to verify that the mechanic did not provide an unnecessary repair by inspecting the defect of the exchanged part.

This example illustrates a situation where an expert knows more about the type of good or service the consumer needs than the consumer himself. The expert seller is able to identify the quality that fits a customer's need best by performing a diagnosis. He can then provide the right quality and charge for it, or he can exploit the informational asymmetry by defrauding the customer.

Goods and services where an expert knows more about the quality a consumer needs than the consumer himself are called

* Dulleck: Johannes Kepler University Linz and University of Vienna. Kerschbamer: University of Innsbruck and CEPR. Previous versions have benefitted from comments by and discussions with Heski Bar-Isaac, Gerhard Clemenz, Winand Emons, Roger Gordon, Ulrich Kamecke, Gilat Levy, John McMillan, Dennis Mueller, Gerhard Orosel, Carolyn Pitchik, Larry Samuelson, Andreas Wagener, Elmar Wolfstetter, Asher Wolinsky, and three anonymous referees. The first author gratefully acknowledges support by the DFG through SFB 373 at Humboldt University Berlin and SFB 303 at the University of Bonn.

credence goods. Michael R. Darby and Edi Karni (1973) added this type of good to Phillip Nelson's (1970) classification of ordinary, search, and experience goods. They mention provision of repair services as a typical example. Other examples include taxicab rides, where a stranger in a city cannot be sure that the driver takes the shortest route to his destination, and medical treatments, where doctors may provide wrong diagnosis to sell the most profitable treatment or where they can try to charge for a more expensive treatment than provided. There are countless other examples of situations where consumers face similar information problems and where they worry about the two basic problems sketched out by the mechanic example: to pay for goods or services they did not receive or to receive goods or services they did not need in the first place.

Consumers' concerns about being defrauded by experts seem not to be unfounded: Winand Emons (1997) cites a Swiss study reporting that the average person's probability of receiving one of seven major surgical interventions is one third above that of a physician or a member of a physician's family. Asher Wolinsky (1993, 1995) refers to a survey conducted by the Department of Transportation estimating that more than half of auto repairs are unnecessary. He also mentions a study by the Federal Trade Commission that documents the tendency of optometrists to prescribe unnecessary treatment. These examples reveal that the informational asymmetry matters. Monetary incentives play a role too. For the case of the health industry, Jon Gruber, John Kim, and Dina Mayzlin (1999) empirically show that the frequencies of cesarean deliveries compared to normal child births react to the fee differentials of health insurance programs. A similar observation has been made by David Hughes and Brian Yule (1992), who document that the number of cervical cytology treatments is

correlated with the fee for this treatment. That treatment can also be affected by supply of expert services is shown by Victor R. Fuchs (1978): "Other things equal, a 10 percent higher surgeon/population ratio results in about a 3 percent increase in the number of operations . . ." (p. 54).

These observations suggest that we need a more profound understanding of the specific problems associated with markets for diagnosis and treatment. Under which conditions do experts have an incentive to exploit the informational problems associated with markets for diagnosis and treatment? What types of fraud exist? What are the methods and institutions for dealing with these informational problems? Under which conditions does the market provide incentives to deter fraudulent behavior? And what happens if all or some of those conditions are violated?

The existing literature on credence goods does not help much to answer these questions. The studies performed up to now cover very different and quite special situations and in total they yield no clear picture regarding the inefficiencies arising from the informational asymmetry. Rather, the results seem to depend sensitively on the specific conditions of the settings considered and the specific assumptions of the models adopted in the analysis. More disturbingly, apparently similar models often lead to contradicting results (see section 2 for details).

The present article presents a model of credence goods that highlights the information problems in markets for experts services and clarifies the variety of results in this area. The model provides a unifying framework for the analysis of the effects of different market institutions and information structures on efficiency, and for the identification of the forces driving the various inefficiency results in the literature. In addition, the model allows us to generalize

some of the existing results and to show the limits of other ones.

The key feature of credence goods is that consumers do not know which quality of a good or service they need. The provision of a too low quality compared to the needed one is insufficient and the provision of a too high quality does not add extra value. With respect to the auto repair example, if a car breaks down and only an adjustment to the motor is needed to put the car back on the road, then the replacement of the entire engine does not add to the utility the customer derives from the repair. He does not even perceive the difference. That is, the customer can *ex post* only observe (but might be unable to verify) whether the problem still exists. If the problem no longer exists, he can not tell whether he got the right or a too high treatment quality. Furthermore, he may not even be able to observe whether a suggested treatment quality was actually provided or not.

Credence goods can give rise to the following two types of problem: (1) The consumer requires a sophisticated, complex, typically expensive intervention but receives a simple, inexpensive treatment and thus forgoes the benefits of the sophisticated intervention. We refer to this kind of inefficiency as *undertreatment*. A second type of inefficiency arises if the consumer requires an inexpensive treatment but receives a sophisticated one. The additional benefits to the consumer from the sophisticated intervention are less than the additional costs. We refer to this kind of inefficient treatment problem as *overtreatment*. (2) In addition to these inefficient treatment problems surrounding credence goods, credence goods can also be associated with pure transfers from consumers to producers, when a consumer requires and receives an inexpensive treatment but is charged for an expensive one. In the long run, such *overcharging* also leads to an inefficiency if consumers postpone car repairs or medical check-ups and the like because

of the high prices they have paid for these services in the past.¹

A first step in the understanding of credence goods is to identify the conditions that determine which of these two basic problems associated with markets for diagnosis and treatment is the more demanding one. To get a first impression, let us assume we own a garage and our customer cannot observe the repair we perform, but that he will only pay us in case his car is functioning well after the repair (thus undertreatment is ruled out). If this is the case and if we plan to defraud him, we will claim that an expensive and difficult repair is needed even when a minor repair fixes the problem. If the customer authorizes us to perform the repair, we will provide the minor treatment and charge for a more expensive one. Overtreatment is a strictly dominated strategy in this setting because the higher cost of the expensive repair does not affect the payment of the customer. Next assume this customer follows the advice to ask for replaced parts and can thus verify the type of service we perform (but he is still technically not very adept and therefore not able to determine whether the provided service was actually needed or not). In this case, we cannot overcharge the customer but we will consider overtreatment whenever we make a higher profit selling an expensive instead of

¹ There are, of course, a number of other potential problems (and therewith a number of other sources of fraud) in the credence goods market. For example, experts may differ in their ability to perform a diagnosis or a treatment, and this ability may be their own private information. Or, the effort bestowed by an expert in the diagnosis stage may be unobservable. We ignore these problems here, not because we regard them as less important, but rather because they seem to be less specific to the credence goods market. For analyses of situations where effort is needed to diagnose the customer and where an expert's effort investment is unobservable, see Wolfgang Pesendorfer and Wolinsky (2003) and Dulleck and Kerschbamer (2005a). The former contribution focuses on the effect that an additional diagnosis (by a different expert) has on the consumer's evaluation of a given expert's effort and the latter paper studies competition between experts and discounters when customers can free ride on experts' advice.

a cheap repair. In other words, overtreatment is only an issue if overcharging is not feasible due to the specific situation.

In general, the legal framework as well as technicalities and individual education and abilities determine whether, and if yes, which of the mentioned problems appears. Overcharging is only possible if the customer cannot control the service provided by the expert. Asking the mechanic for replaced parts is a way to avoid overcharging and, if one is able to inspect the exchanged part, a device to avoid overtreatment too. If a customer is able to perform a diagnosis himself, there might be no credence goods problem at all. For instance, a patient with a medical education might be able to determine whether the treatment proposed by a physician is really necessary. As the evidence cited above reveals, these differences in abilities affect the treatment prescribed in medical environments.

Existing institutions address the informational problems associated with markets for diagnosis and treatment. The problem of undertreatment is most famously controlled for by the Hippocratic Oath of a physician and its counterpart in the law. Warranties provide similar incentives. The separation of doctors and pharmacies is an institution to avoid overtreatment by disentangling the incentives to prescribe drugs from the profit made selling them. The fixed part of a taxi ride tariff provides incentives for the driver to serve many individual customers and therewith not to take longer routes than necessary. To avoid overcharging as well as overtreatment, in many repair industries the chamber of commerce issues standard work-times for some repairs to allow customers to better check upon a workman's bill—by providing a comparison to usual hours for the ordered repair.

We will see below that the market mechanism itself might discipline experts. In the unifying model that we present, experts have an incentive to commit to prices that induce nonfraudulent behavior and full revelation

of their private information if a small number of critical assumptions are satisfied. These conditions are (i) expert sellers face homogeneous customers; (ii) large economies of scope exist between diagnosis and treatment so that expert and consumer are in effect committed to proceed with an intervention once a diagnosis has been made; and (iii) either the type of treatment (the quality of the good) is verifiable or a liability rule is in effect protecting consumers from obtaining an inappropriate inexpensive treatment (or verifiability and liability both hold). In section 3 we will describe these assumptions in more detail and in section 6 we will discuss and review them in the light of the examples given.

These three basic assumptions are central in our simple, unifying model. We show that by relaxing them one by one, a large part of the existing results on inefficiencies in the credence goods market are obtained. Thus, the analysis of our simple model sheds light on the general structure of credence goods markets. This understanding of the general structure can help to make a next step, from the fractional analysis of certain situations and certain institutions (as done in the existing literature) toward a more complete picture of the issues involved. Given this understanding, we are then in a position to judge other situations and other institutions and to assist in the design of new policies and institutions that remove the inefficiencies surrounding credence goods.

The rest of the paper is organized as follows. The next section briefly reviews the existing literature on credence goods. In section 3 we introduce our model. Section 4 shows that market institutions solve the fraudulent expert problem at no cost when conditions (i), (ii), and (iii) hold. In the subsequent section we characterize the inefficiencies that arise if at least one of these conditions is violated. Section 6 builds a link to real world situations by returning to the key motivating examples mentioned in the text and discussing the plausibility of each of

these assumptions in each case. This section also contains a table where the existing literature is categorized according to our assumptions. Section 7 concludes. Some proofs are relegated to the appendix.

2. On the Existing Literature

Before recapitulating the main contributions of the literature within our model, we want to review briefly the existing literature with respect to assumptions and results.

Models of credence goods markets need to make assumptions on at least three inter-related characteristics of the market considered. They have to specify (1) the technology of the suppliers, i.e., the costs for providing different treatments; (2) the degree of competition together with the organization (“microstructure”) of the market; and (3) the information structure of consumers and courts, i.e., whether consumers are able to observe the quality provided and whether the quality provided and the result of treatment can be proven to the courts. We now turn to each of these points and discuss what the literature assumes.

(1) There are several differences in the considered technologies. Some authors assume that experts can serve arbitrarily many customers at constant (Wolinsky 1993 and 1995, Curtis R. Taylor 1995, Jacob Glazer and Thomas G. McGuire 1996, Dulleck and Kerschbamer 2005a and 2005b, Ingela Alger and François Salanié 2003, Yuk-Fai Fong 2005) or increasing (Darby and Karni 1973) marginal cost, in other contributions experts are capacity constrained (Emons 1997 and 2001, Hugh Richardson 1999). Also, some authors consider models where the right treatment fixes the problem for sure (Carolyn Pitchik and Andrew Schotter 1987 and 1993, Wolinsky 1993 and 1995, Taylor 1995, Dulleck and Kerschbamer 2005a and 2005b, Fong 2005), others focus on frameworks

where success is a stochastic function of service input (Darby and Karni 1973, Glazer and McGuire 1996, Emons 1997 and 2001, Richardson 1999).

(2) The analyzed market conditions range from experts having some degree of market power (Pitchik and Schotter 1987, Richardson 1999, Emons 2001, Dulleck and Kerschbamer 2005b, Fong 2005) to competitive frameworks (Wolinsky 1993 and 1995, Taylor 1995, Glazer and McGuire 1996, Emons 1997, Alger and Salanié 2003, Dulleck and Kerschbamer 2005a). Also, in some contributions experts are able to commit ex ante to take-it-or-leave-it prices (Wolinsky 1993, Taylor 1995, Glazer and McGuire 1996, Emons 1997 and 2001, Dulleck and Kerschbamer 2005a and 2005b, Alger and Salanié 2003, Fong 2005), in others prices are determined ex post in a bilateral bargaining process (Wolinsky 1995, Richardson 1999), still others consider models where prices are exogenously given (Darby and Karni 1973, Pitchik and Schotter 1987 and 1993, Kai Sülzle and Achim Wambach 2005).

(3) Without exception all contributions to the credence goods literature implicitly or explicitly impose our condition (iii). That is, either the type of treatment (the input) is assumed to be observable and verifiable (our verifiability assumption) or the result (the output) is supposed to be verifiable and a liability rule is assumed to be in effect protecting consumers from obtaining an inappropriate inexpensive treatment (our liability assumption).² Which of these two conditions is implicitly or explicitly imposed (none of the contributions imposes both) unambiguously determines the problem on which the

² Alger and Salanié (2003) focus on an intermediate case where some but not all inputs needed for an intervention are observable and verifiable.

respective author(s) focus(es). Emons (1997 and 2001), Richardson (1999), and Dulleck and Kerschbamer (2005a and 2005b) explicitly impose the verifiability assumption and study the problems of over- and undertreatment. Pitchik and Schotter (1987 and 1993), Wolinsky (1993 and 1995), Fong (2005), and Sülzle and Wambach (2005) analyze experts' temptation to overcharge customers, implicitly assuming the liability assumption to hold and verifiability to be violated. Taylor (1995) imposes these assumptions explicitly. And in Darby and Karni (1973), the implicit assumption regarding verifiability varies according to whether capacity exceeds demand or vice versa. For the former case, they analyze experts' incentive to overtreat customers, implicitly assuming verifiability to hold. For the latter case, they discuss the incentive to charge for treatments not provided, implicitly assuming verifiability to be violated.

Given the wide range of problems analyzed in the literature, it is not surprising that the proposed solutions exhibit a broad range of different equilibrium behavior: there are pure strategy equilibria in which experts "mistreat" (i.e., under- or overtreat) some (Dulleck and Kerschbamer 2005b) or all (Darby and Karni 1973, Richardson 1999, Alger and Salanié 2003) consumers, and pure strategy equilibria which give rise to excessive search and diagnosis costs (Wolinsky 1993, Glazer and McGuire 1996).³ There are pure and mixed strategy equilibria where the market outcome involves fraud in the form of overcharging of consumers (Pitchik and Schotter 1987

and 1993, Wolinsky 1995, Fong 2005, Sülzle and Wambach 2005). And there are also equilibria where the only inefficiencies in the credence goods market are experts' inefficient capacity levels (Emons 1997). Thus, the overall picture is rather blurred. At least for the nonexpert reader it is fairly difficult to judge which set of conditions drives the presented results.

The unifying model that we present facilitates the identification of the critical assumptions that lead to the striking diversity of results. In addition, it helps to remove some ambiguities in the literature. For instance, Emons (2001) attributes the main difference between his own work and the papers by Pitchik and Schotter (1987 and 1993) and Wolinsky (1993) to the fact that "they all (implicitly) assume unnecessary repairs to be costless whereas our expert needs resources for unnecessary treatments. This implies that overtreatment is always profitable in their set-up. In contrast, the profitability of overtreatment in our model depends on demand conditions and is determined endogenously" (p. 4). Our analysis below suggests that differences in two of the three conditions specified in the previous (and further discussed in the next) section are more important: First, Pitchik and Schotter (1987 and 1993) and Wolinsky (1993) implicitly assume verifiability to be violated and liability to hold whereas Emons (2001) considers the opposite constellation.⁴ Thus, the former papers focus on the problem of overcharging, while the latter studies experts' incentive to over- or undertreat customers. Secondly, Pitchik and Schotter (1987 and 1993) and Wolinsky (1993) assume that the efficiency loss of diagnosing a consumer more than once is rather low

³ In this context, "some" means that one subgroup of customers is treated efficiently while another subgroup receives one and the same treatment independently of the outcome of the diagnosis. By contrast, "all" means that each treated customer gets the same treatment (again, independently of the type of problem he has). In both cases, some consumers might decide to remain untreated.

⁴ Pitchik and Schotter (1993) claim that their formal framework encompasses both the verifiability and the non-verifiability case (p. 818). In our view, their assumptions on the payoff structure (the profit of an expert authorized by a client to perform an expensive treatment is higher if the client needs a cheap rather than the expensive treatment) hardly allow the verifiability interpretation, however.

TABLE 1
UTILITY FROM A CREDECE GOOD

Customer's utility		Customer	
		\underline{c}	\bar{c}
Customer gets	\underline{c}	v	0
	\bar{c}	v	v

(low economies of scope between diagnosis and treatment), while Emons (2001) assumes this loss to be high (profound economies of scope between diagnosis and treatment).

Below we will reproduce the main results of these studies in our framework and discuss which assumptions are crucial to get them.

3. A Basic Model of Credence Goods

To model credence goods, we assume that each customer (he) has either a (minor) problem requiring a cheap treatment \underline{c} or a (major) problem requiring an expensive treatment \bar{c} .⁵ The customer knows that he has a problem but does not know how severe it is. He only knows that he has an ex ante probability of h that he has the major problem and a probability of $(1 - h)$ that he has the minor one. An expert (she), on the other hand, is able to detect the severity of the problem by performing a diagnosis. She can then provide the appropriate treatment and charge for it or she can exploit the information asymmetry by defrauding the customer.

The cost of the expensive treatment is \bar{c} and the cost of the cheap treatment is \underline{c} , with $\bar{c} > \underline{c}$.⁶ The expensive treatment fixes either

problem while the cheap one is only good for the minor problem.

Table 1 represents the gross utility of a consumer given the type of treatment he needs and the type he gets. If the type of treatment is sufficient, a consumer gets utility v . Otherwise he gets 0. To motivate this payoff structure, consider a car with either a minor problem (car needs oil in the engine) or a major problem (car needs new engine), with the outcomes being “car works” (if appropriately treated or overtreated) and “car does not work” (if undertreated).⁷

An important characteristic of credence goods is that the customer is satisfied in three out of four cases. In general, he is satisfied whenever he gets a treatment quality at least as good as the needed one. Only in one case, where he has the major problem but gets the cheap treatment, will he discover ex post what he needed and what he got.

Consumers may differ in their valuation for a successful intervention or in their probability of needing different treatments. We introduce heterogeneous customers in section 5. In section 4 we follow the main body of the literature on credence goods in assuming that customers are identical. We refer to this as the homogeneity assumption (Assumption H).

⁵ Many of results can also be obtained in an extended model that allows for more than two types of problem and more than two types of treatment (see Dulleck and Kerschbamer 2005b for such a model). However, since our goal is to provide a simple unifying framework, we stick to a binary model.

⁶ For convenience, both the type of treatment and the associated cost is denoted by c .

⁷ Of course, not all credence goods have such a simple payoff structure. For instance, in the medical example the payoff for an appropriately treated major disease might differ from that of an appropriately treated minor disease. Similarly, the payoff for a correctly treated minor disease might differ from that of an overtreated minor disease. Introducing such differences would burden the analysis with additional notation, without changing any of the results, however.

ASSUMPTION H (HOMOGENEITY): *All consumers have the same probability h of having the major problem and the same valuation v .*

An expert needs to perform a diagnosis before providing a treatment. Regarding the magnitude of economies of scope between diagnosis and treatment, there are two different scenarios to consider. If these economies are small, separation of diagnosis and treatment or consultation of several experts may become attractive. With profound economies of scope, on the other hand, expert and customer are in effect tied together once the diagnosis is made. In section 5, we take the diagnosis cost d as a measure of the magnitude of economies of scope between diagnosis and treatment and discuss conditions under which expert and consumer are in effect tied together.⁸ In section 4, we work with the following shortcut assumption which we refer to as the commitment assumption (Assumption C).

ASSUMPTION C (COMMITMENT): *Once a recommendation is made, the customer is committed to undergo a treatment by the expert.*

As mentioned before, the focus of the credence goods literature has been twofold: inefficient treatment, either under- or overtreatment, and overcharging. The inefficiency of treatment can be described by referring to table 1. The case of undertreatment is the upper right corner of the table, the case of overtreatment is the lower left corner. Note that overtreatment is not detected by the customer ($v = v$) and hence cannot be ruled out by institutional arrangements. This is not the case with undertreatment; it is detected by the customer ($0 < v$) and might even be verifiable. If this is the case, and if a legal

rule is in effect that makes an expert liable for the provision of inappropriate low quality, we say that the liability assumption (Assumption L) holds.

ASSUMPTION L (LIABILITY): *An expert cannot provide the cheap treatment \underline{c} if the expensive treatment \bar{c} is needed.*

Referring again to table 1, the second potential problem is that the customer might never receive a signal that discriminates between the upper left and the lower right cell of the table. If this is the case, an expert who discovers that the customer has the minor problem can diagnose the major one so that the customer might authorize and pay for the expensive treatment although only the cheap one is provided. This overcharging is ruled out if the customer is able to observe and verify the delivered quality (he knows and can prove whether he is in the top or the bottom row of the table), and we refer to situations in which consumers have this ability as cases where the verifiability assumption (Assumption V) holds.⁹

ASSUMPTION V (VERIFIABILITY): *An expert cannot charge for the expensive treatment \bar{c} if she has provided the cheap treatment \underline{c} .*

Let us now describe the market environment. There is a finite population of $n \geq 1$ identical risk-neutral experts in the credence goods market. Each expert can serve arbitrarily many customers. The experts simultaneously post take-it-or-leave-it prices. Let \bar{p}^i (\underline{p}^i , respectively) denote the price posted by expert $i \in \{1, \dots, n\}$ for the expensive (cheap) treatment \bar{c} (\underline{c}). An expert's profit is the sum of revenues minus costs over the customers she treated. By assumption, an expert provides the appropriate treatment if she is indifferent between providing the appropriate and

⁸ Since provision of treatment without diagnosis is assumed to be impossible, an increase in diagnosis cost means that consulting more than one expert becomes less attractive.

⁹ An undercharging incentive only exists if the price of the expensive treatment is such that customers reject a \bar{c} recommendation, and if the price of the cheap treatment exceeds the cost of the expensive one. Such price combinations are not observed in equilibrium.

providing the wrong treatment, and this fact is common knowledge among all players.¹⁰

There is a continuum with mass one of risk-neutral consumers in the market. Each consumer incurs a diagnosis cost d per expert he visits independently of whether he is actually treated or not.¹¹ That is, a consumer who resorts to r experts for consultation bears a total diagnosis cost of rd . The net payoff of a consumer who has been treated is his gross valuation as depicted in table 1 minus the price paid for the treatment minus total diagnosis cost. The payoff of a consumer who has not been treated is equal to his reservation payoff, which we normalize to zero, minus total diagnosis cost.¹² By assumption, it is always (i.e., even ex post) efficient that a consumer is treated when he has a problem. That is, $v - \bar{c} - d > 0$. Also, by assumption, if a consumer is indifferent between visiting an expert and not visiting an expert, he decides for a visit, and if a

customer who decides for a visit is indifferent between two or more experts he randomizes (with equal probability) among them.

Figure 1 shows the game tree for the special case where a single expert ($n = 1$) courts a single consumer. The variables v , h , \bar{c} , and \underline{c} are assumed to be common knowledge.¹³ At the outset, the expert posts prices \underline{p} and \bar{p} for \underline{c} and \bar{c} respectively. The consumer observes these prices and then decides whether to visit the expert or not. If he decides against the visit, he remains untreated yielding a payoff of zero for both players. If he visits the expert, a random move of nature determines the severity of his problem.¹⁴ Now the expert diagnoses the consumer. In the course of her diagnosis she learns the customer's problem and recommends either the cheap or the expensive treatment. Next the customer decides whether to accept or reject the recommendation. If he rejects, his payoff is $-d$, while the expert's payoff is zero.¹⁵ Under the commitment assumption (Assumption C), this decision node of the consumer is missing. If the consumer is committed or if he accepts under the noncommitment assumption, then the expert provides some kind of treatment and charges for the recommended one.¹⁶ Under the verifiability assumption (Assumption V), this decision node is degenerate: the expert simply provides the recommended treatment.

¹⁰ Introducing some guilt disutility associated with providing the wrong treatment would yield the same qualitative results as this common knowledge assumption provided the effect is small enough to not outweigh the pecuniary incentives.

¹¹ The diagnosis cost d is assumed to include the time and effort cost incurred by the consumer in visiting a doctor, taking the car to a mechanic, etc. It is also assumed to include a fair diagnosis fee paid to the expert to cover her opportunity cost. A more elaborate model would distinguish between a search and diagnosis cost d_s borne directly by the consumer, a diagnosis cost d_e borne by the expert (where $d = d_s + d_e$) and a diagnosis fee p charged by the expert. Since one of our goals is to reproduce many of the existing results on inefficiencies and fraud in credence goods markets in a simple framework and since the bulk of the literature takes diagnosis fees as exogenously given (to the best of our knowledge the only exceptions are Pesendorfer and Wolinsky 2003, Alger and Salanié 2003, and Dulleck and Kerschbamer 2005a), we do not endogenize the price for diagnosis (but rather assume that $p = d_e$) for easier comparison.

¹² Here, the implicit assumption is that the outside option is not to be treated at all. Again, the car example provides a good illustration. A car may be inoperable for a minor or a major reason, with the lack of treatment giving the same outcome ('car does not work') as undertreatment. The medical example behaves differently. For instance, letting a cancerous growth go untreated is much different than letting a benign growth go untreated. See footnote 7 above, however.

¹³ That consumers know experts' costs of providing different treatments is obviously a very strong assumption. Since it is explicitly or implicitly imposed by all authors who work with the verifiability assumption and since we need it to replicate their results, we introduce it here.

¹⁴ Here note that, from a game-theoretic point of view, there is no difference between a model in which nature determines the severity of the problem at the outset and our model where this move occurs after the consumer has consulted an expert (but before the expert has performed the diagnosis).

¹⁵ We take the convention that each consumer can visit a given expert at most once.

¹⁶ Throughout we assume that an expert's agreement to perform the diagnosis means a commitment to provide a treatment even if treatment-provision is not profitable for the expert. This assumption is not important for our results and we mention in footnotes what changes if the expert is free to send off the customer after having conducted the diagnosis.




Figure 1. Game Tree for the Credence Goods Problems

Under the liability assumption (Assumption L), this decision node is degenerate whenever the customer has the major problem: then the expert must provide the expensive treatment. The game ends with payoffs determined in the obvious way.

The game tree for the model with many consumers and a single expert ($n = 1$) can be thought of as simply having many of these single-consumer games going on simultaneously, with the fraction of consumers with the major problem in the market being h . With more than one expert ($n > 1$), the experts simultaneously post prices p^i and \bar{p}^i ($i \in \{1, \dots, n\}$) for c and \bar{c} respectively.¹⁷

¹⁷ In the efficiency results of section 4, the distinction between the setting where there is only one expert and the one where there are many experts is simply a matter of determining whether the surplus in the market is transferred to consumers or experts. For the results in subsection 5.1, experts need to have market power. In our simple model, market power corresponds to the single expert case. In subsection 5.2, we discuss the case of multiple diagnosis, a frequent phenomenon in markets for expert services. To analyze this phenomenon, we need more than one expert.

Consumers observe these prices and then decide whether to undergo a diagnosis by an expert or not and, if yes, by which expert. Thus, with $n > 1$ a consumer's decision against visiting a given expert (the "out" decision in the game tree) doesn't mean that he remains untreated: he might simply visit a different expert. Similarly, a consumer's payoff if he rejects a given expert's treatment recommendation depends on whether he visits a different expert or not.

With commitment, the game just described is a multistage game with observed actions and complete information; see Drew Fudenberg and Jean Tirole (1991, chapter 3).¹⁸ The natural solution concept for such a game is subgame-perfect equilibrium and we will resort to it in section 4. Subgame-perfection loses much of its bite in the noncommitment case where the less-informed customer has to decide whether to

¹⁸ Tirole (1990, chapter 11) refers to such games as "games of almost perfect information."

stay or to leave without knowing whether the better-informed expert has recommended the right or the wrong treatment. To extend the spirit of subgame-perfection to this game of incomplete information, we require that strategies yield a Bayes–Nash equilibrium not only for each proper subgame, but also for continuation games that are not proper subgames (because they do not stem from a singleton information set). That is, we focus on perfect Bayesian equilibria in the non-commitment case.

4. Efficiency and Honesty with Credence Goods

We start by stating our efficiency result:

PROPOSITION 1: *Under Assumptions H (Homogeneity), C (Commitment), and either L (Liability) or V (Verifiability) or both, market institutions solve the fraudulent expert problem at no cost.*

The proof for this result, as well as the intuition behind it, relies on three observations that are reported as Lemma 1–3 below. Lemma 1 discusses the result under the verifiability assumption. With verifiability alone (i.e., without liability), experts find it optimal to charge a uniform margin over all treatments sold and to serve customers honestly as the following result shows.

LEMMA 1: *Suppose that Assumptions H (Homogeneity), C (Commitment), and V (Verifiability) hold, and that Assumption L (Liability) is violated. Then, in any subgame-perfect equilibrium, all consumers are efficiently served under equal markup prices. The equilibrium prices satisfy $p - c = \bar{p} - \bar{c} = v - d - c - h(\bar{c} - c)$ if a single expert provides the good ($n = 1$), and $p - c = \bar{p} - \bar{c} = 0$ if there is competition in the credence goods market ($n \geq 2$).*

PROOF: First note that with verifiability each expert will always charge for the treatment she provided. The treatment quality provided depends upon the type of price vector (tariff) under which the customer is served. There are three classes of tariffs to

consider, tariffs where the markup for the major treatment exceeds that for the minor one ($\bar{p} - \bar{c} > p - c$), tariffs where the markup of the minor treatment exceeds that for the major one ($\bar{p} - \bar{c} < p - c$), and equal markup tariffs ($\bar{p} - \bar{c} = p - c$). A consumer under a tariff in the first class will always get the major, a consumer under a contract in the second class always the minor interventions. Only under contracts in the last class is the expert indifferent between the two types of treatment and, therefore, behaves honestly.¹⁹ Consumers infer experts' incentives from treatment prices. Thus, if a single expert populates the market ($n = 1$), the maximal profit per customer the monopolist can realize with equal markup prices is $v - d - c - h(\bar{c} - c)$; the maximal obtainable profit with tariffs satisfying $\bar{p} - \bar{c} > p - c$ is $v - d - \bar{c}$; and the maximal profit with tariffs satisfying $\bar{p} - \bar{c} < p - c$ is $(1 - h)v - d - c$. Thus, since $v > \bar{c} - c$, the expert will post the proposed equal markup tariff and serve customers honestly. Next suppose that $n > 1$. Further suppose that at least one expert attracts some customer(s) under a tariff (\hat{p}, \hat{p}) that violates the equal markup rule. Then she can increase her profit by switching to an equal markup tariff (\bar{p}, p) satisfying either $\bar{p} - \bar{c} = p - c = \hat{p} - \bar{c} + (1 - h)(\bar{c} - c)$ (if the old tariff had $\hat{p} - \bar{c} > \hat{p} - c$), or $\bar{p} - \bar{c} = p - c = \hat{p} - c + h(v - \bar{c} + c)$ (in the opposite case). Thus, only equal markup prices can attract customers in equilibrium. The result then follows from the observation that $\bar{p} - \bar{c} = p - c = 0$ is the only equal markup consistent with Bertrand competition. ■

¹⁹ The assumption that it is common knowledge among players that experts provide the appropriate treatment whenever they are indifferent plays an important role in Lemma 1 in generating a unique subgame-perfect equilibrium outcome. Without this assumption, there exist other subgame-perfect equilibria which are supported by the belief that all experts who post equal markup prices—or, that experts who post equal markup prices that are too low (in the monopoly case: too high)—deliberately mistreat their customers. We regard such equilibria as implausible (see footnote 10 above) and have therefore introduced the common knowledge assumption which acts as a restriction on consumers' beliefs.




Figure 2. Consumers' Expected Utility with Verifiability

Figure 2 illustrates the equal markup result graphically. In this figure, the term $\tilde{u}(p, \bar{p})$ denotes the expected net payoff of a consumer who resorts to an expert with posted price-vector (p, \bar{p}) . First notice that verifiability solves the problem of overcharging; that is, the seller cannot claim to have supplied the expensive treatment when she actually has provided the cheap one. There remains the incentive to provide the wrong treatment. Such an incentive exists if the intervention prices specified by the tariff are such that providing one of the treatments is more profitable than providing the other. So, if we fix the markup for the cheaper intervention at $p - c$ and increase the markup for the expensive intervention from 0 (as it is done in figure 2), then the expert's incentives remain unchanged over the interval $(0, p - c)$: she will always recommend and provide the cheap treatment at the price p . Consequently, a consumer's expected utility is constant in this interval at $(1 - h)v - d - p$ and there is an efficiency loss from undertreatment of size $h(v - \bar{c} + c)$. Similarly, if we start at $p - c$ and increase $\bar{p} - \bar{c}$, then the expert will always recommend and provide

the expensive treatment at the price \bar{p} , so that the consumer's utility is $v - \bar{p} - d$ implying an efficiency loss from overtreatment of size $(1 - h)(\bar{c} - c)$. Only at the single point where the markup is the same for both treatments ($\bar{p} - \bar{c} = p - c = \Delta$, say), will the expert perform a serious diagnosis and recommend the appropriate treatment. At this point there is no efficiency loss and consumer's expected utility jumps discontinuously upward to $v - p - h(\bar{c} - c) - d$.

Consumers infer the experts' incentives from the intervention prices. So, experts cannot gain by cheating. Consequently, the best they can do is to post equal markup tariffs and to behave honestly. In the monopoly case ($n = 1$), the expert has all the market power. Hence, posted prices are such that the entire surplus goes to the expert. With two or more experts, on the other hand, Bertrand competition drives profits down to zero and consumers appropriate the entire surplus as experts are not capacity constrained in our model.

The verifiability assumption is likely to be satisfied in important credence goods markets, including dental services, automobile

and equipment repair, and pest control. For more sophisticated repairs, where the customer is usually not physically present during the treatment, verifiability is often secured indirectly through the provision of *ex post* evidence. In the automobile repair market, for instance, it is quite common that broken parts are handed over to the customer to substantiate the claim that replacement, and not only repair, has been performed. Similarly, in the historic car restoration market the type of treatment is usually documented step by step in pictures.

Is there any evidence of equal markups in these markets? Many suppliers in the auto repair and historic car restoration market set standard job-completion times and then charge a uniform hourly rate. This might be interpreted as evidence for uniform margins over all treatments sold.

Equal markup prices are also plausible in cases with expert sellers. Computer stores are an obvious example. Customers can control which quality they receive. Other examples of uniform margins are the pricing schemes of insurance brokers and travel agents. The markup insurance brokers and travel agent charge (the margin plus any bonuses to the agent offered by the provider) are similar for all products. Often also a fixed fee is involved, independent of what insurance is bought or which booking is made.

Lemma 1 facilitates the interpretation of the Emons (1997) result on the provision and pricing strategy of capacity constrained experts. Emons considers a model in which each of a finite number of identical potential experts has a fixed capacity. She becomes an active expert by irreversibly devoting this capacity to the credence goods market. Once she has done this, she can use her capacity to provide two types of treatment at zero marginal cost up to the capacity constraint. One type of treatment consumes more units of capacity than the second. Total capacity over all potential experts exceeds the amount necessary to

serve all customers honestly. Emons proposes a symmetric equilibrium in which each potential expert's entry decision is strictly mixed. Thus, active experts may either have to ration their clientele due to insufficient capacity or they may end up with idle capacity. In the former case, they charge prices such that (1) all the surplus goes to the experts and (2) the price for the more capacity-consuming treatment exceeds the price of the second treatment by such an amount that the profit per unit of capacity consumed is the same for both types of treatment. In the latter case, all experts charge a price of zero for both types of treatment. In both cases experts serve customers honestly. This is exactly what our Lemma 1 would predict: the situation where demand exceeds total capacity corresponds in our model to a situation where $n = 1$, since experts have all the market power in that case and since n is our parameter for market power. With all the market power, experts appropriate the entire surplus and consumers get only their reservation payoff of zero. Furthermore, with insufficient capacity, equal markup prices imply a higher price for the more capacity-consuming treatment since the opportunity cost in terms of units of capacity used is higher. By contrast, if capacity exceeds demand, the opportunity cost is zero for both types of treatment. Thus, the price has to be the same for both treatments to yield equal markups. Furthermore, with idle capacity, experts have no market power (which corresponds in our model to a situation where $n \geq 2$). Without market power, competition drives prices down to opportunity costs and customers appropriate the entire surplus.

To summarize, our analysis shows that many specific assumptions made by Emons (e.g., that there is a continuum of customers, that capacity is needed to provide treatments, that success is a stochastic function of the type of treatment provided, etc.) are not important for his efficiency

result.²⁰ What is important, however, is that consumers are homogeneous (see Proposition 2 below)²¹ and that the type of treatment is verifiable (see Proposition 4, below).²²

Let us turn now to the liability case. Under the liability assumption, experts charge a uniform price for both types of treatment and serve customers honestly as the following result shows.

LEMMA 2: *Suppose that Assumptions H (Homogeneity), C (Commitment), and L (Liability) hold, and that Assumption V (Verifiability) is violated. Then, in any subgame-perfect equilibrium, each expert charges a constant price for both types of treatment and efficiently serves her customers. The price charged in equilibrium is given by $\tilde{p} = v - d$ if a single expert provides the good ($n = 1$), and by $\tilde{p} = \underline{c} + h(\bar{c} - \underline{c})$ if there is competition in the credence goods market ($n \geq 2$).²³*

PROOF: Obvious from the discussion below and therefore omitted. ■

Figure 3 depicts the expected utility of a consumer under the conditions of Lemma 2. In the setting under consideration, the customer cannot control the service provided by the expert. In this case, overtreatment is a strictly dominated strategy because the higher cost of the expensive repair does not affect the payment of the customer. Also, undertreatment is no problem because of the liability assumption. So each expert efficiently

serves her customers. There remains the incentive to overcharge, that is, to always claim that an expensive and difficult repair is needed even when a minor treatment fixes the problem. Such an incentive exists whenever $\bar{p} \neq p$. Consumers know this and expect the expert to charge the higher price and then to provide the cheapest sufficient treatment. Thus, their behavior depends on $\tilde{p} = \max\{p, \bar{p}\}$ only. The rest is trivial. In the monopoly case ($n = 1$), the expert has all the market power. Thus, she charges $\tilde{p} = v - d$. With $n > 1$, the price-posting game is a standard Bertrand game. Thus the price charged in equilibrium is $\tilde{p} = \underline{c} + h(\bar{c} - \underline{c})$ by the usual price-cutting argument.

Lemma 2 offers an explanation for the frequently observed fixed prices for expert services in environments where experts have to provide reliable quality because otherwise they are punished by law or by bad reputation. One example is health maintenance organizations (HMOs) providing medical service to members at an individualized, constant price per customer. Our analysis suggests that the schemes offered by the HMOs are cheaper than a health insurance system. Under insurance, the customer does not care about the price after having paid the insurance premium. With the HMO, the company will make sure that the cheapest sufficient quality will be provided.²⁴

Lemma 2 facilitates the identification of the driving forces behind Taylor's (1995) efficiency results. Taylor studies an elaborate model of a durable credence good which may be in one of three states: health, disease, or failure. The good begins in the state of health, passes over to the state of disease, and from there either to the state of failure or, if efficiently treated, back to the state of health. The amount of time spent in each of the first two states is governed by an exponential distribution. When the good enters

²⁰ The only inefficiency that remains in the Emons model is the suboptimal amount of capacity provided in equilibrium. This inefficiency has nothing to do with the credence goods problem, however. It is rather a coordination failure type of inefficiency similar to that arising in the symmetric mixed strategy equilibrium of the grab-the-dollar game prominent in Industrial Organization.

²¹ In the Emons model, the inefficiencies described in subsection 5.1 arise for *any* n since experts have market power whenever capacity is insufficient to cover demand.

²² Assumption C is not important for the Emons result as Lemma 8 below shows.

²³ If the expert is not obliged to treat the customer after having conducted the diagnosis, the price charged in equilibrium changes to $\tilde{p} = \bar{c} + \alpha(v - d - \bar{c})$, where $\alpha = 1$ for $n = 1$ and $\alpha = 0$ otherwise. The rest of Lemma 2 remains unaffected.

²⁴ As a referee pointed out, this may be a too naive view on HMOs. In section 6, we discuss the plausibility of Assumptions L in this context.




Figure 3. Consumers' Expected Utility with Liability

the state of failure it remains there forever. The main problem for the owner of the good is that he never discovers whether the good is healthy or diseased. An expert, on the other hand, can observe this by performing a diagnostic check. If the check reveals the good to be diseased, she can perform the necessary treatment. Taylor shows that, if experts post treatment prices *ex ante*, then they essentially offer fair insurance against the presence of disease by charging a fixed price equal to the expected treatment cost, as our Lemma 2 would predict. However, in the Taylor model, this insurance solution is not efficient since the good also needs maintenance by the owner for survival. If the owner opts for low maintenance, then the maintenance cost per unit of time is low but the good is later more expensive to treat when it becomes diseased. Obviously, in this set up fixed prices are inefficient since they provide poor incentives for owners to perform maintenance. Taylor proposes two alternative ways to solve this problem: *ex post* contracts where experts commit to treatment prices only after having learned the owner's level of care, and multiperiod contracts for settings with repeated interactions between an owner and an expert.

These solutions would dominate the simple fixed price rule in our setup too if we assumed that the good needs costly maintenance and that the owner's level of care is revealed to the expert in the diagnostic check as is assumed by Taylor.

Our framework shows that many details of the more sophisticated Taylor model are not important for the efficiency results. What is important, however, is that a liability rule protects consumers from obtaining insufficient treatment (without liability experts would never cure the good, except for the prospect of repeat business) and that the type of treatment is not verifiable (otherwise equal markup prices would provide incentives for maintenance).

Let us now consider the case where both the liability and the verifiability assumption hold. In this case, a lower markup for the major than for the minor intervention is sufficient to induce nonfraudulent behavior as the following result shows.

LEMMA 3: *Suppose that Assumptions H (Homogeneity), C (Commitment), L (Liability), and V (Verifiability) hold. Then, in any subgame-perfect equilibrium, each expert posts and charges prices yielding a lower (\leq) markup for the more expensive*

treatment and efficiently serves her customers. Posted (and charged) equilibrium prices satisfy $p + h(\bar{p} - p) = v - d$ and $\bar{p} - \bar{c} \leq p - c$ if a single expert provides the good ($n = 1$), and $p + h(\bar{p} - p) = c + h(\bar{c} - c)$ and $\bar{p} - \bar{c} \leq p - c$ if there is competition in the credence goods market ($n \geq 2$).²⁵

PROOF: The proof is similar to that of Lemma 1, the only difference being that the liability rule prevents the expert from profiting by providing the cheap treatment when the expensive one is needed. ■

Under the conditions of Lemma 3, verifiability prevents an expert from claiming to have supplied the major treatment when she actually has provided the minor one (i.e., overcharging is ruled out). In addition, liability prevents her from providing the minor intervention when her customer needs the major one (i.e., undertreatment is ruled out, too). There remains the incentive to provide a major treatment when the consumer only needs the minor one. This overtreatment incentive disappears if prices are such that the markup for the minor treatment exceeds that of the major treatment.

Lemmas 1–3 discussed possible ideal environments that imply that the price mechanism solves the credence goods problem at no cost. This no-fraud result suggests that asymmetric information alone cannot explain experts' cheating.

Table 2 recapitulates the role our four assumptions play for the efficiency result. If consumers are *homogeneous*, price discrimination is no issue. As we will see below, price discrimination in expert markets proceeds along the dimension of quality of advice offered.

The *commitment* assumption prevents consumers from visiting more than one expert. This assumption is important for efficiency if liability holds while verifiability is violated. Why? Because then a constant

price across treatments is necessary to prevent experts from overcharging (see Lemma 2). A constant price across treatments implies that consumers with minor problems subsidize those with major ones. If consumers are not committed to undergo a treatment once a diagnosis has been made and if the cost of getting a second opinion is low, such cross-subsidization becomes infeasible because it invites cream skimming by a deviating expert who sells the minor intervention only at a lower price. Nondeviating experts who charge a constant price across treatments will realize that only consumers who need a major intervention visit them and they will adjust their prices accordingly. The result is specialization, some experts sell only the minor, others only the major intervention. Specialization is inefficient because it implies a duplication of search and diagnosis costs.²⁶

The *verifiability* assumption prevents overcharging, while the *liability* assumption prevents undertreatment. If none of these assumptions holds, each expert will always provide the minor and charge for the major intervention. Thus undertreatment of consumers who need the major treatment results. Consumers anticipate this and if their valuation for a successful treatment is too low they leave the market.

In the next section, we will discuss these (and other) inefficiencies in more detail.

5. Various Degrees of Inefficiencies and Fraud in the Credence Goods Market

In this section, we characterize the inefficiencies that might arise if at least one of the

²⁵ Obviously, if an expert is not obligated to treat a customer after having conducted the diagnosis, equilibrium prices will also satisfy the condition $\bar{p} \geq \bar{c}$.

²⁶ As one of the referees pointed out, a second diagnosis is not necessary in many settings. Our assumption that provision of treatment without diagnosis is impossible should not be taken literally, however, but rather as a short cut for situations in which there exist economies of scope between diagnosis and treatment. As Darby and Karni (1973, footnote 2) have put it, "... it is easier to repair any damage while the transmission or belly is open to see what is wrong, than to put everything back together and go elsewhere to repeat the process for the actual repair."

TABLE 2
ASSUMPTIONS AND THEIR IMPORTANCE FOR THE EFFICIENCY RESULT

Assumption	Effect
Homogeneity	prevents price discrimination
Commitment	prevents consumers from visiting more than one expert
Liability	prevents undertreatment
Verifiability	prevents overcharging

conditions of Proposition 1 is violated. We begin with the homogeneity assumption (Assumption H).

5.1 *Heterogeneous Consumers: Inefficient Rationing and Inefficient Treatment of some Consumer Groups*

Dropping the homogeneity assumption can give rise to two different kinds of inefficiency: inefficient rationing of consumers and/or inefficient treatment of some consumer groups. To show this, we consider a setting in which a single expert ($n = 1$) sells verifiable treatments (Assumption V holds) to heterogeneous consumers (Assumption H is violated). To keep things simple, we concentrate on the case where consumers only differ in their expected cost of efficient treatment. More precisely, we assume that a consumer of type h has the major problem with probability h and the minor problem with probability $(1 - h)$. Since this assumption implies that higher type consumers have a higher expected cost of efficient treatment, we refer to them as high cost consumers. Consumers' types (i.e., the probabilities h) are drawn independently from the same concave cumulative distribution function $F(\cdot)$, with differentiable strictly positive density $f(\cdot)$ on $[0,1]$. $F(\cdot)$ is common knowledge, but a consumer's type is the consumer's private information. If a consumer gets a treatment that fixes his problem, he obtains a type-independent gross utility of v , and if not one of zero, exactly as in our basic model.

Let us start with a setting in which the expert cannot price discriminate among

consumers. Without price discrimination, the expert chooses equal markup prices. Why? Because consumers infer the expert's incentives from the intervention prices. Thus, the expert cannot gain from cheating. Consequently, the best she can do is to post an equal markup tariff and to provide serious diagnosis and appropriate treatment. With an equal markup tariff, the monopolistic expert is interested in two variables only, in the magnitude of the markup and in the number of visiting customers. If consumers are very similar, the expert finds it profitable to serve all consumers. Otherwise, prices are such that some consumers do not consult her even though serving them would be efficient. This is nothing but the familiar monopoly-pricing inefficiency. We record it in Lemma 4.

LEMMA 4: *Suppose that Assumptions C (Commitment) and V (Verifiability) hold, and that Assumptions H (Homogeneity) and L (Liability) are violated. Further suppose that a single expert ($n = 1$) who cannot price-discriminate among customers serves the market, and that consumers differ in their expected cost of efficient treatment only. Then, in the unique subgame-perfect equilibrium, the expert posts and charges equal markup prices ($\bar{p} - \bar{c} = \underline{p} - \underline{c}$). If the difference in expected cost between the best and the worst type is large relative to the efficiency gain of treating the worst type ($\bar{c} - \underline{c} > (v - d - \bar{c})f(1)$), then prices are such that (i) high cost consumers decide to remain untreated ($\bar{p} > v - d$) and (ii) all other types visit the expert ($\underline{p} < v - d$) and*

get serious diagnosis and appropriate treatment. Otherwise all consumers are efficiently served under equal markup prices ($\bar{p} \leq v - d$).

PROOF: From the proof of Lemma 1 we know that, for given net utilities for the consumers, the monopolist's profit is highest with an equal markup tariff. Thus, the monopolist will choose such a tariff and she will provide the appropriate treatment to all of her customers. With an equal markup tariff, the monopolistic expert is interested in two variables only, in the magnitude of the markup ($\bar{p} - \bar{c} = \underline{p} - \underline{c} = \Delta$, say) and in the number of visiting consumers. The result then follows from the observation that the expert's problem is nothing but the familiar monopoly pricing problem for demand curve $D(\underline{p}) = F[(v - \underline{p} - d)/(\bar{c} - \underline{c})]$ and net revenue per customer $\Delta = \underline{p} - \underline{c}$.²⁷ ■

For our next result, we allow the expert to post more than one tariff. Since consumers differ in their expected cost of efficient treatment, the monopolist might want to target a specific tariff for each consumer or at least different tariffs for different consumer groups. However, in the absence of information about the identity of each consumer (the expert only knows the aggregate distribution of probabilities of needing different treatments), the expert must make sure that consumers indeed choose the tariff designed for them and not the tariff designed for other consumers. This puts self-selection constraints on the set of tariffs offered by the monopolist. As our next result shows, the monopolist uses the quality of advice offered as a self selection device.

LEMMA 5: *Suppose that the general conditions of Lemma 4 hold except that the expert can now price discriminate among consumers (rather than being restricted to a single price vector only). Then, in any*

subgame-perfect equilibrium, the expert posts two tariffs, one with equal markups and one where the markup for the major treatment exceeds that for the minor one.²⁸ Both tariffs attract customers and in total all consumers are served. Low cost consumers are served under the former tariff and always get serious diagnosis and appropriate treatment; high cost consumers are served under the latter and always get the major treatment, sometimes inefficiently.

PROOF: The proof parallels Dulleck and Kerschbamer's (2005b) proof of (their) Proposition 2. We reproduce it here in the appendix to this paper. ■

When price discrimination is possible, the expert finds it optimal to provide high-quality diagnosis and appropriate treatment to low-cost consumers only. High-cost consumers are potentially overtreated; that is, they are induced to demand a high-quality intervention without an honest diagnosis.

Why is such a policy optimal? The reason is simple. Under perfect price discrimination, the monopolist would provide high-quality diagnosis and appropriate treatment to all consumers and gain the entire surplus by charging customer specific prices. In the absence of information about the identity of consumers, perfect discrimination is infeasible. With imperfect discrimination, the expert sells serious diagnosis and appropriate treatment at a relative high price to low-cost consumers only. For the rest of the market, this policy is unattractive since the expected price is larger than the valuation of a successful intervention. Offering efficient diagnosis and treatment at a lower expected price to the residual demand is impossible because this would induce low-cost consumers to switch to the cheaper policy. So, the expert

²⁷ That the condition $\bar{c} - \underline{c} > (v - d - \bar{c})f(1)$ is necessary and sufficient for an interior solution is easily verified by checking that the expert's profit is an increasing function of Δ at $\Delta = v - \bar{c} - d$ if and only if this condition is satisfied.

²⁸ The menu may contain some redundant price vectors too, i.e., some tariffs that attract no consumers.

offers, in addition to the expensive efficient diagnosis and treatment policy, a cheaper but also less efficient one. Her difficulty in designing this second policy is to prevent the low-cost segment from choosing the cheaper policy. The solution is to potentially overtreat the residual demand consisting of high-cost consumers; that is, to induce them to demand a high-quality intervention (a new engine) without a serious diagnosis. For low-cost consumers, this policy is unattractive since their problem is most likely to be a minor one, implying that buying the expensive efficient diagnosis and treatment policy still entails a lower expected cost than buying the high quality intervention (without an honest diagnosis) at a bargain price.

While the equilibrium behavior outlined in Lemma 5 is obviously an abstraction and it is probably impossible to point out an industry that features exactly this kind of price discrimination, the result identifies an element that may be present in the conduct of some industries. The IT industry, for example, features some degree of second degree price discrimination: there are PC manufacturers who distribute their equipment through IT warehouses that offer only selected qualities of equipment at a relatively low price and through specialized dealers that offer the entire assortment as well as advice on choosing the right quality; some consumers (presumably the less profitable ones) buy from warehouse sellers, others consult an expert seller and get serious diagnosis and appropriate equipment.

The equilibrium outlined in Lemma 5 is essentially the overtreatment equilibrium of Dulleck and Kerschbamer (2005b). Dulleck and Kerschbamer go on to show that the result changes dramatically when consumers differ in their valuation for a successful intervention (rather than in their expected cost of efficient treatment). In this case, no overtreatment will be observed but undertreatment appears.

Dulleck and Kerschbamer's *undertreatment* result stands in sharp contrast to the findings in another credence goods paper that has the verifiability assumption and heterogeneous consumers.²⁹ In a model with capacity constrained experts who provide procedures to consumers who differ in their valuation for a successful intervention, Richardson (1999) finds that all treated consumers are *overtreated*; that is, they get a high-quality intervention independently of the outcome of the diagnosis. A closer look at his paper reveals different driving forces behind the Richardson overtreatment and the Dulleck and Kerschbamer (over- and undertreatment) results. The Dulleck and Kerschbamer results are driven by a combination of market power and the ability to price discriminate among heterogeneous consumers. Obviously, the type of treatment must also be verifiable (otherwise overtreating would be strictly dominated by overcharging). By contrast, Richardson's findings result from a lack of commitment power. Consider an expert who can *ex ante* precommit only to the price for a low quality basic intervention. At the diagnosis stage, the expert can tell the consumer that the basic intervention is insufficient to cure his problem and inform him about the additional amount he would have to pay if he accepts an upgrade to a more advanced procedure. The consumer knows that he might have a serious problem and that the basic intervention fails in this case. He is therefore prepared to pay some additional amount for a stronger intervention. If this amount exceeds the difference in treatment costs

²⁹ To the best of our knowledge there are only three further contributions with heterogeneous consumers, one of them is the more verbal paper by Darby and Karni (1973), the second is the 1993 article by Pitchik and Schotter, and the third is Fong (2005). In the former two articles, heterogeneity is only used to purify a mixed strategy equilibrium; in the third contribution, consumers' lack of commitment power plays an important role for the result—we therefore relegate the discussion of that article to the next subsection.

(as is the case under Richardson's assumptions), then the expert has an incentive to always recommend a stronger treatment even if the basic intervention would have been sufficient to cure the problem.³⁰

Before proceeding further, notice that the inefficiencies of lemmas 4 and 5 disappear if experts have no market power. With $n > 1$, price-competing experts provide both types of treatment at marginal cost leaving no leeway for inefficiencies of any kind. Note here that the relevant condition is not $n = 1$, but rather that experts have market power in providing treatments. In a model in which capacity is required to serve customers (cf. e.g., Emons 1997 and 2001 or Richardson 1999), experts have market power (independently of n) whenever tight capacity constraints hamper competition. Similarly, consumer loyalty, travel costs together with location, search costs, collusion and many, many other factors might give rise to market power.

We summarize the results of this subsection in the following proposition.

PROPOSITION 2: *Subgame-perfect equilibria exhibiting inefficient rationing and/or inefficient treatment of some consumer groups can exist if Assumption H (Homogeneity) is violated and experts have market power.*

5.2 No Commitment: Overcharging and Duplication of Search and Diagnosis Costs

In this subsection, we drop the commitment assumption. Under certain conditions, this gives rise to two different types of equilibria—overcharging equilibria and specialization equilibria. In both, the credence goods problem manifests itself in inefficiently high search and diagnosis costs as some consumers end up visiting more than one expert and being diagnosed more

than once. We begin with the overcharging scenario. To get the overcharging result studied in the literature, firms' price setting power needs to be restricted. In part (i) of Lemma 6, this is captured by referring to a (legal) rule requiring experts to choose the price for the expensive treatment from a given range of cost-covering prices. Such a rule is, of course, unrealistic. Part (ii) of the lemma shows that the overcharging equilibrium ceases to exist if prices are fully flexible.

LEMMA 6: (i) *Suppose that Assumptions H (Homogeneity) and L (Liability) hold, and that Assumptions C (Commitment) and V (Verifiability) are violated. Further suppose that there is some competition in the credence goods market ($n \geq 2$), that economies of scope between diagnosis and treatment are relatively low ($d < (\bar{c} - \underline{c})(1 - h)$), and that a (legal) rule is in effect requiring experts to choose the price for the expensive treatment from a given range of cost-covering prices ($\bar{p} \in [\bar{c}, \bar{c} + d]$).³¹ Then there exists a symmetric weak perfect Bayesian equilibrium in which experts overcharge customers (with strictly positive probability). In this equilibrium, experts post prices satisfying $\underline{p} = \underline{c} + \Delta$ and $\bar{p} = \bar{c} > \underline{c} + \Delta$. Experts always recommend the expensive treatment if the customer has the major problem, and they recommend \bar{c} with probability $\omega \in (0, 1)$ if the customer has the minor problem. Consumers make at least one, at most two visits. Consumers at their first visit always accept a \underline{c} recommendation, and they accept a \bar{c} recommendation with probability $\alpha \in (0, 1)$ and reject it with probability $(1 - \alpha)$. Consumers who reject, visit a second (different) expert, and on that visit they accept both recommendations with certainty. A customer who accepts to be treated always gets the appropriate treatment.*

³⁰ Here, remember that equal markup prices are necessary to induce an expert to perform serious diagnosis and to provide the appropriate treatment (see the discussion of Lemma 1).

³¹ If experts are not committed to treat their customers after having conducted a diagnosis, then a lower bound for the price of the expensive treatment ($\bar{p} \geq \bar{c}$) emerges endogenously. So only the upper bound for prices ($\bar{p} \leq \bar{c} + d$) is required in this case to prove part (i) of the lemma.

(ii) *The strategy profile sketched in part (i) of this lemma ceases to form part of a weak perfect Bayesian equilibrium, if experts are completely free in choosing prices.*

PROOF: See the appendix. ■

In the setting under consideration, verifiability is violated. In this case, overtreatment is a strictly dominated strategy because the higher cost of the expensive repair does not affect the payment of the customer. Also, undertreatment is no issue because of the liability assumption. So each expert efficiently serves her customers. There remains the incentive to overcharge, that is, to claim that a major intervention is needed when a minor one fixes the problem. Such an incentive exists in the equilibrium characterized in Lemma 6 since $\bar{p} > \underline{p}$. In the equilibrium, experts do not overcharge their customers all the time (but only with strictly positive probability) because recommending the cheap treatment guarantees a positive profit of $\underline{p} - \underline{c} = \Delta > 0$ for sure, while recommending the expensive treatment when only the cheap one is required is like playing in a lottery yielding a payoff of $\bar{p} - \underline{c} > \Delta$ if the consumer accepts, and zero, otherwise. By construction, the expected payoff under the lottery equals Δ , making experts exactly indifferent between recommending honestly and overcharging.³²

The overcharging equilibrium sketched in part (i) of Lemma 6 is essentially the equilibrium outlined by Pitchik and Schotter

³² For experts to be indifferent between recommending the minor and recommending the major intervention when the customer has the minor problem, the probabilities α and ω and the markup Δ must satisfy $\Delta [1 + \omega(1 - \alpha)] = (\bar{c} - \underline{c})[\alpha + \omega(1 - \alpha)]$. To understand this equation, note that recommending the cheap treatment guarantees a profit of $\underline{p} - \underline{c} = \Delta > 0$ for sure, while recommending the expensive treatment when only the cheap one is required is like playing in a lottery yielding a payoff of $\bar{p} - \underline{c} > \Delta$ with probability $[\alpha + \omega(1 - \alpha)]/[1 + \omega(1 - \alpha)]$, and zero otherwise. This probability takes into account that a fraction $1/[1 + \omega(1 - \alpha)]$ of customers are on their first visit (and hence are accepting the \bar{c} recommendation with probability α), while the remaining fraction $\omega(1 - \alpha)/[1 + \omega(1 - \alpha)]$ are on their second visit (accepting the \bar{c} recommendation for sure).

(1987)—and further studied by Sülzle and Wambach (2005)—for a setting with exogenously given payoffs.³³ Wolinsky (1993) argues that the equilibrium might also exist with flexible prices if sufficiently few experts compete for customers. Lemma 6 shows that the overcharging configuration of part (i) continues to form part of a weak perfect Bayesian equilibrium if experts have some freedom to choose prices, but that it ceases to form part of such an equilibrium if prices are fully flexible.³⁴

The reason for why overcharging equilibria with the essential features as outlined in part (i) of the lemma cease to exist if experts are completely free in choosing prices is that the markup $\underline{p} - \underline{c} = \Delta$, which is necessary to support the described behavior as a weak perfect Bayesian equilibrium, invites price undercutting by a deviating expert who specializes in the minor intervention. How can a deviating expert specialize in the minor intervention? In the present model where each expert's decision variables are her posted prices and her recommendation policy, specialization on the minor intervention occurs via the commitment to a tariff which induces customers who are diagnosed to require the major intervention to reject treatment and to visit another expert. How does such a specialization tariff look? In the equilibrium of Lemma 6, all experts post $(\underline{p}, \bar{p}) = (\underline{c} + \Delta, \bar{c})$. Thus, by posting a tariff that has $\bar{p} > \bar{c} + d$, a deviating expert can credibly signal that she will provide reliable diagnosis and only the minor treatment. The point is that, when the customer has the major problem, the expert cannot sell the minor intervention because of liability. Thus, she will reveal his true condition. And, if the

³³ Sülzle and Wambach's (2005) paper is essentially a comparative statics study on the overcharging equilibrium of Lemma 6. The main focus is on the effect of insurance arrangements on the (mixed strategy) equilibrium behavior of consumers and experts.

³⁴ See Andreu Mas-Colell, Michael Whinston, and Jerry R. Green (1995) for the definition of *weak perfect Bayesian equilibrium*. Roger B. Myerson (1991) calls the same concept a *weak sequential equilibrium*.

customer has the minor problem, the expert will not recommend the major treatment because the customer would reject it and he would visit a nondeviating expert (if he accepts he pays $\bar{p} > \bar{c} + d$, if he rejects and visits a nondeviating expert his cost is $\bar{c} + d$). Under the conditions of part (i) of Lemma 6, such specialization tariffs are infeasible. If they become feasible because prices are flexible (as in part (ii) of the lemma) the overcharging configuration ceases to form part of a weak Bayesian equilibrium.

The “equilibrium with fraud” discussed by Darby and Karni (1973) resembles a purified version of the overcharging equilibrium of Lemma 6. In their analysis “. . . increasing the amount of services prescribed on the basis of the diagnosis, increases the probability of entering the customer’s critical regions for going elsewhere [. . .] Taking this consideration into account, the firm will take fraud up to the point where the expected marginal profit is zero” (p. 73). Adapting Lemma 6 to an environment where consumers are heterogeneous with respect to their search cost would yield an equilibrium with these properties (see Pitchik and Schotter 1993). The analogy is not perfect, however, as the problem discussed by Darby and Karni is that of overtreatment and not that of overcharging customers. Overtreatment can only pose problems if the customer can observe and verify the type of treatment he gets (our Assumption V), since, if he cannot, overcharging is always more profitable. But, with verifiability and flexible prices there are no equilibria exhibiting fraud, as Lemma 8 below shows. In other words, to support the Darby and Karni (1973) configuration with fraud as an equilibrium, payoffs again need to be exogenously fixed as they are in their environment.

A result that resembles a degenerate version of the overcharging equilibrium of Lemma 6 is the “no-cheating result” of Fong (2005). In Fong’s basic model, there is a single expert who offers two types of

treatment to a continuum of homogeneous consumers in an environment in which liability holds while verifiability is violated. Since the expert is completely free in choosing prices, our model would predict that she sets a single price for both treatments at the consumers’ net valuation for a successful intervention (that is, $\bar{p} = p = v - d$; see Lemma 2).³⁵ However, this solution is infeasible in Fong’s setup due to a combination of two assumptions: First, consumers’ valuation for a successful intervention is assumed to be small in comparison to treatment costs. Translated to the present context, Fong’s assumption amounts to $v - d < \bar{c} < v$, while we assume $\bar{c} < v - d$. This modification alone would not change anything provided that it is *ex ante* efficient to visit the expert (i.e., $v - d - c - h(\bar{c} - c) \geq 0$): The experts would still offer to fix both types of problem at the price $v - d$ cross-subsidizing the losses made on selling \bar{c} by the gains made on selling c . Fong’s second departure from our assumptions—that the expert has the option to reject a customer after performing a diagnosis—implies that treatment prices cannot be lower than treatment costs and thereby prevents this (and any other) fixed price solution.³⁶ Also, pure strategy equilibria in which the expert posts two different prices cannot exist because in such an equilibrium it must be common knowledge that only one of the prices will be charged by the expert.³⁷ It remains, as the only possibility, an equilibrium in mixed strategies. In such an equilibrium, the expert posts two prices ($p = v - d / (1 - h)$, $\bar{p} = v$) and nevertheless behaves honestly.

³⁵ Since in Fong’s model there is a single expert, Assumption C is not required to get this result.

³⁶ With a fixed price below \bar{c} , the expert turns down consumers with a major problem; with a fixed price above \bar{c} , no consumer will ever consult the expert (as \bar{c} is above $v - d$ by Fong’s first departing assumption).

³⁷ If the \bar{c} recommendation is accepted for sure, the expert will always recommend \bar{c} ; if only c recommendations are accepted, she will always recommend c . Thus, in an equilibrium in pure strategies a price vector with $p \neq \bar{p}$ is equivalent to a fixed price contract.

The intuition behind this “no-cheating result” parallels that for Lemma 6. The expert does not recommend the major intervention when the minor one fixes the problem because customers accept the major treatment with a low enough probability while they accept the minor one for sure.³⁸ The only difference to the overcharging equilibrium in Lemma 6 is that consumers who get a \bar{c} recommendation are not kept indifferent between accepting and rejecting by the expert’s cheating behavior but rather by the fact that the price charged for the major treatment equals exactly their valuation for a successful intervention ($\bar{p} = v$). Although there is no duplication of search and diagnosis costs in Fong’s model, the equilibrium is still inefficient because a fraction of consumers who need the major treatment remain untreated to keep the expert indifferent.

Based on this basic model, Fong derives two overcharging equilibria, both based on identifiable heterogeneity among consumers. To see how they work, assume that there exist two groups of consumers—one with valuation \underline{v} , the other with valuation $\bar{v} > \underline{v}$, where $\bar{v} - d < \bar{c}$ —and that a consumer’s valuation is observable to the expert. Obviously, from the expert’s point of view it would be ideal if she could price discriminate between the two groups. This is assumed to be impossible; that is, the expert is constrained to post a single price vector only. With a single vector, she can implement the no-cheating result for a single customer group only. If the fraction of \bar{v} consumers is large, she will implement the no-cheating result for those consumers and ignore the rest of the market. By contrast, if there are many \underline{v} consumers, the expert tailors the no-cheating tariff to their

valuation ($\underline{p} = \underline{v} - d/(1 - h)$, $\bar{p} = \underline{v}$) and she will sometimes overcharge \bar{v} consumers to keep them indifferent between accepting and rejecting a \bar{c} recommendation (exactly as in our Lemma 6).³⁹

To summarize, our analysis shows that many specific assumptions made by Fong (e.g., that there is no diagnosis cost or that the treatments for the two problems are not substitutable) are not important for his findings. What is important, however, is that liability holds, while verifiability is violated, that consumers are not committed, that the expert is not committed either, and that consumers’ valuation for successful treatment is rather low.

Overcharging equilibria with the essential features as outlined in Lemma 6 cease to exist if there is more than one expert and if experts are free in choosing prices. With fully flexible prices, a continuum of experts, and low economies of scope between diagnosis and treatment, the only perfect Bayesian equilibria that survive when Assumptions H and L hold while Assumptions C and V are violated are specialization equilibria similar to the one outlined in the next lemma. This has been shown by Wolinsky (1993).⁴⁰

LEMMA 7: *Suppose that Assumptions H (Homogeneity) and L (Liability) hold and that Assumptions C (Commitment) and V (Verifiability) are violated. Further suppose*

³⁹ One difference between the overcharging equilibrium of Lemma 6 and Fong’s overcharging equilibrium deserves mentioning: In the equilibrium of Lemma 6, consumers’ incentive to reject a \bar{c} recommendation increases in ω because a higher ω implies that the consumer has a higher probability to pay less for the treatment on his second visit. By contrast, in Fong’s model consumers have no opportunity to visit a second expert. Thus, an additional assumption is needed to keep consumers indifferent; namely that the loss born by the consumer if his problem is left untreated is greater if \bar{c} rather than \underline{c} is required.

⁴⁰ In the Wolinsky (1993) model, where experts commit to an unobservable recommendation policy at the price-posting stage of the game, an additional restriction on beliefs is required to prove the result. In the present model, that requirement is implied by the notion of perfect Bayesian equilibrium.

³⁸ In Fong’s model, where consumers who reject do not reenter the market, for the expert to be indifferent between recommending the minor and recommending the major treatment when the customer has the minor problem, the probability α must satisfy $\alpha = (\underline{p} - \underline{c})/(\bar{p} - \underline{c}) = [\underline{v} - d/(1 - h) - \underline{c}]/(\underline{v} - \underline{c})$.

that there is enough competition in the credence goods market ($n \geq 4$) and that economies of scope between diagnosis and treatment are relatively low ($d < (\bar{c} - \underline{c})(1 - h)$).⁴¹ Then there exists a perfect Bayesian equilibrium exhibiting specialization. In this equilibrium, some experts (at least two) post prices given by $\underline{p} = \underline{c}$ and $\bar{p} > \bar{c} + d$ (we call such experts “cheap experts”) and some other experts (again at least two, “expensive experts”) post prices given by $\underline{p} \leq \bar{p} = \bar{c}$. Cheap experts always recommend the appropriate treatment, expensive experts always the expensive one. Consumers first visit a cheap expert. If this expert recommends \underline{c} , the customer agrees and gets \underline{c} . If the cheap expert recommends \bar{c} , the customer rejects and visits an expensive expert who treats him efficiently.

PROOF: See the appendix. ■

In the specialization equilibrium of Lemma 7, liability (again) solves the problem of undertreatment and the cost differential $\bar{c} - \underline{c}$ that of overtreatment. The incentive to overcharge customers is eliminated because cheap experts lose their customers if they recommend the expensive treatment (see the discussion above).

Is there any supporting evidence for vertical specialization as outlined in Lemma 7 in the market place? Wolinsky (1993) mentions the automobile repair industry as an example of a market featuring some degree of specialization of this kind: there are backyard garages who specialize in minor repairs and there are certified dealer garages who offer the whole spectrum of treatments; some consumers (presumably those with a lower opportunity cost of time) first visit a backyard garage and, if the problem turns out to be a major one, they turn to a certified dealer.

Note that the condition for the strategies described in Lemma 1 to form part of a

perfect Bayesian equilibrium even if the commitment assumption (Assumption C) is not imposed and even if $n \geq 4$ is exactly that the restriction imposed by Lemma 7 on d is violated; that is, the diagnosis cost d must exceed $(1 - h)(\bar{c} - \underline{c})$. To verify this, notice that a deviation that might jeopardize the equilibrium of Lemma 1 is a specialization tariff that has $\underline{p} < \underline{c} + h(\bar{c} - \underline{c})$ and $\bar{p} \geq \underline{c} + h(\bar{c} - \underline{c}) + d$. The expected cost to a consumer who visits the deviator first and, if recommended the expensive treatment, resorts to a nondeviating expert is $d + (1 - h)\underline{p} + h[\underline{c} + h(\bar{c} - \underline{c}) + d]$. Consulting only a nondeviating expert, on the other hand, costs $d + \underline{c} + h(\bar{c} - \underline{c})$. Thus, to attract customers, the deviator must post a price vector with a \underline{p} such that $d + \underline{c} + h(\bar{c} - \underline{c}) \geq d + (1 - h)\underline{p} + h[\underline{c} + h(\bar{c} - \underline{c}) + d]$, which is equivalent to $\underline{p} \leq \underline{c} + h[(\bar{c} - \underline{c}) - d/(1 - h)]$. But, if $d > (1 - h)(\bar{c} - \underline{c})$, then such a price vector doesn't cover cost, and no deviation is profitable.

The equilibrium outlined in Lemma 7 is essentially the specialization equilibrium of Wolinsky (1993), the only difference being that experts can refuse to provide a treatment after having conducted the diagnosis in the Wolinsky model while they cannot refuse in our present framework.⁴² What drives the Wolinsky result is the combination of two assumptions, the assumption that consumers are neither able to observe the type of treatment they need nor the type they get (our Assumption V is violated), and the assumption that experts are liable for providing the cheap treatment when the expensive one is needed (our Assumption L holds).⁴³ Specialization

⁴¹ If experts are not committed to treat their customers after having conducted the diagnosis, the condition $d < (\bar{c} - \underline{c})(1 - h)$ changes to $d < (\bar{c} - \underline{c})(1 - h)/h$.

⁴² Glazer and McGuire (1996) characterize a similar equilibrium for a setting in which there are (by assumption) two types of experts, safe ones who can successfully serve all consumers, and risky ones whose treatment might fail. The focus of their work is on optimal referral from risky to safe experts after risky experts' diagnosis of a consumer's problem.

⁴³ Wolinsky (1993) doesn't explicitly impose the liability assumption. This assumption is implicit in his specification of consumer payoffs, however.

equilibria cease to exist if Assumption L is violated and they also cease to exist if Assumption V holds. We postpone the discussion of the case where neither liability nor verifiability holds to the next subsection and record the rest of the result as Lemma 8.

LEMMA 8: *Suppose that prices are flexible, that consumers are homogeneous (Assumption H) and that the type of treatment is verifiable (Assumption V). Then the equilibria summarized in Lemma 1 (for the case where Assumption L is violated) and Lemma 3 (for the case where Assumption L holds) remain the only perfect Bayesian equilibria even if Assumption C is violated.*

PROOF: The proof is similar to that of Lemma 1 and therefore omitted. ■

Why is it that specialization equilibria involving costly double advice exist if liability holds while verifiability is violated, while they cease to exist if verifiability holds? The point is that verifiability allows experts to set treatment prices close to marginal cost. With prices close to marginal cost, consumers have no incentive to search for a second opinion.

For obvious reasons, perfect Bayesian equilibria exhibiting specialization and perfect Bayesian equilibria in which experts overcharge customers also cease to exist if a single expert serves the market.⁴⁴

We summarize the results of this subsection in the following proposition.

PROPOSITION 3: *Perfect Bayesian equilibria exhibiting specialization and perfect Bayesian equilibria in which experts overcharge their customers might exist if Assumption C (Commitment) is violated.*

5.3 Neither Liability Nor Verifiability: The Credence Goods Market Breaks Down

In this subsection, we consider an environment resembling a hidden action version

of George A. Akerlof's (1970) lemons model: consumers can neither observe the type of treatment they get (Assumption V is violated), nor can they punish the expert if they realize ex post that the type of treatment they received is not sufficient to solve their problem (Assumption L is violated too). Under these adverse conditions, the expert(s) always provide(s) the cheap (and charge(s) for the expensive) treatment. Consumers anticipate this and consult an expert only if the price of the expensive treatment is such that getting the minor intervention at this price for sure increases their expected utility. If there is no $\bar{p} \geq \underline{c}$ that attracts customers, no trade takes place and the credence goods market ceases to exist.⁴⁵

PROPOSITION 4: *Suppose that Assumption H (Homogeneity) holds and that Assumptions V (Verifiability) and L (Liability) are violated. Then there is no perfect Bayesian equilibrium in which experts serve customers efficiently. If consumers valuation v is sufficiently high ($\underline{v} \geq (\underline{c} + d)/(1 - \underline{h})$), then each expert charges a constant price (given by $\bar{p} = (1 - h)v - d$ if a single expert provides the good, and by $\bar{p} = \underline{c}$ if there is competition in the credence goods market) and always provides the cheap treatment. If the consumers' valuation is too low, then the credence goods market ceases to exist.*

PROOF: Obvious and therefore omitted. ■

The assumption that one treatment is more expensive than the second one is important for the negative result in Proposition 4. Without this assumption (that is, if $\bar{c} = \underline{c}$), expert(s) have no incentive to mistreat customers and therefore behave honestly. This helps to explain the Emons (2001) results. In his 2001 contribution, Emons investigates the same basic model

⁴⁴ Since the uniform price charged by the monopolistic expert under the conditions of Lemma 1 does not exceed consumers' gross utility v , they will not quit even if not committed, for their only alternative is to remain without any treatment.

⁴⁵ Proposition 4 relies on the assumption that consumers are able to verify whether some kind of treatment has been provided or not. If this is not the case, each expert always has an incentive to provide no treatment at all and the credence goods market ceases to exist for any value of v (see also the next footnote).

as in the 1997 paper (see our discussion in section 4). That is, again capacity is required to provide treatments to homogeneous consumers, and once a given capacity level has been installed both types of treatment can be provided at zero marginal cost up to the capacity constraint. Again, one type of treatment consumes more units of capacity than the second. Also, there is no liability rule in effect that protects consumers from getting an inappropriate cheap treatment. The major difference between the two Emons contributions is (1) that the 2001 article considers a credence goods monopolist while the earlier paper is about competing experts, and (2) that the 2001 article distinguishes between the cases of verifiable and unverifiable treatment, and between observable and unobservable capacity while the earlier article deals only with the case of verifiable treatment together with observable capacity. The major result of Emons (2001) is that for three out of the four possible constellations the monopolist always behaves honestly. Only when capacity and treatment are both unobservable no trade takes place.

Efficiency with homogeneous consumers and verifiable treatments is exactly what one would expect given our Lemma 8. What is more intriguing is that Emons obtains efficiency even in one of the non-verifiability cases. The intuition for this result is as follows. With observable capacities, the expert can publicly precommit to a technology (i.e., to a capacity level) that allows her to provide the right treatment to all consumers at zero marginal cost ($\bar{c} = \underline{c} = 0$). Consumers observe capacity and, since they know that there is nothing the monopolist can do with her technology but to provide honest services, they trust the expert and get honest treatment. By contrast, with unobservable capacities, the expert has no precommitment technology at hand and the credence goods market

breaks down.⁴⁶ To summarize, our analysis shows that many specific assumptions made by Emons (e.g., that capacity is needed to provide treatments, that success is a stochastic function of the type of treatment provided, etc.) are not important for his findings. What is important, however, for the positive part of his result is that consumers are homogeneous; with heterogeneous customers the effects described in subsection 5.1 above would emerge and inefficiency would appear.

The equilibrium of Proposition 4 is rather extreme and one is tempted to argue in favor of public intervention, e.g., the introduction of a legal rule that makes the expert liable for providing an inappropriately expensive treatment. In reality, liability rules are far from being perfect mechanisms, however (see the discussion in section 6). Is there a way out? In practice, experts might be kept honest by their need to maintain their reputation (if bad reputation spreads, reputation considerations might mitigate the problem even if consumers are expected to buy only once), or by their desire to retain customers who are expected to frequently need a treatment.⁴⁷

6. Discussion: Theory and Real World Examples

We started our analysis by showing that experts have an incentive to commit to

⁴⁶ In the Emons paper, the credence good market breaks down even if consumers' gross valuation v is high. The reason is the Emons assumption that not only the type of treatment is unobservable, but also whether treatment has been provided at all or not. Under this assumption, expert's capacity investment is zero irrespective of consumers' valuation v . If only the type of treatment were unobservable, as is the case in the present paper, and if v were high enough, then the expert would install a capacity level that allows her to sell \underline{c} to all consumers, exactly as one would expect from our Proposition 4.

⁴⁷ Wolinsky (1993) and In-Ück Park (2005) study credence goods models where the reputation mechanism increases efficiency. That the possibility of reputation building may also go in the opposite direction has been shown by Jeffrey C. Ely and Juuso Välimäki (2003) and by Ely, Fudenberg, and David K. Levine (2005) in their papers on "bad reputation."

prices that induce nonfraudulent behavior and full revelation of their private information if a small set of critical assumptions hold. These conditions are (1) expert sellers face homogeneous customers (Assumption H); (2) large economies of scope exist between diagnosis and treatment so that expert and consumer are in effect committed to continue with a treatment once a diagnosis has been made (Assumption C); and (3) either the type of treatment (the quality of the good) is verifiable (Assumption V) or a liability rule is in effect protecting consumers from obtaining an inappropriately inexpensive treatment (Assumption L). This positive result tempts one into thinking that problems in credence goods markets are not very prevalent. However, some of our organizing assumptions sound more innocent than they really are. In this section, we make the link to real world situations by returning to the key motivating examples mentioned in the introduction and discussing the plausibility of Assumptions C, V, and L in each case (strictly speaking Assumption H is never satisfied in practice). This section also provides a table where the existing literature is categorized according to our assumptions.

Let us start by considering the example of car repairs and other *repair services*. With repair services, strict liability is difficult to impose in practice. For instance, an insufficiently repaired item may work for some time and problems may develop later on. To mitigate the undertreatment problem in such a situation, the liability needs to cover a longer period. But during this longer period the item may stop working for reasons unrelated to the expert's behavior. Also, an extended liability period introduces a moral hazard problem on the consumers' side in situations where the eventual performance of the product depends not only on the expert's but also on the consumer's behavior: If the consumer is fully compensated for a breakdown during the extended liability period, then he has no incentive to take care

of the product (see Taylor 1995 for a model capturing this feature).

The verifiability assumption poses similar problems. For instance, the advice that opens the article urging customers of repair services to ask the expert for replaced parts is sound only if the consumer can verify that the parts actually came from his product and if he has enough knowledge to identify the parts as those he is charged for. The latter problem is pronounced by the fact that most real-world repair settings feature a large number of possible treatments (Is the presented part an oil or an air filter? Did my brakes get their drums ground and new brake shoes or just new brake shoes?). Some customers may have enough expertise to secure verifiability, others do not. Thus, technical expertise, or expert's expectation of its existence on the consumer's side, may affect market outcomes.⁴⁸ The existing literature has ignored consumers' heterogeneity in expertise so far. We think this is an interesting area for future research.

Next consider the commitment assumption. This assumption may be a reasonable approximation to reality for complicated problems where repair is more or less a by-product of diagnosis and where an additional diagnosis adds nothing to the information that is revealed during the repair process anyway (for example, if the problem is a broken part in the gear box of a car). However, it may be violated in many other cases.

With respect to *taxicab rides*, the taximeter, registering time and distance travelled, ensures verifiability and paying on reaching your destination ensures liability. This implies that the only problem to worry about is overtreatment, that is, the driver's incentive to take a circuitous route. Obviously, such an incentive only exists if the driver

⁴⁸ Casual observation suggests that seemingly less technically able customers face a higher risk of becoming victim of overtreatment than seemingly more able customers. See, for example, the article in the *Los Angeles Times* from the 23rd of February 2000, "Your Wheels, More Women Are Beginning To Take A Peek Under The Hood" stating that women are more likely to be defrauded by mechanics.

feels that her passenger is a stranger in the city (and it is therefore often helpful to demonstrate one's knowledge of the city by giving the driver details on which route to take). Thus, again heterogeneity in consumers' knowledge—some consumers need recommendation and treatment; others can self-diagnose the problem and need only the expert's treatment—is an important feature of the market. For market segments with many strangers, our model would predict that drivers commit to prices where the markups for shorter routes exceed those for longer ones (Lemma 3). The usual two-part tariff consisting of a fixed fee and a meter approximates such a tariff. The approximation is not perfect, however. An important aspect of the taxicab-rides market is that the cost of selling \bar{c} instead of \underline{c} varies with the time of the day: during peak times the opportunity cost for the driver to use a circuitous route is high, while it approximates zero during off-peak times. To incorporate such variations in opportunity cost into a time-invariable two part tariff, the variable part of the tariff has to be set equal to zero so that consumers essentially pay a fixed price. A fixed price over all destination is infeasible of course, given the large heterogeneity in rides. However, for more homogeneous market segments fixed price contracts may well be optimal. And in some segments they indeed exist; for instance, in many cities taxicabs charge a fixed price for trips to and from the airport and from one city to the next, independently of the exact pick-up location and/or the concrete destination.

Finally consider the commitment assumption. In the case of taxicab rides, this assumption is implicitly enforced by not asking the driver for the route he plans to take.

Medical treatments offer the most complicated and maybe the most important environment. Strict liability is difficult to impose in this context. With many sicknesses there is no sufficient treatment, with others success is only random. Thus, a failing treatment is no perfect signal of undertreatment. Also,

treatment success is often impossible or very costly to measure for a court, while still being observed by the consumer (how can one prove the presence/absence of pain, for instance). In such a situation, a patient may misreport treatment success to sue the physician for compensation. Similarly, the physician may claim that the treatment was successful, even if she knows it was not. With individual physicians, the Hippocratic Oath (and/or intrinsic motivation) may work reasonable well as a substitute for formal liability. With profit seeking institutions, this solution does not work. This may explain the complaints against HMOs in the United States that they often provide inexpensive treatment where expensive ones are needed. The undertreatment problem is worsened by the fact that HMOs have sought, and sometimes received, relief from liability through government intervention.⁴⁹

Verifiability is also a too demanding assumption with medical treatments. Patients usually are either physically not able to observe treatment—as during an operation—or simply lack the education to verify the treatment delivered by a physician (did I get a root canal and a crown or just a crown from the dentist?).

The commitment assumption (as a shortcut for situations where there are large economies of scope between diagnosis and treatment) seems reasonable for some but not for all cases. For instance, in surgery the assumption of profound economies of scope is reasonable, whereas for the prescription and preparation of drugs it is not.

Insurance coverage seems to be a peculiarity of the market for medical treatments

⁴⁹ We thank a referee for suggesting this discussion. He/she mentions federal laws—ERISA—that preempt state suits; and state and (proposed) federal limits on medical malpractice awards. A (possibly too naive) alternative explanation for the undertreatment complaints against HMOs is that they result from HMOs' incentive to provide an inexpensive treatment whenever it is sufficient. By doing so, they prescribe the inexpensive treatment more often than it is prescribed under an insurance scheme.

and asks for some attention. On the one hand, patients have less incentive to properly take care of themselves because health risks are covered by insurance. On the other hand, they are less worried about overtreatment as they do not carry the full social costs. For treatments with positive side effects (for example, vitamin treatments or massages), patients may even actively ask for overtreatment. But, given that overtreatment is often associated with negative side effects (for example, a chemo therapy in the case of cancer), patients' interests are at least in some cases realigned with parsimonious treatment. On the supply side, insurance companies may have an interest to collude with physicians in expensive cases such that some patients get inefficiently undertreated as long as this cannot be proven. All these indirect incentive effects are not addressed in our model.

Table 3 provides a summary of our findings and categorizes the literature. First suppose that the *homogeneity and the commitment assumption hold*. Then the market provides incentives to deter fraudulent behaviour provided *either the verifiability or the liability assumption is satisfied*. If a liability rule protects consumers from obtaining insufficient treatment while the type of intervention is not verifiable then experts are tempted to overcharge customers, while to overtreat customers is a strictly dominated strategy because the higher cost of the major intervention does not affect the payment of the customer. In this case, a fixed price agreement is an efficient way to solve the credence goods problem (Case 1). Similarly, if verifiability holds so that overcharging is no problem while liability is violated (so that experts might not only be tempted to over- but also to undertreat consumers) then experts' commitment to equal markup prices—i.e., to tariffs where the differences in intervention prices reflect the differences in treatment costs—provides incentives for nonfraudulent behaviour (Case 2). If consumers can observe and ver-

ify the type of treatment they get (verifiability holds) and punish the expert if they realize ex post that the type of intervention they received was insufficient to solve their problem (liability holds too), then the only remaining problem is overtreatment. In this case, experts' commitment to prices where the markup for the minor intervention exceeds (weakly) that for the major one induces truthful behaviour (Case 3).

If the *homogeneity assumption is dropped*, different consumers might be treated differently: High quality advice and appropriate treatment might be provided to the most profitable market segment only. Less profitable consumers might be induced to demand either unnecessary or insufficient procedures (Case 4).

Dropping the commitment assumption (and assuming low economies of scope between diagnosis and treatment) opens up problems of multiple diagnoses. Customers who are not happy with the repair price proposed by the first expert they visit can search for a second opinion. This search for a second opinion makes fraudulent behaviour by the first expert less attractive (she knows that she will lose the business of some consumers she diagnoses as requiring a major intervention), but it also leads to duplication of search and diagnosis costs. As table 3 reveals, dropping the commitment assumption (and assuming low economies of scope) alone is not enough to get equilibria characterized by multiple diagnosis. The type of treatment has also to be nonverifiable. Why is it that specialization equilibria involving costly double advice arise when Assumptions C and V are violated (as in Case 5 in the table) while efficiency prevails when Assumption C alone or Assumptions C and L are violated (Cases 2 and 3)? Under nonverifiability an expert who has recommended the major intervention to a consumer with a minor problem will not have to provide it anyway. Consequently, overtreatment is not a problem in this setting. There remains the incentive to overcharge. In an efficient equilibrium, this incentive is

TABLE 3
 CREDENCE GOODS: ASSUMPTIONS, PREDICTIONS, AND LITERATURE

Case	Assumptions* dropped	Equilibrium characterized by	Equilibrium prices	Case treated in Section (S), Lemma (L), Prop. (P)	Literature
1	V	Full Efficiency and	$\bar{p} = p$	S 4, L 2	Emons (1997)
2	L; L and C	Truthful Revelation	$\bar{p} - \bar{c} = p - c$	S 4, L 1; S 5.2, L 8	Taylor (1995)
3	none; C	of Private information	$\bar{p} - \bar{c} \leq p - c$	S 4, L 3; S 5.2, L 8	
Price Discrimination:					
4	H and L	Over- and/or Undertreatment (exp. having market power)	$\bar{p} - \bar{c} \neq p - c$ for some customers	S 5.1, L 5	Dulleck & Kerschbamer (2005b)
Specialization:					
5	C and V	Duplication of Search and Diagnosis Cost	some exp.: $p = c$, $\bar{p} = \infty$ other exp.: $p \leq \bar{p} = \bar{c}$	S 5.2, L 7	Wolinsky (1993) Glazer&McGuire (1996)
Overcharging:					
6	C and V	Duplication of search and diagnosis cost (when prices inflexible); Some Consumers Untreated (when valuation low and monopolistic expert not committed)	$\bar{p} = \bar{c} > p > c$ $\bar{p} = v > p = v - \frac{d}{(1-h)}$	S 5.2, L 6	Pitchik&Schotter (1989) Wolinsky (1993) Stilzle&Wambach (2003) Fong (2005)
Lemons Problem:					
7	L and V+	Undertreatment/ MarketBreak Down	constant price	S 5.3, P 4	Akerlof (1970) Emons (2001)

* - H: Homogeneity - all customers have the same v and h

- C: Commitment - customers are committed to undergo treatment once a diagnosis has been performed

- V: Verifiability - customers observe and are able to verify to courts the quality of treatment received (rules out overcharging)

- L: Liability - experts cannot provide c when \bar{c} is needed (rules out undertreatment)

+ - and the following combinations of Assumptions dropped: C, L&V; H, L&V; C, H, L&V

removed by experts' commitment to charge a constant price across treatments (Case 1 in the table). The reason why this solution is not attainable with low economies of scope is that the constant price across treatments implies that consumers with minor losses subsidize those with major losses. If consumers are not committed and if the cost of getting a second opinion is low (low economies of scope) this cross-subsidization invites ex ante cream skimming by a deviating expert who specializes in the minor intervention (by committing to prices that induce the customer to reject the expensive recommendation with certainty, see Case 5 in the table). By contrast, if the

type of intervention is verifiable, treatment prices can be set close enough to marginal cost to deter cream skimming (Case 2). Dropping assumptions C and V might also lead to overcharging equilibria (Case 6 in the table). In such an equilibrium, liability removes the under-, the cost differential between the interventions the overtreatment incentive. So, only experts' temptation to overcharge consumers remains. In equilibrium experts do not overcharge their customers all the time (but only with strictly positive probability) because recommending the cheap treatment guarantees a small positive profit ($p - c > 0$) for sure (because customers

accept such a recommendation with certainty), while recommending the expensive treatment when only the cheap one is required is like playing in a lottery yielding a high payoff ($\bar{p} - \underline{c} > p - \underline{c}$) if the consumer accepts, and zero, otherwise. By construction, the expected payoff under the lottery equals $p - \underline{c}$, making the expert exactly indifferent between recommending honestly and overcharging. As we have shown in subsection 5.2 this overcharging equilibrium ceases to exist if experts are completely free in choosing prices. Then a deviating expert can again specialize on the minor intervention and thereby attract all consumers on their first visit.

Whenever *neither verifiability nor liability holds*, an expert has an incentive to provide low quality but to charge for high quality. In this situation we have Akerlof's (1970) market for lemons problem and the market may break down all together (Case 7 in table 3).

7. Conclusions

Information asymmetries in expert–customer relationships are an everyday life phenomenon. Education and experience give experts the ability to diagnose the exact needs of customers who themselves are only able to detect a need but not the most efficient way to satisfy it. The information problems in markets for diagnosis and treatment suggest that experts may be tempted to defraud customers. The present article has shown that the price mechanism alone is sufficient to solve the fraudulent expert problem if a small set of critical assumptions are satisfied and that most existing results on inefficiencies and fraud in credence goods markets can be reproduced by dropping one of those assumptions.

Our systematic approach is new. Previous work has fostered the impression that the equilibrium behavior of experts and consumers in the credence goods market delicately depends on the details of the model. By contrast, the present paper has shown that the results for the majority of

the specific models can be reproduced in a simple unifying framework.

Our analysis suggests that market institutions solve the fraudulent expert problem at no cost if (1) expert sellers face homogeneous customers, (2) large economies of scope exist between diagnosis and treatment so that expert and consumer are in effect committed to continue with a treatment once a diagnosis has been made, and (3) either the type of treatment is verifiable or a liability rule is in effect protecting consumers from obtaining an inappropriate inexpensive treatment.

We have shown that inefficient rationing and inefficient treatment of some consumer groups may arise if condition (i) fails to hold, that equilibria involving overcharging of customers or specialization of experts—both leading to a duplication of search and diagnosis costs—may result if condition (ii) is violated, and that the credence goods market may break down altogether if condition (iii) doesn't hold.

Our model might be considered restrictive in several respects. It rests on the assumption that there are only two possible types of problem, that only two types of treatment exist, that treatment costs are observable, that posted prices are take-it-or-leave-it prices, that experts can diagnose a problem perfectly, and so on. This is certainly a justified criticism. Nevertheless, our simple model is sufficient to derive most results of that class of models that have been the focus of research in the credence goods literature.⁵⁰ Thus, our simple model provides a useful benchmark for the development of more general frameworks which allow for an assessment of the robustness of the results.

⁵⁰ Our model is insufficient, however, to reproduce the Wolinsky (1995) result, which relies on the assumption that posted prices are bargaining prices. It is also insufficient to provide Taylor's (1995) theoretical micro-foundations of several features observed in markets for diagnosis and treatment.

With respect to robustness, a comparison of common experience with the results of this paper suggests that something is missing from existing models of credence goods.⁵¹ As discussed earlier, one common problem in such markets is that of overtreatment, that is, that a more expensive treatment is provided than is needed. Suppose for example, that our car needs either a new fuse or a new engine. A common fear is that the garage will sell us an engine, even when the fuse would be sufficient, because it is more profitable to replace the engine. Some of us would suspect this will be the case even if we can tell whether the car is fixed when the garage is done (and we refuse payment if the car is not fixed) and even though we can see whether the engine has been replaced, i.e., even if liability and verifiability are satisfied. Proposition 1 of this article tells us that such fraud cannot occur if liability and verifiability, as well as homogeneity and commitment, are satisfied; whereas common experience suggests the opposite. What accounts for this dissonance? The answer appears to be that, in such situations, prices should adjust so that the firm can make just as much profit on selling us the fuse (when we only need the fuse) as on selling us the engine. What does this mean? For experts facing tight capacity constraints (for example, work time) the requirement is that the markup on the engine must exceed the markup on the fuse by such an amount that the profit per unit of capacity consumed is the same for both types of intervention. Such markups are likely to be feasible. By contrast, if experts hold idle capacities then the absolute markup has to be the same for both types of treatment. But, are consumers really prepared to pay the same absolute markup for a fuse as for an engine? Common experience tells us that they are not, to the contrary, most of us would punish an expert who charges a high absolute margin on a fuse with malicious

gossip. Is there a way out? For some of us a feasible solution might be to follow the day-to-day advice mentioned in the beginning of the article: to ask the mechanic to put the replaced part in the back of the car and to inspect the defect of this part. It remains to hope that future research comes up with solutions for this problem that are both more efficient (in the sense that consumers do not have to investigate exchanged parts) and feasible also for the less technically adept. For the meantime, a valuable advice might be to consult an expert featuring a lengthy queue of customers waiting for service.

Appendix

Proofs of lemmas 5, 6, and 7 follow.

PROOF OF LEMMA 5: The proof proceeds in four steps. In *Step 1* we show that any arbitrary menu of tariffs partitions the type-set into (at most) three subintervals delimited by cut-off values h_{10} , h_{12} , and h_{02} , with $0 \leq h_{10} \leq h_{12} \leq h_{02} \leq 1$ and either $h_{12} = h_{02}$ or $h_{02} = 1$ (or both) such that (i) the optimal strategy of types in $[0, h_{10})$ is to choose a tariff where the markup for the minor treatment exceeds that for the major one, (ii) the optimal strategy of types in (h_{10}, h_{12}) is to decide for an equal markup tariff, and (iii) the optimal strategy of types in $(h_{12}, 1)$ is either to choose a tariff where the markup for the major treatment exceeds that for the minor one ($h_{02} = 1$), or to remain untreated ($h_{12} = h_{02}$).⁵² Our strategy is then to show in *Step 2* that an optimal price-discriminating menu cannot have $h_{10} = h_{12}$ (that is, there

⁵²The borderline types h_{10} and h_{12} are indifferent between the strategies of the types in the adjacent intervals (whenever such intervals exist). Here note that we allow for $h_{12} = 1$ (all consumers are served and no consumer chooses a tariff where the markup for the major treatment exceeds that for the minor one), for $h_{10} = h_{12}$ (no consumer is attracted by an equal markup tariff), and for $h_{10} = 0$ (no consumer is attracted by a tariff where the markup for the minor treatment exceeds that for the major one). Price discrimination requires, however, that at least two of the three relations hold as strict inequalities.

⁵¹We thank an anonymous referee for suggesting this discussion.

must be an equal markup tariff which attracts a strictly positive measure of types), to show next (in *Step 3*) that $h_{10} = 0$ whenever $h_{10} < h_{12}$ (that is, the expert has never an incentive to post a menu where both an equal markup tariff and a tariff with a higher markup for the cheap treatment attract types), and to show in the end (*Step 4*) that the expert has indeed always a strict incentive to cover a strictly positive interval by a tariff where the markup for the major treatment exceeds that for the minor one ($h_{12} < h_{02} = 1$).

Step 1: First note that any arbitrary menu of tariffs can be represented by (at most) three variables, by the lowest $\Delta_{02} \equiv \bar{p} - \bar{c}$ from the class of tariffs where the markup for the major treatment exceeds that for the minor one (we denote the lowest Δ_{02} in this class by Δ'_{02}), by the lowest $\Delta_{10} \equiv \underline{p} - \underline{c}$ from the class of tariffs where the markup for the minor treatment exceeds that for the major one (we denote the lowest Δ_{10} in this class by Δ'_{10}), and by the lowest equal markup $\Delta_{12} \equiv \bar{p} - \bar{c} = \underline{p} - \underline{c}$ from the class of all equal markup tariffs in the menu (denoted by Δ'_{12}).⁵³ To see this, note that each possible tariff is member of exactly one of these three classes and that a customer who decides for a tariff in a given class will always decide for the cheapest one.⁵⁴ An immediate implication is that each menu of tariffs partitions the type-set into the above mentioned three subintervals. This follows from the fact that the expected utility under the Δ'_{02} tariff is type-independent (implying

that either $h_{12} = h_{02}$, or $h_{02} = 1$, or both), while the expected utility under both, the Δ'_{12} tariff and the tariff where the markup for the minor treatment exceeds that for the major one, is strictly decreasing in h , and from $v > \bar{c} - \underline{c}$ (implying that the Δ'_{10} -function is steeper than the Δ'_{12} -function).

Step 2: To see that $h_{10} < h_{12}$, suppose to the contrary that $h_{10} = h_{12}$. Then $h_{10} > 0$, since $h_{10} = h_{12} = 0$ is incompatible with price-discrimination (and since—by Lemma 4—a non-price-discriminating expert will always decide for an equal markup tariff). But such a menu is strictly dominated since the Δ'_{10} tariff can always be replaced by a tariff with equal markups of $\Delta_{12} = \Delta'_{10} + h_{10}(v - \bar{c} + \underline{c})$; the latter attracts exactly the same types as the replaced one and yields a strictly higher profit.

Step 3: To see that $h_{10} = 0$ whenever $h_{10} < h_{12}$, suppose to the contrary that $0 < h_{10} < h_{12}$. Then the expert's profit is strictly increased by removing all tariffs where the markup for the minor treatment exceeds that for the major one from the menu. This follows from the observation that (by the monotonicity of the expected utility—in h —under equal markup tariffs) all types in $[0, h_{10})$ switch to Δ'_{12} when all tariffs where the markup for the minor treatment exceeds that for the major one are removed from the menu, and from the fact that the expected profit per customer is strictly higher under Δ'_{12} than under Δ'_{10} whenever $0 < h_{10} < h_{12}$, since $\Delta'_{12} \leq \Delta'_{10}$ is incompatible with the shape of expected utilities ($\Delta'_{12} \leq \Delta'_{10}$ implies that $v - d - \underline{c} - h(\bar{c} - \underline{c}) - \Delta'_{12} > (1 - h)v - d - \underline{c} - \Delta'_{10}$ for all $h > 0$ contradicting $h_{10} > 0$). Thus, $h_{10} = 0 < h_{12} \leq h_{02} \leq 1$. So, if price discrimination is observed in equilibrium, it is performed via a menu that contains two tariffs, one with equal markups and a tariff where the markup for the major treatment exceeds that for the minor one.⁵⁵

⁵³ An immediate implication of this observation is that successful price discrimination requires that some types are mistreated with strictly positive probability. Why? Since at least two tariffs must attract a positive measure of consumers and since only one of them can be an equal markup tariff.

⁵⁴ Under a tariff where the markup for the major treatment exceeds that for the minor one, neither the consumer nor the expert cares about the associated \underline{p} . All tariffs in this group that have the same Δ_{02} can therefore be thought off as being a single tariff without any loss in generality. The argument for tariffs where the markup for the minor treatment exceeds that for the major one is symmetric.

⁵⁵ The menu might contain some redundant vectors too, which can safely be ignored, however.

Step 4: We now show that the expert has always a strict incentive to post such a menu. Consider the equal markup tariff posted by the expert under the conditions of Lemma 4. The markup in this vector is at least $\Delta_{12} = v - d - \bar{c}$, in an interior solution even higher. First suppose that the monopolist's maximization problem under the conditions of Lemma 4 yields an interior solution (i.e., $\Delta_{12} > v - d - \bar{c}$). Then the expert can increase her profit by posting a menu consisting of two tariffs, the one chosen under the conditions of Lemma 4, and a second tariff, where the markup for the major treatment exceeds that for the minor one, with $\bar{p} = v - d$. The latter vector guarantees each type an expected utility equal to the reservation utility of 0. Thus, all types that remain untreated under the conditions of Lemma 4 will opt for it since they are indifferent. Also, all types served under the conditions of Lemma 4 still choose the equal markup vector since $v - d - \underline{c} - h(\bar{c} - \underline{c})$ is strictly decreasing in h . Hence, since $v - d > \bar{c}$, and since all types in $[0,1]$ have strictly positive probability, the expert's expected profit is increased.⁵⁶ Now suppose that the monopolist's maximization problem under the conditions of Lemma 4 yields the corner solution $\Delta_{12} = v - d - \bar{c}$. Then again the monopolist can increase her profit by posting a menu consisting of two tariffs, a tariff where the markup for the major treatment exceeds that for the minor one, with $\bar{p} = v - d$, and an equal markup tariff that maximizes $\pi(\Delta_{12}) = \Delta_{12}F[(v - d - \underline{c} - \Delta_{12})/(\bar{c} - \underline{c})] + (v - d - \bar{c})(1 - F[(v - d - \underline{c} - \Delta_{12})/(\bar{c} - \underline{c})])$. Since $\pi(\Delta_{12})$ is strictly increasing in Δ_{12} at $\Delta_{12} = v - d - \bar{c}$ an interior solution is guaranteed. ■

⁵⁶ Here note that the expert can do even better by increasing Δ_{12} . This follows from the observation that the expert's trade-off under the conditions of Lemma 4 is between increasing the markup charged from the types in the segment of served customers and losing some types to the unprofitable segment of not served consumers, while the trade-off here is between increasing the markup charged from the types in the segment of customers served under the more profitable equal markup vector and losing some types to the segment of customers served under the less profitable Δ'_{12} tariff.

PROOF OF LEMMA 6 Part (i): First notice that, under the conditions of Lemma 6, liability prevents undertreatment, and the cost differential ($\bar{c} - \underline{c}$) prevents overtreatment. So, each expert's provision policy is trivial and we will therefore focus in the rest of the proof on experts' price-posting and recommendation policy and on consumers' visiting and acceptance decisions. For the recommendation and acceptance behavior described in part (i) of the lemma to potentially form part of a (weak) Perfect Bayesian Equilibrium (henceforth PBE), the equilibrium probabilities ω^* and α^* , and the markup Δ^* must satisfy

$$(1) \quad d = (\bar{c} - \underline{c} - \Delta^*) \frac{(1 - \omega^*)\omega^*(1 - h)}{h + (1 - h)\omega^*}$$

and

$$(2) \quad \Delta^* = (\bar{c} - \underline{c}) \frac{\alpha^* + \omega^*(1 - \alpha^*)}{1 + \omega^*(1 - \alpha^*)}.$$

The *first* of these two equations guarantees that consumers who get a \bar{c} recommendation are indifferent between accepting and rejecting,⁵⁷ the *second equation* guarantees that experts are indifferent between recommending \underline{c} and recommending \bar{c} when the consumer has the minor problem.⁵⁸

⁵⁷ If a consumer rejects, he incurs an additional diagnosis cost of d for sure. His benefit is to pay less for the treatment on his second visit with probability $[(1 - \omega^*)\omega^*(1 - h)]/[h + (1 - h)\omega^*]$ because he has the minor problem with probability $[\omega^*(1 - h)]/[h + (1 - h)\omega^*]$ given that the expert has recommended \bar{c} .

⁵⁸ Recommending the cheap treatment guarantees a profit of $\Delta^* > 0$ for sure, while recommending the expensive treatment when only the cheap one is required is like playing in a lottery yielding a payoff of $\bar{p} - c > \Delta^*$ with probability $[\alpha^* + \omega^*(1 - \alpha^*)]/[1 + \omega^*(1 - \alpha^*)]$, and zero otherwise. This probability takes into account that a fraction $1/[1 + \omega^*(1 - \alpha^*)]$ of customers are on their first visit (and hence, are accepting the \bar{c} recommendation with probability α^*), while the remaining fraction $\omega^*(1 - \alpha^*)/[1 + \omega^*(1 - \alpha^*)]$ are on their second visit (accepting the \bar{c} recommendation for sure). Here notice that a consumer who is indifferent between accepting and rejecting a \bar{c} recommendation on his first visit will accept the \bar{c} recommendation on his second visit because the probability of needing the minor treatment is lower on the second than on the first visit.

Obviously, consumers who get a \underline{c} recommendation will accept with certainty. Also, since $\underline{p} < \bar{c}$, and since a \underline{c} recommendation is always accepted, experts will always recommend \bar{c} when the consumer has the major problem. So, if prices are exogenously fixed at $(\underline{p}, \bar{p}) = (\underline{c} + \Delta^*, \bar{c})$ and if Δ^* is such that the probabilities α^* and ω^* satisfying equations (1) and (2) above are in $(0,1)$, then the behavior described in the lemma forms part of a PBE.⁵⁹

To show that this behavior continues to form part of a weak PBE even if experts are free to choose \underline{p} while \bar{p} can vary within the range specified in part (i) of the lemma, we have to specify out-of-equilibrium beliefs and strategies that support the equilibrium. Each expert's strategy consists of a price vector (\underline{p}, \bar{p}) and a recommendation policy for all possible price vectors. Each consumer's strategy consists of a visiting strategy for each set of posted price vectors and an acceptance behavior for each price vector. We first specify consumers' beliefs and acceptance behavior and experts' recommendation policy for each subgame beginning at the moment a consumer visits an expert characterized by her price vector (\underline{p}, \bar{p}) . Then we specify consumers' visiting and experts' price posting strategy. In the end we verify that we have indeed identified a weak PBE.

Let ω_i denote the probability that the expert recommends treatment \bar{c} given that consumer's problem is diagnosed to be $i \in \{m, M\}$, where m stands for minor, and M for major. Also, let $\bar{\mu}(\underline{\mu})$ be the probability a consumer assigns to the event that he has the major problem given that the expert has recommended treatment \bar{c} (\underline{c}). Similarly, let $\bar{\alpha}(\underline{\alpha})$ be the probability that a consumer accepts the recommendation \bar{c} (\underline{c}). Finally, let k denote a consumer's expected cost if he follows the proposed equilibrium strategy (that is, $k = (1 - h)(1 - \omega^*)(\underline{c} + \Delta^*) + [h + (1 - h)\omega^*]\bar{c} + d =$

$(\underline{c} + \Delta^*) + (\bar{c} - \underline{c} - \Delta^*)[h + (1 - h)\omega^*] + d$) and let π denote experts' profit per customer if they follow the proposed equilibrium strategy (that is, $\pi = (1 - h)[1 + \omega^*(1 - \alpha^*)]\Delta^*$). Suppose that consumers' beliefs are correct at the proposed price vector and that consumers' beliefs at other price vectors are given by (i) $\bar{\mu}(\underline{p}, \bar{p}) = 1$ and $\underline{\mu}(\underline{p}, \bar{p}) = 0$ iff $\underline{p} \leq d + (1 - \omega^*)(\underline{c} + \Delta^*) + \omega^*\bar{c}$ and $\bar{p} \in [\bar{c}, \bar{c} + d]$; and (ii) $\bar{\mu}(\underline{p}, \bar{p}) = h$ and $\underline{\mu}(\underline{p}, \bar{p}) = 0$ in any other case. Further suppose that consumers' acceptance decisions are given by (i) $\underline{\alpha}(\underline{p}, \bar{p}) = 1$ iff $\underline{p} \leq d + (1 - \omega^*)(\underline{c} + \Delta^*) + \omega^*\bar{c}$, and $\underline{\alpha}(\underline{p}, \bar{p}) = 0$ otherwise; and (ii) $\bar{\alpha}(\underline{p}, \bar{p}) = 1$ iff either $\underline{p} \leq d + (1 - \omega^*)(\underline{c} + \Delta^*) + \omega^*\bar{c}$ and $\bar{p} \leq \bar{c} + d$, or $\underline{p} > d + (1 - \omega^*)(\underline{c} + \Delta^*) + \omega^*\bar{c}$ and $\bar{p} \leq k$, and $\bar{\alpha}(\underline{p}, \bar{p}) = 0$ otherwise. Also suppose that deviating experts' recommendation policy is given by $\omega_m(\underline{p}, \bar{p}) = \omega_M(\underline{p}, \bar{p}) = 1$.⁶⁰ Finally suppose that consumers' visiting strategy prescribes not to visit a deviating expert, and that experts' price-posting strategy prescribes not to deviate to a price vector different from the proposed one.

Now we verify that we have indeed identified a weak PBE. First observe that consumers' acceptance strategies are indeed optimal given their beliefs: If a single expert deviates, the proposed price vector is still available since at least two experts post this vector in equilibrium. The expected cost to the consumer under that vector is k with prior beliefs, it reduces to $d + (1 - \omega^*)(\underline{c} + \Delta^*) + \omega^*\bar{c}$ if the consumer believes to need the cheap treatment for sure, and it increases to $\bar{c} + d$ if the consumer believes to need the expensive treatment for sure. Given this, consumers' acceptance strategies

⁵⁹ Here note that for any $(\bar{c}, \underline{c}, d, h)$ with $d < (\bar{c} - \underline{c})$ $(1 - h)$ there always exists a Δ^* such that both α^* and ω^* are in $(0,1)$.

⁶⁰ Here note that for out-of-equilibrium price-vectors satisfying $\underline{p} \leq d + (1 - \omega^*)(\underline{c} + \Delta^*) + \omega^*\bar{c}$ and $\bar{p} \in [\bar{c}, \bar{c} + d]$, the proposed beliefs are not consistent with the proposed equilibrium strategies. Sequential rationality and full consistency of beliefs at all information sets (full PBE) would require that for all out-of-equilibrium price-vectors with these properties experts and consumers play mixed strategies similar to the ones specified in Lemma 6.

are indeed optimal given their beliefs. Next observe that $\omega_m(p, \bar{p}) = \omega_M(p, \bar{p}) = 1$ is indeed optimal for a deviating expert, either because $\bar{\alpha}(p, \bar{p}) = 1$ and $\bar{p} \geq \bar{c}$, or because $\underline{\alpha}(p, \bar{p}) = \bar{\alpha}(p, \bar{p}) = 0$. Given deviating experts' recommendation policy and $\bar{p} \geq \bar{c}$, consumers obviously prefer to ignore any deviating expert. This establishes that no deviation to an alternative price vector will attract any customers and so (and since $\pi > 0$) no deviation will occur.

Part (ii): To verify part (ii) of the lemma it suffices to specify a profitable deviation. Consider a deviation to a price vector (p, \bar{p}) with $p \in (\underline{c}, d + (1 - \omega^*)(\underline{c} + \Delta^*) + \omega^*\bar{c})$ and $\bar{p} > \bar{c} + d$. First observe that, for a consumer served under such a price vector, it is optimal to set $\underline{\alpha}(p, \bar{p}) = 1$ and $\bar{\alpha}(p, \bar{p}) = 0$ for any belief that he might have.⁶¹ Given consumers' acceptance behavior, sequential rationality for the expert requires that she chooses $\omega_m = 0$ and $\omega_M = 1$.⁶² Given deviating experts' recommendation- and consumers' acceptance-behavior, the expected cost to a first-time consumer who plans to visit the deviator is $\hat{k} = d + (1 - h)p + h(d + \bar{c})$.⁶³ So, if $\hat{k} < k$ then the deviator attracts all first-time consumers and serves them at an expected profit of $\hat{\pi} = (1 - h)(p - \underline{c})$ per customer. The inequality $\hat{k} < k$ is equivalent to $p < \underline{c} + \Delta^* + \omega^*(\bar{c} - \underline{c} - \Delta^*) - dh/(1 - h)$, which (using the indifference equations (1) and (2) above) is again equivalent to $p < \underline{c} + \Delta^* + d\omega^*/[(1 - \omega^*)(1 - h)]$. Now suppose the deviating expert sets $p = \underline{c} + \Delta^* + d\omega^*/[(1 - \omega^*)(1 - h)] - \epsilon$, with ϵ

⁶¹ This follows from the fact that the only alternative price vector available if a single expert deviates is $(p, \bar{p}) = (\underline{c} + \Delta^*, \bar{c})$ and that the cost to the consumer under that vector is $d + (1 - \omega^*)(\underline{c} + \Delta^*) + \omega^*\bar{c}$ if the consumer believes to need the cheap treatment for sure while it is $\bar{c} + d$ if the consumer believes to need the expensive treatment for sure.

⁶² This follows from $p > \underline{c}$ (so that it is indeed optimal to recommend \underline{c} to a consumer with the minor problem) and $p < \bar{c}$ (so that it is never optimal to recommend \underline{c} to a consumer who needs \bar{c}), where the latter inequality is implied by consumers' indifference between accepting and rejecting a \bar{c} recommendation under $(p, \bar{p}) = (\underline{c} + \Delta^*, \bar{c})$ and by $p < d + (1 - \omega^*)(\underline{c} + \Delta^*) + \omega^*\bar{c}$.

strictly positive but smaller than $d\omega^*/[(1 - \omega^*)(1 - h)]$. Then she attracts all consumers and, thus, makes an expected profit $\hat{\pi} > (1 - h)\Delta^*$. The expected profit per expert in the proposed symmetric overcharging equilibrium is $\pi/n = (1 - h)[1 + (1 - \alpha^*)\omega^*]\Delta^*/n$ which is strictly less than $(1 - h)2\Delta^*/n$ implying that $\hat{\pi} < \pi/n$ even for $n = 2$. This establishes that the *weak PBE* sketched in part (i) of the lemma ceases to exist if experts are completely free in choosing prices since a profitable deviation always exists. ■

PROOF OF LEMMA 7: Again, liability prevents undertreatment, and the cost differential $(\bar{c} - \underline{c})$ prevents overtreatment. So each expert's provision policy is again trivial and we will therefore again focus on experts' price-posting and recommendation policy and on consumers' visiting and acceptance decisions. To show that the proposed strategies are part of a (*full*) *PBE* we have to specify reasonable out-of-equilibrium beliefs and strategies that support the equilibrium. Using the notation introduced in the proof of Lemma 6, we first specify consumers' beliefs and acceptance behavior and experts' recommendation policy for each subgame beginning at the moment a consumer visits an expert characterized by her price vector (p, \bar{p}) . Then we specify consumers' visiting and experts' price posting strategy. In the end we verify that we have indeed identified a *PBE*.

Let k denote the expected cost to the customer if he follows the proposed equilibrium strategy; that is, $k = d + (1 - h)\underline{c} + h(\bar{c} + d)$. Suppose that *consumers' beliefs* are correct at the proposed price vectors and that consumers' beliefs at other price vectors are given by (i) $\underline{\mu}(p, \bar{p}) = 0$ and $\bar{\mu}(p, \bar{p}) = 1$ if either $p = \bar{p} \leq \underline{c} + d$, or $\bar{p} > \bar{c} + d$ and $p < \bar{c}$; (ii) $\underline{\mu}(p, \bar{p}) = h$ and $\bar{\mu}(p, \bar{p}) = 1$ if $\bar{p} > \bar{c} + d$ and $p > \bar{c}$; (iii) $\underline{\mu}(p, \bar{p}) = 0$ and

⁶³ With probability $(1 - h)$ the consumer has the minor problem and gets treated for the price p by the deviator; with probability h he has the major problem, the deviator recommends \bar{c} , the consumer rejects, visits a nondeviator and gets the right treatment for the price \bar{c} .

$\underline{\mu}(p, \bar{p})$ is such that a customer who gets a \bar{c} recommendation is indifferent between accepting and rejecting whenever $\bar{p} \in (k, \bar{c} + d)$ and $p < \underline{c} + d$; and (iv) $\underline{\mu}(p, \bar{p}) = 0$ and $\bar{\mu}(p, \bar{p}) = h$ in any other case. Also suppose that *new consumers' (first time visitors') acceptance decisions* are given by (i) $\underline{\alpha}(p, \bar{p}) = 1$ for $p \leq k$ and $\underline{\alpha}(p, \bar{p}) = 0$ for $p > k$ whenever $\bar{p} > \bar{c} + d$ and $p > \bar{c}$; and $\underline{\alpha}(p, \bar{p}) = 1$ for $p \leq \underline{c} + d$ and $\underline{\alpha}(p, \bar{p}) = 0$ for $p > \underline{c} + d$ whenever either $\bar{p} \leq \bar{c} + d$ or $p \leq \bar{c}$; and (ii) $\bar{\alpha}(p, \bar{p}) = 1$ whenever $\bar{p} < k$; $\bar{\alpha}(p, \bar{p})$ is such that an expert who observes that the customer has the minor problem is exactly indifferent between recommending \underline{c} and recommending \bar{c} whenever $\bar{p} \in (k, \bar{c} + d)$ and $p < \underline{c} + d$; and $\bar{\alpha}(p, \bar{p}) = 0$ for any other price vector. Further suppose that the *acceptance decisions of consumers on their second visit* are given by (i) $\underline{\alpha}(p, \bar{p}) = 1$ for $p \leq k$ and $\underline{\alpha}(p, \bar{p}) = 0$ for $p > k$ whenever $\bar{p} > \bar{c} + d$ and $p > \bar{c}$; and $\underline{\alpha}(p, \bar{p}) = 1$ for $p \leq \underline{c} + d$ and $\underline{\alpha}(p, \bar{p}) = 0$ for $p > \underline{c} + d$ whenever either $\bar{p} \leq \bar{c} + d$ or $p > \bar{c}$; and (ii) $\bar{\alpha}(p, \bar{p}) = 1$ whenever $\bar{p} \leq \bar{c} + d$ and $\bar{\alpha}(p, \bar{p}) = 0$ otherwise. Also suppose that *experts recommend* in accordance with consumers' beliefs. Finally suppose that *consumers visiting strategy* prescribes not to deviate from the proposed visiting behavior provided no deviating expert offers either $p < \underline{c}$ and $\bar{p} > \bar{c} + d$, or $\bar{p} < \bar{c}$, and that *experts' price-posting strategy* prescribes not to deviate to price vectors different from the proposed ones.

Now we verify that we have indeed identified a PBE. First observe that customers beliefs reflect experts' incentives: The expert recommends honestly (i.e., $\omega_m = 0$ and $\omega_M = 1$) if $\bar{p} < \underline{c} + d$ and $p = \bar{p}$, and if $\bar{p} > \bar{c} + d$ and $p < \bar{c}$, either since $\underline{\alpha}(p, \bar{p}) = \bar{\alpha}(p, \bar{p}) = 1$ and $p = \bar{p}$ (in the former case) or since $\underline{\alpha}(p, \bar{p}) = 1, \bar{\alpha}(p, \bar{p}) = 0$ and $p < \bar{c}$ (in the latter case); the expert always recommends \underline{c} (i.e., $\omega_m = 0$ and $\omega_M = 0$) if $\bar{p} > \bar{c} + d$ and $p > \bar{c}$, since $\bar{\alpha}(p, \bar{p}) = 0$ and $p > \bar{c}$; the expert recommends \bar{c} if the consumer has the major problem and she randomizes between

recommending \underline{c} and recommending \bar{c} if he has the minor problem (i.e., $\omega_m \in (0, 1)$ and $\omega_M = 1$) whenever $\bar{p} \in (k, \bar{c} + d)$ and $p \in (\underline{c}, \underline{c} + d)$ since $\underline{\alpha}(p, \bar{p}) = 1$ and $\bar{\alpha}(p, \bar{p})$ is such that she is exactly indifferent between both recommendations; and the expert recommends the expensive treatment (i.e., $\omega_m = 1$ and $\omega_M = 1$) in all other cases, either because $\bar{\alpha}(p, \bar{p}) = 1$ and $\bar{p} > p$, or because $\underline{\alpha}(p, \bar{p}) = \bar{\alpha}(p, \bar{p}) = 0$. Next observe that customers' strategies are optimal given their beliefs. First consider *consumers' acceptance strategies*. If a single expert deviates, the proposed equilibrium offers are still available since at least two experts make each offer. Thus, the above described acceptance strategies are optimal given consumers' beliefs. Next consider *new consumers' visiting strategy*. If no expert deviates, the relevant alternatives are (i) to visit a cheap expert first and to reject the \bar{c} recommendation as proposed, or (ii) to visit an expensive expert first and to accept the \bar{c} recommendation. The former strategy has an expected cost of $d + (1 - h)\underline{c} + h(\bar{c} + d)$, the latter a cost of $\bar{c} + d$, while the benefit is the same. Since $(\bar{c} - \underline{c})(1 - h)/h > (\bar{c} - \underline{c})(1 - h) > d$, the former cost is strictly lower and customers' visiting strategy is optimal if no expert deviates. Given that the equilibrium offers are still available if a single expert deviates, and given the above specified beliefs, no customer has an incentive to visit a deviating expert if her posted prices do not satisfy either $p < \underline{c}$ and $\bar{p} > \bar{c} + d$, or $\bar{p} < \bar{c}$. Finally observe that no expert has an incentive to deviate. Deviations to prices satisfying $p < \underline{c}$ and $\bar{p} > \bar{c} + d$ are unattractive since only the \underline{c} recommendation is accepted. Deviations to price vectors where $p > \underline{c}$ and $\bar{p} \geq \bar{c}$ are unattractive since they attract no customers. Price vectors where $\bar{p} < \bar{c}$ are unprofitable too, if they only attract customers who first visit a cheap expert, and, if recommended the expensive treatment, resort to the deviator. The expected cost to the customer of this latter strategy is $d + \underline{c} + h(\hat{p} + d - \underline{c})$ where \hat{p} is the price

posted by the deviator for the expensive treatment. Going directly to the deviator and accepting her recommendation, on the other hand, costs at most $\hat{p} + d$. Thus, in order to avoid being visited only by consumers with the major problem, the deviation must satisfy $d + \underline{c} + h(\hat{p} + d - \underline{c}) > d + \hat{p}$, which is equivalent to $\hat{p} < \underline{c} + dh/(1 - h)$. To cover expected treatment cost the price \hat{p} must also satisfy $\hat{p} \geq \underline{c} + h(\bar{c} - \underline{c})$. But, $\underline{c} + h(\bar{c} - \underline{c}) > \underline{c} + dh/(1 - h)$ since $(1 - h)(\bar{c} - \underline{c}) > d$. This proves that no deviation by an expert is profitable. ■

REFERENCES

- Alger, Ingela, and François Salanié. 2003. "A Theory of Fraud and Over-Consumption in Experts Markets." Boston College. Mimeo.
- Akerlof, George A. 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*, 84(3): 488–500.
- Darby, Michael R., and Edi Karni. 1973. "Free Competition and the Optimal Amount of Fraud." *Journal of Law and Economics*, 16(1): 67–88.
- Dulleck, Uwe, and Rudolf Kerschbamer. 2005a. "Experts vs. Discounters—Competition and Market Unravelling When Customers Do Not Know What They Need." University of Vienna. Mimeo.
- Dulleck, Uwe, and Rudolf Kerschbamer. 2005b. "Price Discrimination Via the Choice of Distribution Channels." University of Vienna Department of Economics Working Paper 0511.
- Ely, Jeffrey C., Drew Fudenberg, and David K. Levine. 2005. "When is Reputation Bad?" UCLA. Levine's Working Paper Archive.
- Ely, Jeffrey C., and Juuso Välimäki. 2003. "Bad Reputation." *Quarterly Journal of Economics*, 118(3): 785–814.
- Emons, Winand. 1997. "Credence Goods and Fraudulent Experts." *RAND Journal of Economics*, 28(1): 107–19.
- Emons, Winand. 2001. "Credence Goods Monopolists." *International Journal of Industrial Organization*, 19(3-4): 375–89.
- Fong, Yuk-Fai. 2005. "When Do Experts Cheat and Whom Do They Target?" *Rand Journal of Economics*, 36(1): 113–30.
- Fuchs, Victor R. 1978. "The Supply of Surgeons and the Demand for Operations." *Journal of Human Resources*, 13: 35–56.
- Fudenberg, Drew, and Jean Tirole. 1991. *Game Theory*. Cambridge and London: MIT Press.
- Glazer, Jacob, and Thomas G. McGuire. 1996. "Price Contracts and Referrals in Markets for Services." Tel Aviv University Working Paper 10/96.
- Gruber, Jon, John Kim, and Dina Mayzlin. 1999. "Physician Fees and Procedure Intensity: The Case of Cesarean Delivery." *Journal of Health Economics*, 18(4): 473–90.
- Hughes, David, and Brian Yule. 1992. "The Effect of Per-Item Fees on the Behaviour of General Practitioners." *Journal of Health Economics*, 11(4): 413–37.
- Mas-Colell, Andreu, Michael Whinston, and Jerry R. Green. 1995. *Microeconomic Theory*. New York: Oxford University Press.
- Myerson, Roger B. 1991. *Game Theory: Analysis of Conflict*. Cambridge and London: Harvard University Press.
- Nelson, Phillip. 1970. "Information and Consumer Behavior." *Journal of Political Economy*, 78(2): 311–29.
- Park, In-Uck. 2005. "Cheap-Talk Referrals of Differentiated Experts in Repeated Relationships." *Rand Journal of Economics*, 36(2): 391–411.
- Pesendorfer, Wolfgang, and Asher Wolinsky. 2003. "Second Opinions and Price Competition: Inefficiency in the Market for Expert Advice." *Review of Economic Studies*, 70(2): 417–37.
- Pitchik, Carolyn, and Andrew Schotter. 1987. "Honesty in a Model of Strategic Information Transmission." *American Economic Review*, 77(5): 1032–36.
- Pitchik, Carolyn, and Andrew Schotter. 1988. "Honesty in a Model of Strategic Information Transmission: Correction." *American Economic Review*, 78(5): 1164.
- Pitchik, Carolyn, and Andrew Schotter. 1993. "Information Transmission in Regulated Markets." *Canadian Journal of Economics*, 26(4): 815–29.
- Richardson, Hugh. 1999. "The Credence Good Problem and the Organization of Health Care Markets." Texas A&M University. Mimeo.
- Sülzle, Kai, and Achim Wambach. 2005. "Insurance in a Market for Credence Goods." *Journal of Risk and Insurance*, 72(1): 159–76.
- Taylor, Curtis R. 1995. "The Economics of Breakdowns, Checkups, and Cures." *Journal of Political Economy*, 103(1): 53–74.
- Tirole, Jean. 1988. *The Theory of Industrial Organization*. Cambridge and London: MIT Press.
- Wolinsky, Asher. 1993. "Competition in a Market for Informed Experts' Services." *RAND Journal of Economics*, 24(3): 380–98.
- Wolinsky, Asher. 1995. "Competition in Markets for Credence Goods." *Journal of Institutional and Theoretical Economics*, 151(1): 117–31.