

2022-09-22

Equilibrium information in credence goods; Oslo Metropolitan University seminar presentation

Liu, Ting and Ma, Ching-To Albert, Equilibrium Information in Credence Goods. Available at SSRN: <https://ssrn.com/abstract=4326242> or <http://dx.doi.org/10.2139/ssrn.4326242>
<https://hdl.handle.net/2144/46575>

"Downloaded from OpenBU. Boston University's institutional repository."

Equilibrium Information in Credence Goods

Ting Liu
Department of Economics
Stony Brook University
ting.liu@stonybrook.edu

Ching-to Albert Ma
Department of Economics
Boston University
ma@bu.edu

July 2022

Abstract

We study credence goods in a general model. A consumer may suffer a loss which is a continuous random variable. Privately observing the loss value, an expert can provide a repair at a price to eliminate the consumer's loss. All perfect-Bayesian equilibria are inefficient, in that some losses are not repaired. In closed form, we derive a pooling equilibrium (where losses are inferred to be in an interval), and a separating equilibrium (where losses are precisely inferred). If the expert can acquire an information structure on losses, the first best is achieved by a binary signal. Results are robust when cost and loss are random and correlated, and when there are multiple experts.

Keywords: credence goods, experts, separating equilibrium, pooling equilibrium, information acquisition

JEL: D80, D82, D83

Acknowledgment: We thank Jacobo Bizzotto, Lester Chan, Winand Emons, Sanghoon Kim, Henry Mak, Chiara Margaria, Juan Ortner, and seminar participants at Boston University, Stony Brook University, International Industrial Organization Conference 2022, Midwest Theory Conference 2021, Econometric Society Summer Meetings 2022, International Game Theory Festival 2022, and Webinar Series on Credence Goods and Expert Markets for their suggestions and comments.

1 Introduction

Credence goods and services have qualities and values unobserved by consumers before and after purchase or consumption. For example, car owners rely on mechanics for diagnosis and treatment. After a repair, car owners often cannot ascertain whether the performed services are necessary. Consumers face the same challenge when they need appliance repair, financial advice, healthcare consultation, and legal service.

Since Darby and Karni (1973) introduced the credence-good concept, a large literature has focused on sellers' or experts' incentives to defraud clients. Most existing literature analyzes experts' recommendation strategies and market efficiency in a stylized model with two problems and two treatments.¹ Although the binary model is convenient, it is too simplistic for studying complex problems, which are now common due to advanced technologies. According to automobile.com, computer systems have made new vehicles increasingly complex, which in turn creates complex repairs.² For example, a car's check engine light can indicate more than ten problems, from a loose fuel cap to a defective catalytic converter.³ Different problems reduce a car's performance to different degrees, so the consumer's willingness to pay for repair depends on which fault is related to the check engine light. Obviously this complexity arises in markets other than automobile repair.

We propose a general model to study complex credence-good problems. Our paper departs from the literature in two main aspects. First, we expand the binary model to allow for a continuum of losses. The continuous loss setup can approximate any finite number of losses, and is more tractable than a setup with a finite number of losses. Second, we endogenize the expert's diagnosis precision. Most of the extant literature has assumed that the expert obtains a perfect signal about the loss. We allow the expert to choose an information structure to find out about the loss.

We ask two sets of questions for each of the two aspects. What new insights can the continuous loss model throw on complex problems? Are the insights of the binary model robust in the continuum model? Is

¹Pesendorfer and Wolinsky (2003) study a model with a continuum of problems. We elaborate on the difference between their model and ours in the literature review.

²"Managing vehicle complexity with a seamless diagnostics experience."

³"11 Reasons for the Check Engine Light," bumper.com

the expert better off with more accurate diagnostic information? How much diagnostic information should the expert acquire?

Our basic model is the following. A consumer realizes that she may suffer from a loss, but cannot tell the severity, so must consult an expert. The expert makes a diagnosis and recommends a price for repair. The consumer then decides whether to accept the expert’s offer. If she accepts the offer, the loss will be avoided. This corresponds to an assumption of “Liability” in the literature—the expert is liable for the repair. Our main model analyzes the case that the expert’s diagnosis is perfect. We then extend the model to analyze endogenous diagnosis precision.

The new insight of our model is the existence of a pure-strategy pooling equilibrium. Here, the expert recommends a price if the consumer’s loss is below a threshold and refuses to provide repair otherwise; the consumer accepts the (equilibrium) price with probability one and rejects any higher prices. The recommended price is the consumer’s average valuation conditional on losses being below the threshold, and this price is also the expert’s repair cost at the threshold. Despite the complexity due to continuum of losses, the pooling equilibrium is remarkably simple. The expert’s equilibrium price offer is based on the marginal repair cost; the consumer’s acceptance decision is based on average valuations conditional on losses below the threshold. This contrast is reminiscent of the lemons problem à la Akerlof (1970). Because of the discrepancy between marginal and average, the pooling equilibrium will leave some high losses unrepaired, even though it is efficient to do so. The pooling equilibrium has not appeared in the extant literature because the binary model excludes any loss between the highest and the lowest levels.

The main insight of the binary model has been that consumers use probabilistic rejection to discipline experts.⁴ If a consumer trusts an expert completely and always accepts his recommendation, the expert will always recommend the expensive treatment, yielding the consumer a net loss. Hence, to sustain an equilibrium the consumer must randomize between accepting and rejecting the expensive treatment. The consumer’s probabilistic rejection keeps the expert in check. Fong (2005) shows the existence of an honest equilibrium in which the consumer perfectly infers her losses from the expert’s recommendations, but the

⁴See Pitchik and Schotter (1987), Wolinsky (1993), Fong (2005), Fong et al. (2018, 2022).

consumer rejects the expensive treatment with a positive probability. Moreover, the honest recommendation strategy maximizes profits and efficiency.

We also derive a unique separating equilibrium which generalizes the honest equilibrium in the binary model. In the separating equilibrium, the expert charges the consumer her loss from the problem; the consumer rejects each price with a positive probability, and the rejection probability increases in the offered price. The expert's equilibrium profit and repair likelihood decrease in the loss. The technique for characterizing the separating equilibrium is similar to the envelope argument commonly used in contract theory. Because the expert has incentives to exaggerate the loss, incentive compatibility implies the expert earning higher information rents when losses are lower. To support the separating equilibrium, the consumer's rejection probability is a sophisticated function of the price offer.⁵ The consumer's mixed strategy may seem appealing theoretically but is hard to implement in practice. Although the insight of the binary model is robust, its application to complex problems is arguably limited.

The pooling equilibrium is more compelling than the separating equilibrium. First, the pooling equilibrium is a pure-strategy equilibrium, whereas the separating equilibrium requires the consumer to reject each price according to carefully calibrated probabilities. Second, the consumer earns a zero ex ante surplus in the pooling equilibrium and a zero ex post surplus in the separating equilibrium. In the separating equilibrium, no consumer should complain about being cheated. In the pooling equilibrium, consumers with low losses will realize, if they eventually learn about their losses, that they have been overcharged—the equilibrium price being above the loss value. But consumers with high losses will discover that the price they have paid is below the loss value. The same expert, ex post, may be praised by some consumers but complained against by others.⁶ This is consistent with the anecdotal evidence that experts often have mixed reviews.⁷

The welfare comparison between pooling and separating equilibria coincides with the equilibrium profit comparison because consumers always have zero equilibrium surplus. Both separating and pooling equilibria

⁵Formally, this probability is the solution of a differential equation.

⁶The credence good model is identical to one in which losses become observed if repairs are not sought, but repair is no longer feasible once losses have occurred.

⁷Consumer Federation of America finds that auto and home improvement are the top two consumer complaints in 2020. https://consumerfed.org/press_release/nations-top-10-consumer-complaints/

result in insufficient repairs. The separating equilibrium leads to under-treatment at all losses (except for the lowest one). In the pooling equilibrium, losses are repaired if and only if they are below the threshold. Hence, the pooling equilibrium yields more surplus than the separating equilibrium for low losses and vice versa for high losses. We characterize the condition under which the pool equilibrium is more profitable and efficient than the separating equilibrium, and the condition depends on the distribution of the losses and the shape of the cost function. The result that the pooling equilibrium may dominate the separating equilibrium stands in sharp contrast to the well-known result of the binary model that honesty maximizes profit and efficiency (Fong (2005)).

We also study the expert's information acquisition decision in the general model. We let the expert choose an information structure, which generates a signal about the consumer's problem. After observing the signal, the expert makes the price offer. The information structure is public information, whereas the signal is the expert's private information. The information structure can be interpreted as diagnosis devices or tests. For example, a car mechanic can use code readers or scan tools for diagnosis. Code readers report fault codes; scan tools report detailed diagnostics for each code and is more accurate than scan tools. A car owner can verify whether a code reader or a scan tool is used for diagnosis but cannot interpret the codes. Hence, the diagnosis is the mechanic's private information. Similarly, a dentist can use radiography or the more accurate Cone-Beam Computed Tomography (CBCT) for x-rays. Whereas patients know which technology is used, they cannot interpret the scans.

The expert will optimally choose imperfect diagnostic information even when information acquisition is costless. In fact, the expert can achieve the first-best profit with a binary-signal information structure, which shows whether the loss is below or above a certain cutoff. Because the expert's price is based on the signal, the optimal price depends on whether the loss is above or below that cutoff. Likewise, when the consumer receives a price offer, she knows that it is based on a binary signal. In other words, each player bases decisions on average losses and costs conditional on binary signals, so the Akerlof's lemons problem is avoided. The insight is that the expert should devote resources to acquire coarse diagnostic information, not perfect information.

Our results are robust against (i) two dimensions of uncertainty, (ii) multiple experts, and (iii) a more general cost function. Under two-dimensional uncertainty, the consumer’s loss and the repair cost are both random, but positively correlated. The previous single dimensional model assumes that the cost is a deterministic and increasing function of loss; in that case, information about loss is equivalent to information about cost. For the two-dimensional, loss-cost uncertainty model, the expert’s equilibrium price offer can only be a function of the expert’s privately observed cost. In equilibrium, the expert’s information about the consumer’s loss cannot be used to set prices. The intuition is that the expert only cares about his own cost and whether his price offer will be accepted, never the consumer’s loss. The two dimensional model thus reduces to the single dimensional model.

We also consider the case of multiple experts with consumer search. We show that both the pooling and the separating equilibria in the single-expert market remain equilibria in the multiple-expert market. The idea is that any expert will follow the same equilibrium strategy as in the single-expert model. The consumer recognizes that experts will make the same price offer, so visiting more experts does not pay. In sum, the separating and the pooling equilibria in the single-expert model are robust. Then we discuss the case in which it is inefficient to fix the problem when the loss is either very high or very low.⁸ We show that the pooling equilibrium is robust but the separating equilibrium fails to exist. Moreover, the pooling equilibrium features over provision of treatment for low losses and under provision of treatment for high losses.

The rest of the paper is organized as follows. The next subsection is a literature review. Section 2 presents the model. Section 3 derives properties that hold for all equilibria. We characterize pooling and separating equilibria in closed form. Next, Section 4 is about expert acquiring information. Section 5 has a first subsection that presents multiple pooling equilibria, but introduces a natural refinement to reject all but the one with the highest equilibrium price; then the next subsection characterizes what we call “hybrid” equilibria. Section 6 considers three extensions: a two-dimensional model in loss and cost, multiple experts, and an alternative cost function. We offer some concluding remarks in Section 7.

⁸For example, an automobile’s very minor problems may not be worth the repair cost, and the same may be true for very serious problems. Repairs are cost effective for intermediate conditions.

1.1 Related literature

Darby and Karni (1973) introduced “credence goods” and discussed firms’ incentives to mislead consumers to raise demand. The subsequent literature analyzes sellers’ strategic behavior in game-theoretic models with binary losses. Dulleck and Kerschbamer (2006) and Balafoutas and Kerschbamer (2020) are comprehensive literature surveys.

The extant literature has taken two directions. The first line of the literature assumes that the expert is liable for repairing consumers’ problems once consumers agree to price offers; the expert must then work to eliminate the loss. Papers along this line of research include Pitchik and Schotter (1987), Wolinsky (1993, 1995), Taylor (1995), Fong (2005), Liu (2011), Fong et al. (2018, 2021). The second line assumes that only the expert’s repair or treatment expense is verifiable; the expert must carry out the promised work, but is not responsible for the loss elimination. This line of research includes Emons (1997), Alger and Salanie (2006), Dulleck and Kerschbamer (2009), Fong et al. (2014), Bester and Dahm (2018), and Chen et al. (2021).

The importance of the distinction between these two lines of work is this. Where cost is nonverifiable, cheating or hiding loss information can be regarded as costless because the expert only performs the minimal repair but is free to inflate the loss. However, when cost is verifiable, the expert will have to carry out wasteful work to get a higher price, so cheating or hiding loss information can be regarded as costly. We assume that the expert is liable for the repair, and that treatment cost is not verifiable, so our work belongs to the first line of the literature. We do extend the model to one with uncertain cost and loss; there neither loss nor cost is verifiable.

The binary model under the liability assumption has never exhibited any nontrivial pure-strategy equilibria. However, we show that in our general model pooling equilibria always exist, and they can be more profitable and efficient than a separating equilibrium. This finding draws a sharp contrast to the idea that the expert makes the most profit by telling the truth (for example, Fong (2005)).

Most of the existing literature assumes that experts can perfectly diagnose the consumer’s problem at zero cost. There is a small but growing literature on the expert’s incentives to acquire diagnostic information by costly hidden actions; these papers include Pesendorfer and Wolinsky (2003), Dulleck and Kerschbamer

(2009), Bester and Dahm (2018), and Chen et al. (2021). In particular, Pesendorfer and Wolinsky (2003) used a model with a continuum of losses to study experts' information acquisition incentives. Given a continuum of losses, an expert knows the true loss if and only if effort is used to acquire information. If two experts agree, both must have acquired information, so the consumer then can pick one among them randomly for treatment. These papers study an environment in which treatment is verifiable and information acquisition is costly. Our paper here complements them because we assume that treatment is not verifiable and information acquisition is costless. We are unaware of a result for the first-best outcome when the expert may choose to acquire a general information structure. Our paper also shows that acquiring too much information may diminish profits.

A more recent literature has implemented tests of credence-good models with data from the laboratory and from the real world. Most of the laboratory experiments are based on the binary model (Dulleck et al. (2011), Kerschbamer et al. (2015), Mimra et al. (2016), Balafoutas et al. (2021)). Hence, the subject chooses between whether to make an honest recommendation or not. Our framework permits a general strategy space for the expert; in such an experiment, subjects would have a rich set of ways to mislead consumers. In field experiments, subjects of course are quite unrestricted in their strategies; see, for example, Schneider (2012), Balafoutas et al. (2013), Liu et al. (2019), and Gottschalk et al. (2020), and our general model does allow such rich strategy sets. Our work therefore provides a foundation for future experiments and empirical analyses.

Our work is broadly related to the literature on trading under asymmetric information. Bagwell and Riordan (1991) and Wolinsky (1983) study how firms can signal their goods' qualities by prices. In both papers, there exists a separating equilibrium in which a high price signals a high quality. A common driving force for the separating equilibrium is that some consumers are informed about the quality prior to consumption. In our model we also identify a separating equilibrium in which a high price signals a high loss. Here, however, consumers are never informed about their losses before accepting the expert's offer or after the repair. Kim (2012) studies endogenous market segmentation as a way to alleviate information asymmetry. In Kim (2012), sellers send a cheap-talk message to buyers before trade takes place. Kim constructs an equilibrium in which some low-quality sellers reveal their quality, whereas some other low-quality sellers

pool with high-quality sellers. In our paper, the seller’s recommendation is the price, which directly affects the consumer’s payoff. We do have hybrid equilibria in which some seller types reveal information through prices, but other seller types pool together by offering the same price.

2 General credence-good model

A risk-neutral consumer has been enjoying a utility B from a good or a service. However, it has come to light, say from a fault indicator, that the baseline utility B may be reduced. The potential utility or monetary loss is denoted by ℓ and is randomly distributed on the strictly positive support $[\underline{\ell}, \bar{\ell}]$, with an absolutely continuous distribution⁹ F , and we denote its density by $f \equiv F'$. Any potential loss can be avoided by a risk-neutral expert’s repair. For now, we assume that the consumer interacts with one expert only, but in Subsection 6.2, we let the consumer interact with many experts. If the loss turns out to be ℓ , the expert can incur a cost $C(\ell)$ to eliminate it, where $C : [\underline{\ell}, \bar{\ell}] \rightarrow \mathbb{R}_+$ is a strictly increasing and differentiable function. The distribution of ℓ and the cost function of ℓ , respectively F and C , as well as the baseline utility B , are assumed to be common knowledge. We assume $\ell > C(\ell)$ if and only if $\ell < \widehat{\ell}$, with $\underline{\ell} < \widehat{\ell} \leq \bar{\ell}$, which says that it is efficient to treat only problems with losses lower than $\widehat{\ell}$. In Subsection 6.3, we discuss the case in which it is inefficient to treat problems with very low or very high losses.

The credence-good nature of the model is this. The expert perfectly observes the potential loss ℓ but the consumer does not observe it once her problem is repaired.¹⁰ In Section 5, we endogenize the expert’s information acquisition. The function C exhibits a monotone relationship between the expert’s cost and loss, so we assume that the consumer never learns the expert’s repair cost either; in other words, repair cost is not verifiable. In Subsection 6.1, we will relax the monotonicity assumption; there, costs and losses are random, but positively correlated. The only verifiable event here is that the consumer gets to enjoy the utility B upon a repair. The expert’s offer to repair the good at some price is a contract that lets the consumer continue to

⁹The absolute continuity of F ensures that its derivative exists almost everywhere for $\ell \in [\underline{\ell}, \bar{\ell}]$.

¹⁰If the repair is not done, the consumer’s utility eventually becomes $B - \ell$. At that point, the consumer can infer the loss value, but repair at that time has become infeasible.

enjoy B upon paying some repair price.¹¹

If the expert repairs the loss ℓ at a price p , his profit is $p - C(\ell)$; otherwise, profit is zero. If the consumer accepts the repair offer at price p , her utility is $B - p$; otherwise utility will become $B - \ell$. Because loss ℓ belongs to $[\underline{\ell}, \bar{\ell}]$, we can let prices belong to $[\underline{\ell}, \bar{\ell}]$ without any loss of generality. We denote the expert's refusal to repair by the notation of a price set at $+\infty$. The extensive form between the consumer and the expert is as follows:

Stage 1: A consumer visits the expert, who then observes the consumer's potential loss $\ell \in [\underline{\ell}, \bar{\ell}]$, drawn from distribution F . The expert chooses between refusing to repair and offering to repair at a price $p \in [\underline{\ell}, \bar{\ell}]$.

Stage 2: If the consumer agrees to the repair, she pays the expert p , and the expert incurs cost $C(\ell)$ for the repair. If the consumer rejects the expert's offer, she suffers the loss.

The expert's strategy is a price function that maps the consumer's losses to prices: $P : [\underline{\ell}, \bar{\ell}] \rightarrow [\underline{\ell}, \bar{\ell}] \cup \{+\infty\}$, with $P(\ell) \in [\underline{\ell}, \bar{\ell}]$ denoting the expert's offer to eliminate loss ℓ and $+\infty$ denoting the expert's refusal to treat the consumer. The consumer's strategy is a probability acceptance function: $\alpha : [\underline{\ell}, \bar{\ell}] \rightarrow [0, 1]$, with $\alpha(p)$ denoting the consumer's probability of accepting a repair offer $p \in [\underline{\ell}, \bar{\ell}]$. Given the expert's strategy, the consumer updates her belief about her loss according to Bayes rule when it is possible. We characterize perfect-Bayesian equilibria.

Each price constitutes an information set. There can be many unreached information sets in an equilibrium. For example, suppose that an equilibrium price function specifies that $P([\underline{\ell}, \tilde{\ell})) = p_1$ and $P([\tilde{\ell}, \bar{\ell}]) = p_2$, so the consumer responds to prices p_1 and p_2 on the equilibrium path. Prices other than p_1 and p_2 never get offered in equilibrium. Yet, the consumer's equilibrium strategy must still specify her acceptance probability at these prices. Clearly, the consumer's responses at off-equilibrium price offers may support many equilibria, but we will impose belief restrictions later to rule out many.

We do not include a price-posting stage in the extensive form, which has been commonly used in earlier

¹¹The expert does not have the option of repairing a fraction of the loss ℓ by incurring some fraction of the cost $C(\ell)$. This assumption is consistent with the postulate that only the baseline utility B is verifiable.

papers. Here, the expert does not restrict himself to offering a subset of prices before he learns the loss. In complex repairs, it is uncommon that experts would make binding estimates before diagnosis. For example, according to NBCNews, few hospitals in the U.S. post prices of common procedures online, despite a federal requirement to do so since 2021. More important, a combination of procedures may be needed for a repair, so posting prices of each procedure does not necessarily indicate the total repair cost. Car mechanics or technicians may have hourly labor rates, but only decide on the required hours for repair after diagnosis.

3 Equilibria

The continuous loss model admits a large set of equilibria because the expert’s pricing strategy is a function of losses and the consumer’s acceptance strategy is a function of prices. We begin the analysis by deriving necessary conditions on players’ equilibrium strategies and the expert’s equilibrium profit. Guided by the equilibrium conditions, we then derive in closed form the pooling and separating equilibria.

3.1 Perfect-Bayesian equilibrium strategies

Consider any perfect-Bayesian equilibrium (P, α) , consisting of the expert’s pricing function $P : [\underline{\ell}, \bar{\ell}] \rightarrow [\underline{\ell}, \bar{\ell}] \cup \{+\infty\}$ and the consumer’s acceptance probability function $\alpha : [\underline{\ell}, \bar{\ell}] \rightarrow [0, 1]$. (We omit the consumer’s response—the empty set—when the expert refuses to offer treatment.) Let $Range(P) \subset [\underline{\ell}, \bar{\ell}]$ denote the set of expert’s equilibrium prices for treatment. Those prices in $[\underline{\ell}, \bar{\ell}] \setminus Range(P)$ are off-path prices or unreached information sets. The acceptance function α is defined over all elements of $[\underline{\ell}, \bar{\ell}]$ which may be a strict superset of $Range(P)$.

Next, for any given equilibrium strategy profile (P, α) , define a function $\tilde{\alpha}_P : [\underline{\ell}, \bar{\ell}] \rightarrow [0, 1]$ by $\tilde{\alpha}_P(\ell) \equiv \alpha(P(\ell))$. The function $\tilde{\alpha}_P$ tracks the acceptance probability on the equilibrium path according to strategy profile (P, α) , over all loss values; it does so by compounding the equilibrium acceptance probability and the equilibrium price function. Let $\Pi(\ell) \equiv \tilde{\alpha}_P(\ell) \times \{P(\ell) - C(\ell)\}$ denote the equilibrium profit from repairing ℓ . The expert’s expected equilibrium profit is $\int \Pi(\ell) dF(\ell)$. We now state properties valid for all equilibria; all proofs of results (unless omitted) are in the Appendix.

Proposition 1 *Any perfect-Bayesian equilibrium strategy profile (P, α) has the following properties:*

- i) The equilibrium acceptance probability $\tilde{\alpha}_P(\ell)$ is weakly decreasing.*
- ii) The price $P(\ell)$ is weakly increasing if $\tilde{\alpha}_P(\ell) > 0$, $\ell \in [\underline{\ell}, \bar{\ell}]$.*
- iii) Equilibrium profit $\Pi(\ell)$ is weakly decreasing in $\ell \in [\underline{\ell}, \bar{\ell}]$ and is strictly decreasing if $\tilde{\alpha}_P(\ell) > 0$ for $\ell \in [\ell_1, \ell_2] \subset [\underline{\ell}, \bar{\ell}]$.*

Any perfect-Bayesian equilibrium (P, α) must satisfy the following no-deviation conditions: at ℓ' and ℓ

$$[P(\ell') - C(\ell')] \tilde{\alpha}_P(\ell') \geq [P(\ell) - C(\ell')] \tilde{\alpha}_P(\ell) \quad (1)$$

and

$$[P(\ell) - C(\ell)] \tilde{\alpha}_P(\ell) \geq [P(\ell') - C(\ell)] \tilde{\alpha}_P(\ell'). \quad (2)$$

Part i) of Proposition 1 is obtained by adding these two inequalities and collecting terms, a common step in asymmetric-information games. Part ii) says that the expert must not decrease his price when the repair cost goes up if the offer may be accepted. Together, these are intuitive results. Reducing prices when costs increase would just be against profit maximization. Parts i) and ii) together say that the consumer must not play an equilibrium strategy that results in higher acceptance probabilities when prices increase. Otherwise, the expert would deviate to recommending the prices meant for high losses when the consumer's loss is low. As the loss and cost increase, the price increases but the acceptance rate decreases. Part iii) says that the equilibrium profit must decrease, and this is due to incentive compatibility.

According to Proposition 1, the equilibrium price is (i) constant in losses, (ii) strictly increasing in losses, or (iii) alternating between constant and increasing in losses. The consumer cannot distinguish her losses which are charged the same price in case (i) and can perfectly infer her losses from the price in case (ii).

In what follows, we characterize the first two types of equilibria and begin by defining pooling and separating equilibria in terms of equilibrium allocations with positive-probability repairs. We postpone the discussion of the third type of equilibria to Subsection 5.2.

Definition 1 (Pooling equilibrium) *An equilibrium (P, α) is said to be a pooling equilibrium if for any ℓ and ℓ' in $[\underline{\ell}, \bar{\ell}]$ for which $\alpha(P(\ell)) > 0$ and $\alpha(P(\ell')) > 0$, $P(\ell) = P(\ell')$.*

Definition 2 (Separating equilibrium) *An equilibrium (P, α) is said to be a separating equilibrium if for any ℓ and ℓ' in $[\underline{\ell}, \bar{\ell}]$, $\ell \neq \ell'$, for which $\alpha(P(\ell)) > 0$ and $\alpha(P(\ell')) > 0$, then $P(\ell) \neq P(\ell')$.*

In both definitions, we ignore those prices that will never be accepted in an equilibrium. A pooling price is one that will be accepted with some positive probability, but does not vary with losses. Separating prices are those that will be accepted with some positive probability, but vary according to losses. Notice that the requirement that prices vary in a separating equilibrium and that prices remain constant in a pooling equilibrium is for *all* losses in $[\underline{\ell}, \bar{\ell}]$ for which equilibrium acceptance may happen. Hence, these two definitions are not exhaustive. Continuum of equilibria that are neither separating nor pooling do exist. In anticipation, we call these *hybrid equilibria*, which will be defined and analyzed in Subsection 5.2

3.2 A pooling equilibrium

To begin, we set up some preliminary concepts for a pooling equilibrium. First, we denote the expected loss by $\mu \equiv \int \ell dF(\ell)$. Next, we define conditional average loss values for a loss interval by $AL : [\underline{\ell}, \bar{\ell}] \times [\underline{\ell}, \bar{\ell}] \rightarrow [\underline{\ell}, \bar{\ell}]$ where

$$AL(\ell_1, \ell_2) \equiv \frac{\int_{\ell_1}^{\ell_2} x dF(x)}{F(\ell_2) - F(\ell_1)}.$$

Of particular relevance here is the expected loss conditional on the loss below ℓ , $AL(\underline{\ell}, \ell)$. Clearly, $AL(\underline{\ell}, \ell) < \ell$, and $\lim_{\ell \rightarrow \bar{\ell}} AL(\underline{\ell}, \ell) = \mu$.

We define a set of losses LP by

$$LP \equiv \{\ell : AL(\underline{\ell}, \ell) \geq C(\ell)\},$$

where LP can be taken as an acronym of losses in pooling equilibria. The set LP denotes all losses each of which has a conditional expected repair value $AL(\underline{\ell}, \ell)$ larger than cost $C(\ell)$. Now, both $AL(\underline{\ell}, \ell)$ and $C(\ell)$ are increasing, so may intersect multiple times as ℓ varies between $\underline{\ell}$ and $\bar{\ell}$. Hence, the set LP may consist of disjoint loss intervals. However, $C(\widehat{\ell}) = \widehat{\ell} > AL(\underline{\ell}, \widehat{\ell})$, so for losses above $\widehat{\ell}$, $C(\ell)$ must over take $AL(\underline{\ell}, \ell)$. Finally, we define a threshold ℓ^* by

$$\ell^* \equiv \max\{\ell : AL(\underline{\ell}, \ell) \geq C(\ell)\}.$$

Obviously $AL(\underline{\ell}, \ell) < C(\ell)$ for $\ell > \ell^*$; moreover, $\ell^* < \widehat{\ell}$ because $C(\ell^*) = AL(\underline{\ell}, \ell^*) < \ell^*$ and $C(\ell) < \ell$ if and only if $\ell < \widehat{\ell}$. Figure 1 graphs an example of $AL(\underline{\ell}, \ell)$ and $C(\ell)$ in which they cross three times. There the set LP consists of those losses between $\underline{\ell}$ and ℓ_1 and between ℓ_2 and ℓ^* , the maximum loss in LP .

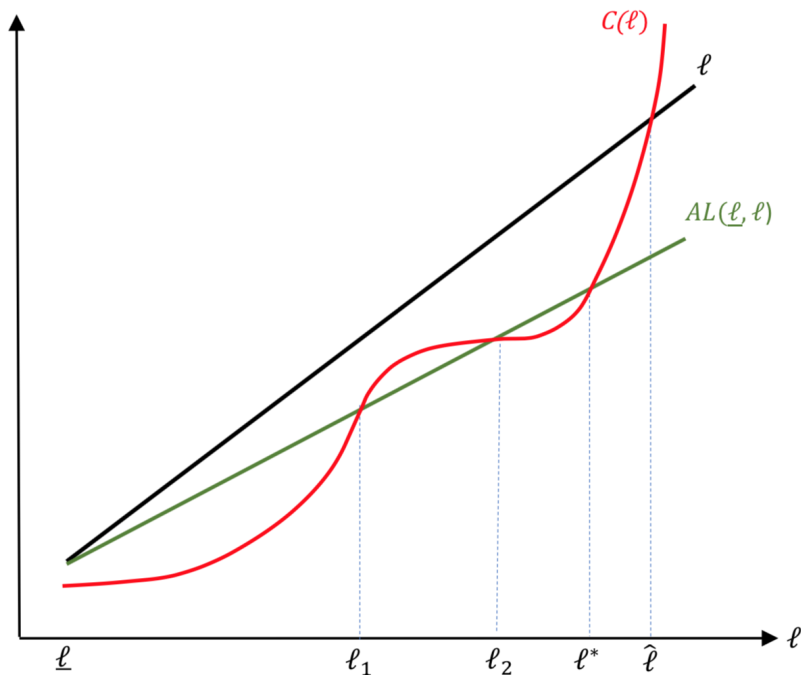


Figure 1: Conditional expected loss value $AL(\underline{\ell}, \ell)$ and cost $C(\ell)$

In a pooling equilibrium, the expert must set the same price for all repairs that happen with a positive probability. Given that treatment cost $C(\ell) > \ell$ for $\ell > \widehat{\ell}$, there does not exist a pooling equilibrium in which all losses are repaired. (In such a (candidate) pooling equilibrium, the expert's price is at most μ , the consumer's expected loss. Because $\mu < \bar{\ell} < C(\bar{\ell})$, the expert will refuse to repair those losses that cost more than μ .) In the next proposition, we characterize an equilibrium in which the expert offers a pooling price for losses lower than the threshold ℓ^* . Many prices in the range $[\underline{\ell}, \widehat{\ell}]$ will be off path, but we only need to impose a minimum restriction on the consumer's belief upon a price deviation.

Proposition 2 *There is a pooling equilibrium indexed by ℓ^* (where $AL(\underline{\ell}, \ell^*) = C(\ell^*)$). The expert's strategy*

is

$$P(\ell) = \begin{cases} AL(\underline{\ell}, \ell^*) & \text{if } \underline{\ell} \leq \ell \leq \ell^* \\ +\infty & \ell^* < \ell \leq \bar{\ell} \end{cases}.$$

The consumer's strategy is

$$\alpha(p) = \begin{cases} 1 & \text{if } p \leq AL(\underline{\ell}, \ell^*) \\ 0 & \text{if } p > AL(\underline{\ell}, \ell^*) \end{cases}.$$

When the consumer receives a price offer different from $AL(\underline{\ell}, \ell^*)$, she believes that her loss is drawn from $[\underline{\ell}, \ell^*]$ according to $F(\ell)$




Figure 2: A pooling equilibrium

There is a single equilibrium price $AL(\underline{\ell}, \ell^*)$ offered by the expert that results in equilibrium transactions, so this is a pooling equilibrium. Essentially, the consumer assumes (correctly) that the expert offers $AL(\underline{\ell}, \ell^*)$ for $\ell \in [\underline{\ell}, \ell^*]$ and that no transaction can occur at higher losses. The consumer accepts all prices below $AL(\underline{\ell}, \ell^*)$ and rejects higher prices. Against the consumer's best response, the expert's best response, indeed, is to offer price $AL(\underline{\ell}, \ell^*)$ for $\ell \in [\underline{\ell}, \ell^*]$. At $\ell = \ell^*$, the expert makes zero profit because $AL(\underline{\ell}, \ell^*) = C(\ell^*)$. Part iii) of Proposition 1 says that the expert cannot make any profit for higher losses. Therefore, no transaction will happen for losses greater than ℓ^* . The pooling equilibrium is illustrated in Figure 2. There

are other pooling equilibria indexed by losses different from ℓ^* (such as ℓ_1 and ℓ_2), but they fail a natural restriction on the consumer's off-equilibrium beliefs. In Section 5, we state the refinement (Definition 3) and the result (Proposition 5).

3.3 A separating equilibrium

To begin constructing a separating equilibrium, we present the following key lemma.

Lemma 1 *Suppose that the equilibrium price function P is strictly increasing over an interval of losses on the equilibrium path, say (ℓ_1, ℓ_2) , then it must be $P(\ell) = \ell$ if the consumer accepts $P(\ell)$ with some positive probability.*

Lemma 1 says that in a separating equilibrium, the price must equal the consumer's loss, making her just indifferent between accepting and rejecting the offer. Clearly, if the consumer can infer her loss from the price offer, she will pay at most her loss. If the expert offers to repair a problem at an equilibrium price strictly less than the loss, the consumer will accept it with probability one. However, this will create an incentive problem. The expert would offer this price for all problems with lower losses because he could surely earn this higher price. The argument implies that the consumer must reject the expert's price with a positive probability, so $P(\ell) = \ell$. Using this property, can we find the acceptance probability function to support a separating equilibrium?

Proposition 3 *There is a unique separating equilibrium (which involves the threshold $\widehat{\ell}$, defined by $C(\widehat{\ell}) = \widehat{\ell}$). The expert's strategy is*

$$P(\ell) = \begin{cases} \ell & \text{if } \underline{\ell} \leq \ell \leq \widehat{\ell} \\ +\infty & \text{if } \widehat{\ell} < \ell \leq \bar{\ell} \end{cases},$$

and the consumer's strategy is

$$\alpha(\ell) = \begin{cases} \exp \left\{ -\int_{\underline{\ell}}^{\ell} \frac{dx}{x - C(x)} \right\} & \text{if } \underline{\ell} \leq \ell \leq \widehat{\ell} \\ 0 & \text{if } \widehat{\ell} < \ell \leq \bar{\ell} \end{cases}.$$

Uniqueness follows from Lemma 1 and the consumer's acceptance function. Given the expert's strategy, the consumer is indifferent between accepting and rejecting the offer. Hence, the acceptance strategy $\alpha(\ell)$ is

a best response. Given the acceptance strategy $\alpha(\ell)$, the expert strictly prefers to set $P(\ell) = \ell$ at loss ℓ . The rest of the proof of the Proposition consists of deriving the necessary and sufficient condition on the consumer acceptance probability function to support the expert's separating price function. The construction of the equilibrium acceptance probability function is by means of the envelope condition to obtain the equilibrium profit's derivative, a familiar method in contract design. Because choosing $P(\ell) = \ell$ is optimal, the derivative of equilibrium profit $\Pi(\ell) = (\ell - C(\ell))\alpha(\ell)$ only depends on how it changes with respect to cost:

$$\Pi'(\ell) = -C'(\ell)\alpha(\ell) < 0. \tag{3}$$

The expert's equilibrium profit must be strictly decreasing in the separating equilibrium, consistent with Proposition 1. To prevent the expert from gaining by offering a higher price when the loss is small, the consumer must decrease acceptance probabilities as losses (and prices) increase.

The equilibrium in Proposition 3 does not have any unreached information set in the relevant price range $[\underline{\ell}, \widehat{\ell}]$. The separating equilibrium, therefore, is robust against any refinement that restricts beliefs off the equilibrium path. The equilibrium does require the consumer to fine tune the acceptance probability according to the solution of a differential equation. This is demanding, perhaps unrealistic, and suffers from the usual criticism of a player randomizing between multiple optimal choices to support a rival's optimal action. Nevertheless, Proposition 3 does confirm some insights from the binary-loss models in the literature: i) credence goods do not necessarily yield equilibrium lies (Fong (2005)), and ii) the consumer rejecting offers does discipline the expert's information advantage ((Wolinsky (1993), Fong (2005), Dulleck and Kerschbamer (2006))).

The contrasts between the separating equilibrium and the pooling equilibrium are quite striking. First, the consumer infers her loss from price in the separating equilibrium but only knows that her loss is below or above a threshold in the pooling equilibrium. Second, the consumer plays a mixed strategy in the separating equilibrium but plays a pure strategy in the pooling equilibrium. Third, in the separating equilibrium, the expert repairs each loss with a positive probability (declining as loss increases), and makes a strictly positive profit (also declining as loss increases) for losses efficient to repair; in the pooling equilibrium, the expert only repairs losses between $\underline{\ell}$ and ℓ^* , and profit goes to 0 as the loss value goes to $\ell^* < \bar{\ell}$. Fourth, the consumer has a

zero ex-post payoff in the separating equilibrium; by contrast, in the pooling equilibrium, the consumer has a zero ex-ante payoff, but ex post, if the loss eventually became known, her payoff might turn out to be positive or negative. These results also compare squarely with those in the literature. A separating equilibrium is synonymous with truth-telling, or lack of lying. By contrast, a pooling equilibrium is synonymous with lying. Both truth-telling and lying are equilibrium behavior in the model with a continuum loss. We believe that the pooling equilibrium is more plausible than the separating equilibrium because it is a simple, pure strategy equilibrium. Furthermore, as we show in Subsection 6.3, a separating does not exist when cost is higher than loss at low loss values, so may not be robust. The pooling equilibrium does not suffer from this problem.

We now compare the expert's equilibrium profits between pooling and separating equilibria. Because consumer surplus is zero in both equilibria, comparing profits across the two equilibria is the same as comparing social surpluses. Depending on the loss distribution and cost functions, pooling equilibrium profit may be higher or lower than separating equilibrium profit. In general, for a given cost function C and a given distribution F , we have

Corollary 1 *The pooling equilibrium in Proposition 2 yields more social surplus (and profit) than the separating equilibrium in Proposition 3 if and only if*

$$\int_{\underline{\ell}}^{\ell^*} F(\ell)C'(\ell)d\ell \geq \int_{\underline{\ell}}^{\widehat{\ell}} F(\ell)C'(\ell)\alpha(\ell)d\ell, \quad (4)$$

where $\alpha(\ell) = \exp \left\{ -\int_{\underline{\ell}}^{\ell} \frac{dx}{x - C(x)} \right\}$.

Both the pooling and the separating equilibria are inefficient. In the pooling equilibrium, losses below ℓ^* are efficiently repaired but any losses above ℓ^* are left untreated. In comparison, in the separating equilibrium, losses below $\widehat{\ell}$ are recovered with a positive probability below one. Corollary 1 says that the pooling equilibrium is more efficient and profitable than the separating equilibrium if the two cutoffs ℓ^* and $\widehat{\ell}$ are sufficiently close. This can happen if the cost function is sufficiently convex. The following Figures 3 illustrates these possibilities.

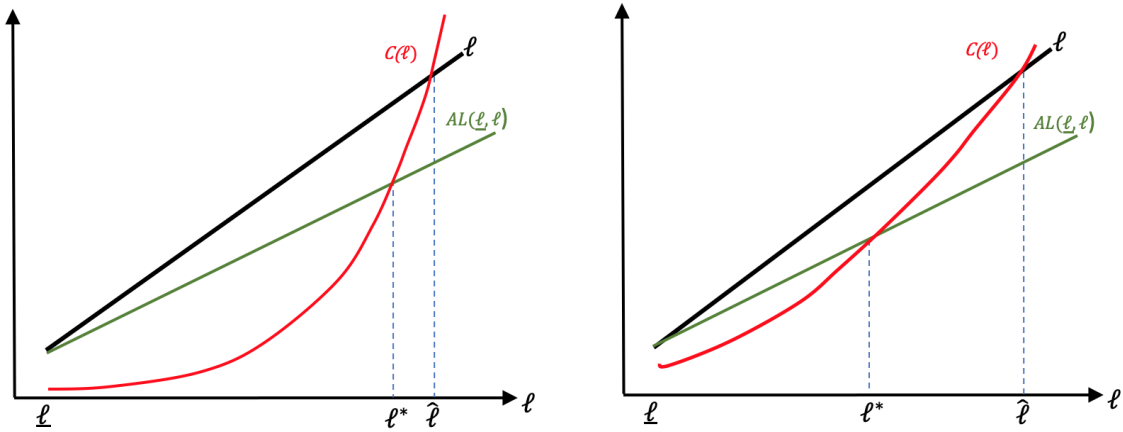


Figure 3: Two examples comparing pooling and separating equilibria

In these panels the consumer's loss distribution F , and hence the function $AL(\underline{\ell}, \ell)$, stay the same, but the cost function C differs. In the left panel, the value of ℓ^* is high, near $\hat{\ell}$, so the pooling equilibrium profit is quite close to the first best, higher than the separating equilibrium profit. In the right panel, the value of ℓ^* is low, near $\underline{\ell}$, so the profit in the pooling equilibrium is quite low, and the separating equilibrium must yield higher profits.

We work out two examples that verify Corollary 1:

Example 1: Loss ℓ is uniformly distributed on $[1, 3]$ and $C(\ell) = \frac{1}{2}\ell^2$. It follows that $\hat{\ell} = 2$, $\ell^* = \frac{1 + \sqrt{5}}{2}$, $\alpha(\ell) = \left(\frac{2}{\ell} - 1\right)$ for $\ell \leq \hat{\ell}$. From straightforward computation, the left-hand-side of (4) is 0.1348 and the right-hand-side is 0.0833.

Example 2: Loss ℓ is uniformly distributed on $[\underline{\ell}, \bar{\ell}]$, $\underline{\ell} > 0$, and $C(\ell) = b\ell - k$, with $b > 1$ and $k \in ((b-1)\underline{\ell}, (b-1)\bar{\ell})$. In this case, $\hat{\ell} = \frac{k}{b-1}$, $\ell^* = \frac{\underline{\ell} + 2k}{2b-1}$ and $\alpha(\ell) = \left[\frac{k - (b-1)\ell}{k - (b-1)\underline{\ell}}\right]^{-\frac{1}{b-1}}$. The difference between the left-hand side and the right-hand side of (4) is $\frac{[k - (b-1)\underline{\ell}]^2}{(\bar{\ell} - \underline{\ell})(2b-1)^2} > 0$. Note that $k - (b-1)\underline{\ell}$ is the surplus from repairing the lowest loss $\underline{\ell}$. Hence, the profit difference between pooling and separating equilibria increases in this surplus.

Corollary 1 and the above two examples show results that would be unavailable in a binary model. Fong (2005) shows that the expert's profit is maximized by the separating equilibrium. If there are only two losses,

say minor $\underline{\ell}$ and major $\bar{\ell}$, and the average loss is smaller than the major repair cost $C(\bar{\ell})$, the only pooling equilibrium has the expert offering price $\underline{\ell}$ if and only if the loss is minor. The binary setting has lumpy equilibrium inefficiency characteristics. However, for continuous losses, all losses up to ℓ^* are repaired, and this often yields more profits than the separating equilibrium.

Propositions 1 to 2 together show that *all* equilibria are inefficient. Efficiency requires all losses $\ell \in [\underline{\ell}, \widehat{\ell}]$ be repaired with probability one. Proposition 1 says that the equilibrium acceptance rate is weakly decreasing in loss. Hence, the only possible candidate for an efficient equilibrium is a pooling equilibrium. Nevertheless, because the highest loss threshold for repair in any pooling equilibrium ℓ^* , is strictly less than $\widehat{\ell}$, there does not exist such an efficient pooling equilibrium. We now turn to an alternative information structure for a fully efficient equilibrium.

4 Expert information acquisition

We have taken as given that the expert perfectly observes the consumer's loss before he makes a price offer. In this section, we study the expert's information acquisition decision. We show that the expert is better off being imperfectly informed about the consumer's loss even when information acquisition is costless. We modify the game by adding a "Stage 0," in which the expert picks an information structure or a diagnostic test. This is defined by a set of signals $S \subset \mathfrak{R}$ and a joint distribution, $J : S \times [\underline{\ell}, \bar{\ell}] \rightarrow [0, 1]$, so $J(s, \ell)$ measures the probability that the signal is less than s and the loss is less than ℓ . To be a valid test, we require: i) $\int_S dJ(s, \ell) = J(\sup(S), \ell) = F(\ell)$ and ii) $J(s, \bar{\ell}) \equiv G(s)$ yields a probability distribution over S . The requirement i) implies that $\int_{\underline{\ell}}^{\bar{\ell}} \int_S \ell dJ(s, \ell) = \int_{\underline{\ell}}^{\bar{\ell}} \ell dF(\ell) = \mu$. We do rule out the uninteresting case where J is the product of its two marginal distributions, the case of uninformative signals. The model in Section 2 has a perfectly informative test.¹² We let the information structure or diagnostic test be public information once it has been chosen. A cost can be associated with an information structure, and including this is straightforward. For now we just assume that an information structure is costless.

¹²That is, there is an increasing, bijective function $\phi : [\underline{\ell}, \bar{\ell}] \rightarrow [\underline{s}, \bar{s}]$ such that $F(\ell) = G(s)$ where $s = \phi(\ell)$ and $G(s) \equiv \int_{\underline{\ell}}^{\bar{\ell}} dJ(s, \ell)$. The simplest case is where ϕ is the identity map.

We augment the game in Section 2 as follows. In Stage 0, the expert chooses an information structure or a diagnostic test. In Stage 1, the consumer visits the expert who then privately receives a signal from the test. The expert makes a price offer based on the private signal, or refuses to provide a repair. In Stage 2, upon receiving a price offer, the consumer chooses between accepting or rejecting the offer.

We assume that the consumer's actual loss ℓ will be revealed to the expert during the repair even if the test is imperfect. When the expert provides treatment for loss ℓ , he still has to incur the cost $C(\ell)$ eventually because he is liable for repair outcomes. More information arrives during the repair process (but after the price has been agreed upon); for example, car mechanics usually learn more about the cost after they disassemble parts during the repair.

Recall that the first-best allocation is one with repair for all losses up to $\widehat{\ell}$ (at which point, $\widehat{\ell} = C(\widehat{\ell})$). If the diagnostic test is perfect, the analysis is the same as in Section 3. Neither the pooling equilibrium in Proposition 2 nor the separating equilibrium in Proposition 3 can yield the first best. The following diagnostic test, however, can implement the first best. Let the test have binary signals: $S = \{s_L, s_H\}$, where $s_L < s_H$. The signal takes the value s_L when loss is below $\widehat{\ell}$, and s_H otherwise. Thus, the signal is s_L if and only if the loss is above cost; the signal s_L indicates that repair is efficient. We write down its definition as follows. It is somewhat more familiar to use conditional densities. So let $f = F'$, and $j(s_L, \ell) = \frac{\partial J(s_L, \ell)}{\partial \ell}$ and $j(s_H, \ell) = \frac{\partial J(s_H, \ell)}{\partial \ell}$. Then in terms of joint densities, we have

joint density	$j(s_L, \ell)$	$j(s_H, \ell)$
$\underline{\ell} \leq \ell \leq \widehat{\ell}$	$f(\ell)$	0
$\widehat{\ell} \leq \ell \leq \bar{\ell}$	0	$f(\ell)$
signal probabilities	$F(\widehat{\ell}) = \Pr(S = s_L)$	$1 - F(\widehat{\ell}) = \Pr(S = s_H)$

(5)

and the corresponding conditional densities, $k(\ell | S)$, are

$$k(\ell | S = s_L) = \begin{cases} \frac{f(\ell)}{F(\widehat{\ell})} & \text{for } \ell \leq \widehat{\ell} \\ 0 & \text{for } \ell > \widehat{\ell} \end{cases} \quad \text{and} \quad k(\ell | S = s_H) = \begin{cases} 0 & \text{for } \ell \leq \widehat{\ell} \\ \frac{f(\ell)}{1 - F(\widehat{\ell})} & \text{for } \ell > \widehat{\ell} \end{cases} \quad (6)$$

Upon observing signal s_L , the expert's expected repair cost is $\frac{\int_{\underline{\ell}}^{\widehat{\ell}} C(\ell) dF(\ell)}{F(\widehat{\ell})} \equiv AC(\underline{\ell}, \widehat{\ell})$, the average

cost of repairing losses below $\widehat{\ell}$. By the definition of $\widehat{\ell}$, $\ell > C(\ell)$ for all $\ell < \widehat{\ell}$. Hence, $AL(\underline{\ell}, \widehat{\ell}) > AC(\underline{\ell}, \widehat{\ell})$. Given that the information structure is public information, if the consumer knew that the expert received the signal s_L , the maximum acceptable price would be $AL(\underline{\ell}, \widehat{\ell})$.

Proposition 4 *The expert can implement the first best and extract the entire surplus from trade by the information structure in (5) or (6). The expert's strategy is*

$$P(s) = \begin{cases} AL(\underline{\ell}, \widehat{\ell}) & \text{if } s = s_L \\ +\infty & \text{if } s = s_H \end{cases},$$

and the consumer's strategy is

$$\alpha(p) = \begin{cases} 1 & \text{if } p \leq AL(\underline{\ell}, \widehat{\ell}) \\ 0 & \text{if } p > AL(\underline{\ell}, \widehat{\ell}) \end{cases}.$$

The consumer always believes that the expert has received the signal s_L upon being made a price offer for repair. In words, the expert offers repair at price $AL(\underline{\ell}, \widehat{\ell})$ if and only if the signal is s_L ; the consumer always accepts an offer if it does not exceed $AL(\underline{\ell}, \widehat{\ell})$.

The proof of this Proposition is obvious, and omitted. The expert's equilibrium expected profit is $[AL(\underline{\ell}, \widehat{\ell}) - AC(\underline{\ell}, \widehat{\ell})] F(\widehat{\ell}) = \int_{\underline{\ell}}^{\widehat{\ell}} [\ell - C(\ell)] dF(\ell)$, which is the first-best surplus. However, ex post, the repair cost for some losses may be above the equilibrium price. Indeed, $C(\widehat{\ell}) = \widehat{\ell} > AL(\underline{\ell}, \widehat{\ell})$, so for losses close to $\widehat{\ell}$, the expert will make a loss ex post, but of course the expert will make a profit for low losses. A similar observation holds for the consumer: at low losses, the consumer will over pay, but at high losses, the consumer will earn a surplus. On average, the consumer earns a zero expected payoff, all ex-ante surplus being expropriated by the expert.

Clearly, the information structure in (6) acts like it could only implement a pooling equilibrium, but this equilibrium is first best, so more efficient than the one in Proposition 2. In Proposition 4, the offer is based on the average cost; in Proposition 2, it is based on the actual or marginal cost. The strategy profile in Proposition 4 would not constitute an equilibrium if the expert had perfect information; the expert would refuse to repair losses with (marginal) costs above the pooling price $AL(\underline{\ell}, \widehat{\ell})$. Anticipating the expert's cream skimming, the consumer would not accept that price offer. By contrast, under the information structure in (6), the consumer and the expert are on a symmetric footing. All the expert manages to learn is whether

the average cost is below the average loss. Using a binary signal, the expert commits to using average-cost information to make price offers, and achieves the first best.

Proposition 4 is not meant to report an optimism that the simple diagnostic test entirely solves credence-good problems. The analytical result notwithstanding, the extensive form does require the expert’s choice of a diagnostic test be common knowledge. It may be possible for an expert to install a piece of equipment or design a test to screen whether the consumer’s problem severity is above or below a certain level. For example, a car mechanic can use a basic diagnostic scanner which only indicates whether the problem is caused by a failure in engine before quoting price. A doctor can order a laboratory test which screens whether the patient has a cancer. However, in practice, consumers may not understand how a diagnostic test actual works, or how a signal is to be interpreted. Proposition 4 succinctly shows that more information may impede economic transaction because the extra information will become private; this is in addition to more information may risk a cognitive overload

5 Other equilibria

5.1 Multiple pooling equilibria and a belief refinement

The equilibrium in Proposition 2 may not be the only pooling equilibrium. We can apply the same construction and argument for any $\tilde{\ell} \in LP$ where $AL(\underline{\ell}, \tilde{\ell}) = C(\tilde{\ell})$ (such as points ℓ_1 and ℓ_2 in Figure 2). We simply replace ℓ^* in Proposition 2 (and its proof) by $\tilde{\ell}$ verbatim. We record this as a corollary (proof omitted).

Corollary 2 *There may be other pooling equilibria indexed by $\tilde{\ell} \in LP$ at which $AL(\underline{\ell}, \tilde{\ell}) = C(\tilde{\ell})$, and, in each of these equilibria, the strategies and beliefs are the same as those in Proposition 2 with ℓ^* there replaced by $\tilde{\ell}$.*

A pooling equilibrium indexed by $\tilde{\ell} < \ell^*$ is supported by the off-equilibrium belief that the consumer’s loss belongs to the interval $[\underline{\ell}, \tilde{\ell})$, irrespective of the price. This off-equilibrium passive belief is not very sophisticated or convincing. We now present a belief refinement, and then shows that it eliminates all pooling equilibria except for the one indexed by ℓ^* .

Definition 3 (Nonnegative Profit Principle) *Upon receiving an off-equilibrium price offer p , the consumer believes that the loss is drawn from $\{\ell : p \geq C(\ell)\}$ according to F . A pooling equilibrium is said to satisfy the Nonnegative Profit Principle if the consumer's strategy specifies a best response against off-equilibrium price p based on the belief that loss has been drawn from $\{\ell : p \geq C(\ell)\}$ according to F .*

According to the Nonnegative Profit Principle, the consumer believes that the expert's price at least covers cost. Offering a price that would not cover cost is a dominated strategy, and the consumer would not believe that the expert would do that. Apart from this requirement, the consumer cannot pin down other motives, so we assume that she believes losses are drawn from a truncated distribution of losses.

Proposition 5 *The pooling equilibrium in Proposition 2 satisfies the Nonnegative Profit Principle. A pooling equilibrium indexed by $\tilde{\ell} < \ell^*$ in Corollary 2 does not satisfy the Nonnegative Profit Principle.*

Figure 4 illustrates how the Nonnegative Profit Principle eliminates all pooling equilibria but the one indexed by ℓ^* . Consider the pooling equilibrium indexed by ℓ_1 in the figure. The equilibrium price is $AL(\underline{\ell}, \ell_1)$. Suppose that the expert offers the off-equilibrium price p' . Based on the Nonnegative Profit Principle, the consumer believes that her loss belongs to $[\underline{\ell}, \ell']$, where $p' = C(\ell')$, so her conditional expected loss $AL(\underline{\ell}, \ell')$ is higher than the price p' . Hence, the consumer will accept p' with probability one. Given the consumer's response to the off-equilibrium price, the expert will deviate to offering p' for losses lower than ℓ' . The Nonnegative Profit Principle effectively allows an expert to convince the consumer to accept a higher price. Proposition 5 says that among all pooling equilibria, the one in Proposition 2 is most compelling.

5.2 Hybrid equilibria

Our analysis has focussed on the separating and the pooling equilibria. There are many other equilibria, indeed, a continuum of them. First we give a definition for equilibria that are neither pooling nor separating.

Definition 4 (Hybrid equilibrium) *An equilibrium (P, α) is said to be a hybrid equilibrium if there are disjoint intervals (ℓ_1, ℓ_2) and (ℓ_3, ℓ_4) , both in $[\underline{\ell}, \bar{\ell}]$, such that for any ℓ and ℓ' in (ℓ_1, ℓ_2) for which $\alpha(P(\ell)) > 0$ and $\alpha(P(\ell')) > 0$, $P(\ell) \neq P(\ell')$, and for any ℓ and ℓ' in (ℓ_3, ℓ_4) for which $\alpha(P(\ell)) > 0$ and $\alpha(P(\ell')) > 0$,*




Figure 4: The Nonnegative Profit Principle

$P(\ell) = P(\ell')$. (There may well be many intervals of the sort like (ℓ_1, ℓ_2) and (ℓ_3, ℓ_4) .)

In a hybrid equilibrium, over some loss interval, prices that may be accepted vary with losses, but over some other loss interval, a price that may be accepted does not vary with losses. There may also be many of these separating and pooling intervals. Now we construct a continuum of hybrid equilibria. First, we present a corollary of Proposition 3, which characterizes a separating equilibrium. However, loss values may be revealed over a subset of $[\underline{\ell}, \widehat{\ell}]$, not necessarily for all losses. In fact, the proof of Proposition 3 shows that explicitly. The following presents an acceptance function that supports separation over a loss interval (proof omitted).

Corollary 3 *In an equilibrium, the expert offers $P(\ell) = \ell$ for each $\ell \in [\ell_1, \ell_2] \subset [\underline{\ell}, \widehat{\ell}]$ if and only if the*

consumer accepts offer ℓ with probability $\alpha(\ell) = \alpha(\ell_1) \exp \left\{ - \int_{\ell_1}^{\ell} \frac{dx}{x - C(x)} \right\}$, where $\alpha(\ell_1)$ is the acceptance probability of offer $P(\ell_1) = \ell_1$.

The next result presents a family of hybrid equilibria. In each, the expert offers a pooling price for low losses, and then switches to a separating price for high losses; in between there is jump in the offer. The consumer accepts the pooling price with probability one, and then follows the acceptance probabilities prescribed by Corollary 3.

Proposition 6 For each $\tilde{\ell} \in \{\ell : AL(\underline{\ell}, \ell) > C(\ell)\}$ (which is the interior of LP), there is a hybrid equilibrium. The expert's strategy is

$$P(\ell) = \begin{cases} AL(\underline{\ell}, \tilde{\ell}) & \text{if } \ell \in [\underline{\ell}, \tilde{\ell}] \\ \ell & \text{if } \ell \in (\tilde{\ell}, \hat{\ell}] \\ +\infty & \text{if } \ell \in (\hat{\ell}, \bar{\ell}] \end{cases} .$$

The consumer's strategy is

$$\alpha(p) = \begin{cases} 1 & \text{if } p \in [\underline{\ell}, AL(\underline{\ell}, \tilde{\ell})] \\ 0 & \text{if } p \in (AL(\underline{\ell}, \tilde{\ell}), \tilde{\ell}) \\ \alpha(\tilde{\ell}) \exp \left\{ - \int_{\tilde{\ell}}^p \frac{dx}{x - C(x)} \right\}, \text{ with } \alpha(\tilde{\ell}) = \frac{AL(\underline{\ell}, \tilde{\ell}) - C(\tilde{\ell})}{\tilde{\ell} - C(\tilde{\ell})} & \text{if } p \in [\tilde{\ell}, \hat{\ell}] \\ 0 & \text{if } p \in (\hat{\ell}, \bar{\ell}] \end{cases} .$$

If the consumer is offered a price at or below $\tilde{\ell}$, she believes that the loss is drawn from $[\underline{\ell}, \tilde{\ell}]$ according F ; if she is offered any price at or above $\hat{\ell}$, she believes that the loss is $\hat{\ell}$.

This proposition presents a continuum of “hybrid” equilibria; they amalgamate equilibria in Propositions 3 and 2. Each hybrid equilibrium has a pooling interval, for those losses between $\underline{\ell}$ and $\tilde{\ell}$, with the pooling price $AL(\underline{\ell}, \tilde{\ell})$ that is always accepted by the consumer, and then a separating interval, for losses above $\tilde{\ell}$. The expert makes a strictly positive profit at loss $\tilde{\ell}$ because $AL(\underline{\ell}, \tilde{\ell}) > C(\tilde{\ell})$ by the definition of $\tilde{\ell}$. Therefore, there is scope for more transactions, which happen to be separating prices. At the cutoff loss $\tilde{\ell}$, the expert is indifferent between recommending the pooling price $AL(\underline{\ell}, \tilde{\ell})$ and the higher separating price $\tilde{\ell}$, by the choice of the acceptance probability $\alpha(\tilde{\ell})$. But there is a discontinuous jump from the pooling price

$AL(\underline{\ell}, \tilde{\ell})$ to the lowest separating price $\tilde{\ell}$. As $\tilde{\ell}$ goes to $\underline{\ell}$, the hybrid equilibrium converges to the separating equilibrium in Proposition 3, and as $\tilde{\ell}$ goes to ℓ^* , the hybrid equilibrium converges to the pooling equilibrium in Proposition 2.

Figure 5 illustrates Proposition 6. There, the expert recommends the pooling price, $AL(\underline{\ell}, \tilde{\ell})$, for losses up to $\tilde{\ell}$. At $\tilde{\ell}$, the price offers jump up to the 45-degree line. Off-equilibrium prices are those below $AL(\underline{\ell}, \tilde{\ell})$, between $AL(\underline{\ell}, \tilde{\ell})$ and $\tilde{\ell}$, and above $\hat{\ell}$. At off-equilibrium prices between $AL(\underline{\ell}, \tilde{\ell})$ and $\tilde{\ell}$, the consumer continues to believe that the loss is drawn from the pooling set of losses $(\underline{\ell}, \tilde{\ell})$, so rejects them.




Figure 5: A hybrid equilibrium

As we have shown in Proposition 5, the “passive” off-equilibrium belief at prices between $AL(\underline{\ell}, \tilde{\ell})$ and $\tilde{\ell}$ is implausible. The consumer should not ignore the expert’s information advantage and pricing strategy. In fact, we can reject these hybrid equilibria.¹³

¹³We are aware that hybrid equilibria might resemble two-part tariffs: a base fee for diagnosis and minor repairs, and then a higher fee for more complicated repairs. We are agnostic about whether such two-part tariffs can be

Proposition 7 *Hybrid equilibria in Proposition 6 do not satisfy the Nonnegative Profit Principle.*

The proof of this Proposition follows closely that of Proposition 5, so we will only give an outline here. Consider a hybrid equilibrium indexed by $\tilde{\ell}$ where $AL(\underline{\ell}, \tilde{\ell}) > C(\ell)$, as in Figure 5. The pooling price of the hybrid equilibrium is $AL(\underline{\ell}, \tilde{\ell})$; the profit from repairing a problem with loss $\tilde{\ell}$ is $AL(\underline{\ell}, \tilde{\ell}) - C(\tilde{\ell}) > 0$. Suppose that the expert offer price is $AL(\underline{\ell}, \tilde{\ell}) + \varepsilon$. For a sufficiently small $\varepsilon > 0$, the set $\{\ell : AL(\underline{\ell}, \tilde{\ell}) + \varepsilon \geq C(\ell)\}$ has an upper bound $\tilde{\ell} + \delta$, where $AL(\underline{\ell}, \tilde{\ell}) + \varepsilon = C(\tilde{\ell} + \delta)$; see Figure 6. At $\tilde{\ell} + \delta$, we have $AL(\underline{\ell}, \tilde{\ell} + \delta) > AL(\underline{\ell}, \tilde{\ell}) + \varepsilon$, which says that the consumer's willingness to pay is higher than the price if she believes that the loss is drawn from $[\underline{\ell}, \tilde{\ell} + \delta]$ according to F . This belief is what the Nonnegative Profit Principle prescribes, so the consumer should accept the price $AL(\underline{\ell}, \tilde{\ell}) + \varepsilon$. The same argument can be applied to rule out any other hybrid equilibrium.¹⁴

More complicated hybrid equilibria can be constructed in a similar fashion. An equilibrium can start with separating prices, turns to pooling prices, and then reverts to separating prices—as long as monotonicity in Proposition 1 and no-deviation conditions are satisfied.¹⁵ We can certainly write down conditions on the primitives for such constructions, but they will be little more than applying Proposition 1 and modifying strategies in Proposition 6.

We use the model primitives in Example 1 above to construct a complicated hybrid equilibrium. Again, suppose that ℓ is uniformly distribution on $[1, 3]$ and $C(\ell) = \frac{1}{2}\ell^2$. Then $AL(\ell_1, \ell_2) = \frac{1}{2}[\ell_1 + \ell_2]$. We construct an equilibrium in which the expert recommends the separating price in two disjoint intervals of losses and

convincingly rationalized by our basic model. A full analysis is beyond the scope here.

¹⁴A formal argument is the following. For a given $\tilde{\ell} \in \{\ell | AL(\underline{\ell}, \ell) > C(\ell)\}$, define $\ell_m \equiv \min\{\ell | AL(\underline{\ell}, \ell) = C(\ell) \text{ and } \ell > \tilde{\ell}\}$. Because $\tilde{\ell} < \ell_m$, $AL(\underline{\ell}, \tilde{\ell}) < AL(\underline{\ell}, \ell_m)$. There exists a small ε such that $AL(\underline{\ell}, \tilde{\ell}) + \varepsilon < AL(\underline{\ell}, \ell_m)$. Following the definition of δ and ℓ_m , $C(\tilde{\ell} + \delta) < C(\ell_m)$ and hence $\tilde{\ell} < \tilde{\ell} + \delta < \ell_m$. Again, by the definition of ℓ_m and $\tilde{\ell}$, $AL(\underline{\ell}, \tilde{\ell} + \delta) > C(\tilde{\ell} + \delta) = AL(\underline{\ell}, \tilde{\ell}) + \varepsilon$. This argument shows that upon recommended the off-the-equilibrium-path price $AL(\underline{\ell}, \tilde{\ell}) + \varepsilon$, the consumer's expected loss is $AL(\underline{\ell}, \tilde{\ell} + \delta) > AL(\underline{\ell}, \tilde{\ell}) + \varepsilon$. Hence, the consumer will accept the price.

¹⁵We can also construct equilibria that are made up of only many pooling regions.




Figure 6: Rejecting hybrid equilibria

the pooling price in between. Consider the following expert strategy:

$$P(\ell) = \begin{cases} \ell & \text{if } \ell \in [1, \ell_1] \\ \frac{1}{2}[\ell_1 + \ell_2] & \text{if } \ell \in (\ell_1, \ell_2] \\ \ell & \text{if } \ell \in (\ell_2, \widehat{\ell}] \\ +\infty & \text{if } \ell \in (\widehat{\ell}, 3] \end{cases} .$$

The expert recommends the separating price for $\ell \in [1, \ell_1] \cup (\ell_2, \widehat{\ell}]$ and the pooling price $\frac{1}{2}[\ell_1 + \ell_2]$ for $\ell \in (\ell_1, \ell_2]$. Let $\alpha(\ell) = \alpha(\ell_0) \exp \left\{ - \int_{\ell_1}^{\ell} \frac{dx}{x - C(x)} \right\} = \alpha(\ell_0) \left[\frac{2}{\ell} - 1 \right]$, where ℓ_0 is the beginning loss value at which separation occurs. The consumer's acceptance probability is $\alpha(\ell) = \left[\frac{2}{\ell} - 1 \right]$ for $\ell \in [1, \ell_1]$. It follows that the expert's profit at the threshold ℓ_1 is $\pi(\ell_1) = \left[\frac{2}{\ell_1} - 1 \right] \left[\ell_1 - \frac{1}{2}\ell_1^2 \right]$. For ℓ_1 , the profit at the pooling price is $A \times \left[\frac{1}{2}[\ell_1 + \ell_2] - \frac{1}{2}\ell_1^2 \right]$, where A is the probability that the consumer accepts the pooling price. Probability A is chosen to make the expert just indifferent between recommending the separating

price ℓ_1 and the pooling price at ℓ_1 . So,

$$\left[\frac{2}{\ell_1} - 1 \right] \left[\ell_1 - \frac{1}{2}\ell_1^2 \right] = A \times \left[\frac{1}{2}[\ell_1 + \ell_2] - \frac{1}{2}\ell_1^2 \right]. \quad (7)$$

The expert's profit from pooling at threshold ℓ_2 is $A \times [\frac{1}{2}[\ell_1 + \ell_2] - \frac{1}{2}\ell_2^2]$, whereas profit from separating at price ℓ_2 is $\alpha(\ell_2) [\ell_2 - \frac{1}{2}\ell_2^2]$. For the expert to be indifferent between recommending the pooling price and the separating price at ℓ_2 . we have

$$A \times \left[\frac{1}{2}[\ell_1 + \ell_2] - \frac{1}{2}\ell_2^2 \right] = \alpha(\ell_2) \left[\ell_2 - \frac{1}{2}\ell_2^2 \right]. \quad (8)$$

Given ℓ_1 and ℓ_2 , equations (7) and (8) define a two-unknown-two-variable linear equation system. We can solve for A and $\alpha(\ell_2)$ given ℓ_1 and ℓ_2 . In fact, we have

$$A = \frac{\left[\frac{2}{\ell_1} - 1 \right] \left[\ell_1 - \frac{1}{2}\ell_1^2 \right]}{\left[\frac{1}{2}[\ell_1 + \ell_2] - \frac{1}{2}\ell_1^2 \right]}$$

$$\alpha(\ell_2) = \frac{A \left[\frac{1}{2}[\ell_1 + \ell_2] - \frac{1}{2}\ell_2^2 \right]}{\ell_2 - \frac{1}{2}\ell_2^2} = \frac{\left[\frac{2}{\ell_1} - 1 \right] \left[\ell_1 - \frac{1}{2}\ell_1^2 \right]}{\left[\frac{1}{2}[\ell_1 + \ell_2] - \frac{1}{2}\ell_1^2 \right]} \times \frac{\left[\frac{1}{2}[\ell_1 + \ell_2] - \frac{1}{2}\ell_2^2 \right]}{\left[\ell_2 - \frac{1}{2}\ell_2^2 \right]}.$$

It can be verified that both A and $\alpha(\ell_2)$ are less than one. We also require the boundary condition for pooling to make nonnegative profits: $\frac{1}{2}[\ell_1 + \ell_2] - \frac{1}{2}\ell_2^2 \geq 0$. For example, if $\ell_1 = 1$, then the largest pooling range, given by ℓ_2 , is $\frac{1}{2}[1 + \ell_2] = \frac{1}{2}\ell_2^2$, which gives the maximum ℓ_2 to be $\frac{1 + \sqrt{5}}{2} = 1.618$ as in Example 1. As ℓ_1 increases from 1, the maximum ℓ_2 to sustain pooling increases.

The two-parameter example includes the separating equilibrium in Proposition 3, and the pooling equilibrium in Proposition 2, as special cases, respectively at $(\ell_1 = 1, \ell_2 = 2)$ and $(\ell_1 = 1, \ell_2 = \frac{1}{2}[1 + \sqrt{5}])$. We can construct other hybrid equilibria following similar constructions.

6 Robustness and extensions

6.1 Random losses and costs

Now we let a consumer's loss and cost be both random. Denote the cost support by $[\underline{c}, \bar{c}]$, a positive interval. Then the random variable $(\ell, c) \in [\underline{\ell}, \bar{\ell}] \times [\underline{c}, \bar{c}]$ has a joint distribution $F : [\underline{\ell}, \bar{\ell}] \times [\underline{c}, \bar{c}] \rightarrow [0, 1]$ and a density

function $f(\ell, c) > 0$, $(\ell, c) \in [\underline{\ell}, \bar{\ell}] \times [\underline{c}, \bar{c}]$. (Do note that we abuse notation by using the same F to denote a random vector of losses and costs.) In other words, now a consumer's type has two dimensions described by (ℓ, c) . A single dimensional model is one where ℓ and c are perfectly correlated, so we would be able to write the repair cost as (a deterministic function) $C(\ell)$, as in Section 2. The increasing C assumption in Section 2 now is replaced by the assumption that ℓ and c are positively correlated. The extensive form is rewritten here:

Stage 1: A consumer visits the expert, who then observes the consumer's potential loss $\ell \in [\underline{\ell}, \bar{\ell}]$ and repair cost $c \in [\underline{c}, \bar{c}]$, drawn with the joint distribution F . The expert chooses between offering to repair at a price $p \in [\underline{\ell}, \bar{\ell}]$ and refusing to repair.

Stage 2: If the consumer agrees to the repair, she pays the expert p , and the expert incurs cost c for the repair.

The expert's strategy is a map: $P : [\underline{\ell}, \bar{\ell}] \times [\underline{c}, \bar{c}] \rightarrow [\underline{\ell}, \bar{\ell}]$, which prescribes a price offer when the provider observes consumer (ℓ, c) . A consumer's strategy is the same as before: an acceptance probability of the offered price. The consumer never gets to observe loss ℓ or cost c .

In a perfect-Bayesian equilibrium, the consumer updates belief about the (ℓ, c) distribution, but her utility depends only ℓ and the price. The expert does not directly care about the consumer's loss because profit only depends on the repair cost and the price. Due to these preferences, in fact, the two-dimensional model is isomorphic to the one-dimensional model. The precise relationship is given by the following.

Proposition 8 *In any perfect-Bayesian equilibrium, the expert's strategy $P : [\underline{\ell}, \bar{\ell}] \times [\underline{c}, \bar{c}] \rightarrow [\underline{\ell}, \bar{\ell}]$ is (almost everywhere) independent of ℓ whenever the price $P(\ell, c)$ is accepted with a positive probability.*

The expert's equilibrium price is only dependent on his cost, never on the consumer's loss. It is the price-cost margin that generates profit; any price function varying with losses cannot be sustained as an equilibrium strategy. The main thrust of Proposition 8 is this. In Section 2, we have assumed a deterministic relationship between loss and cost, and that the expert observes the loss. Obviously, it makes no difference if the expert

observes cost (instead of loss). In that case, the pricing strategy is a function of cost. From the (separating or pooling) equilibrium strategy, the consumer will infer losses, and all results remain the same.¹⁶

Now Proposition 8 says that for a two-dimensional model, any equilibrium price function depends on cost only. The consumer's equilibrium inference therefore concerns the expected losses conditional on cost. Let $L(c)$ denote the expected losses conditional on c :

$$L(c) \equiv \frac{\int_{\underline{\ell}}^{\bar{\ell}} \ell dF(\ell, c)}{\int_{\underline{\ell}}^{\bar{\ell}} dF(\ell, c)}.$$

A positive correlation between loss and cost implies that $L(c)$ is increasing. Hence the function L acts like the inverse of the cost function C in Section 2. In other words, the deterministic, one-dimensional model in Section 2 is the reduced form of the stochastic, two-dimensional model here: they yield the same set of equilibria. Our use of a deterministic cost function is without loss of generality.

6.2 Many experts

Now we consider a market with many consumers and many experts. There are $N \geq 2$ identical experts and a consumer with a loss drawn from F . The consumer randomly chooses an expert. The expert perfectly observes the consumer's loss and makes a price offer. The consumer then decides whether to accept the price. If the consumer accepts the price, the expert must perform the repair. If the consumer rejects, she can exit the market or repeat the search process until she visits all experts. The consumer must pay a small search cost to visit an expert, except for the first visit. The search cost captures the consumer's disutility from waiting. Following the literature (Wolinsky (1993)), we assume that an expert does not know a consumer's search history. Hence, the expert cannot make a price offer contingent on the consumer's past searches or price offers, if any.

We claim that the separating equilibrium in Proposition 3 and the pooling equilibrium in Proposition

¹⁶Suppose that $\ell = D(c)$ where $D \equiv C^{-1}$ is the inverse of the cost function in Section 2. The expert's strategy now is $P : [\underline{c}, \bar{c}] \rightarrow [\underline{\ell}, \bar{\ell}]$, with $\underline{c} = C(\underline{\ell})$ and $\bar{c} = C(\bar{\ell})$. Given a price function P , the consumer will try to infer the loss. For example, if $P(c) = c + k$, then $c = P(c) - k$ and the inferred loss is $D(P(c) - k)$. More generally, if P is strictly increasing over a range of c , then it is invertible, so the inverted value of c can be used to infer loss ℓ . If P is a constant function over a range of c , then at that constant, the consumer will infer that losses belong to an interval.

2 remain equilibria in the market with multiple experts. Precisely, it is an equilibrium for each expert and each consumer to play their respective strategies in Proposition 3; likewise, this is also true for Proposition 2. Recall that the monopolist expert plays pure strategies in each of the separating and pooling equilibria. If all experts play the same pure strategy, a consumer expects to have the same payoff from visiting a different expert, but must incur a search cost. In equilibrium, therefore, a consumer does not search, which reinforces an expert's ability to maintain monopoly power. Hence, the strategy profiles in the monopoly market will continue to constitute an equilibrium in the market with multiple experts.

To elaborate on the argument, suppose that all experts play the same strategy as in the separating equilibrium in Proposition 3. Given that experts will recommend the same price $P(\ell) = \ell$ and that search involves a small cost, a consumer will not solicit a second opinion. Since the consumer is indifferent between accepting the first expert's recommendation and exiting the market, it is a best response for the consumer to accept price ℓ with probability $\alpha(\ell)$ as in the separating equilibrium in Proposition 3. If the consumer does not search for more opinions, each expert is dealing with an inexperienced consumer. Given the consumer's acceptance function $\alpha(\ell)$, it is optimal for an expert to adopt the pricing strategy $P(\ell) = \ell$ for the same argument in Proposition 3.

We have claimed that an equilibrium outcome in the single-expert model remains an equilibrium outcome in the many-expert model, and it is reminiscent of the Diamond (1971) paradox (which asserts that when consumers face any positive search cost, the unique equilibrium has competitive firms setting the monopoly price). In our model, the consumer faces uncertainty about their losses and searches for second opinions. In contrast, Diamond considers the case that the consumer knows his valuation and only searches for cheaper prices. Because of the different modeling assumptions on consumer's information, our result is quite different from Diamond. First, there are multiple (pooling and separating) equilibria in our many-expert model, but there is a unique equilibrium in the Diamond setting. Second, an expert here obtains superior information seeing a consumer; firms in the Diamond setting learn nothing. Third, we have only claimed the existence of a symmetric equilibrium: all experts follow the same pooling or separating equilibrium strategy in the single-expert model. Indeed, there might be asymmetric equilibria in which different experts use different pricing strategies and consumers search for different prices.

6.3 Alternative cost function: efficient repair only at medium losses

We have made the assumption that $C(\ell) < \ell$ for all $\ell < \widehat{\ell}$. A different scenario may have costs higher than losses at low and high loss values; that is $C(\ell) < \ell$ if and only if $\ell \in [\underline{\ell}, \widehat{\ell}]$ where $\underline{\ell} < \ell_0 < \widehat{\ell} < \bar{\ell}$, but we maintain the assumption that C is an increasing function.¹⁷ How do results change under this cost function specification? First, the separating equilibrium will fail to exist. Clearly, if the expert's price offers at ℓ between $\underline{\ell}$ and ℓ_0 reveal the losses, then all these prices must be below cost $C(\ell)$. The expert must reject them; otherwise the expert would have been better off refusing to offer treatment. Given this, Part i) of Proposition 1 says that there cannot be any repair for any higher ℓ . But this means that there is no acceptance for any losses, so a separating equilibrium fails to exist.

Second, the pooling equilibrium in Proposition 2 remains robust, and the equilibrium price is still $AL(\underline{\ell}, \ell^*)$. Here, of course, losses between $\underline{\ell}$ and ℓ_0 will be repaired inefficiently. It follows that the pooling equilibrium involves service over-provision at low losses and under-provision at high losses.

Third, we can continue to construct a continuum of hybrid equilibria, as in Proposition 6, and they still will be eliminated by the nonnegative profit principle by the same argument in Proposition 7. However, any hybrid equilibrium cannot have $P(\ell) = \ell$ for low values of ℓ , so the more complicated example hybrid equilibrium in the previous subsection would not exist under the modified cost function.

7 Conclusion

We study complex credence-good problems. The consumer's problem may result in a continuum of losses. An equilibrium consists of two functions: the expert's pricing function and the consumer's acceptance function. The model has two functions as primitives: the consumer's loss distribution, and the expert's cost. We characterize properties that must hold for all perfect-Bayesian equilibria. In particular, we derive pooling and separating equilibria in closed forms. All equilibria are inefficient, but full efficiency can be restored if the expert gets to choose the diagnostic test.

¹⁷For example, an automobile's very minor problems may not be worth the repair cost, and the same may be true for very serious problems. Repairs are cost effective for intermediate conditions.

We have left open a host of other issues in the market of credence goods. First, we have assumed risk neutrality. Obviously consumers may be risk averse, so the expert's repair may be a way to dampen utility fluctuations. Will the consumer's demand for insurance give the expert a higher leverage to expropriate surplus? On the other hand, the expert may also have to face risks if the diagnostic test cannot fully reveal the loss level.

Second, we have assumed limited contracting options. Only the only verifiable event is the restoration of the baseline utility, so partial repair is ruled out. In practice, partial repairs are quite common. An analysis of partial repair will have to be based on a larger set of verifiable outcomes.

Third, we have used a static approach to model information about losses. Information about potential losses, such as in health and house repair settings, may be coming along during treatment or repair. Renegotiation about price may have to happen because the ultimate repair costs may change.

Fourth, we consider the market environment in which the treatment cost is not verifiable. Our framework can be used to analyze the case when treatment cost is verifiable but the expert is not liable for treatment outcome.

Appendix

Proof of Proposition 1: We begin with i). Consider two losses ℓ and ℓ' . Adding the no-deviation constraints in (1) and (2) we have

$$\begin{aligned} -C(\ell')\tilde{\alpha}_P(\ell') - C(\ell)\tilde{\alpha}_P(\ell) &\geq -C(\ell')\tilde{\alpha}_P(\ell) - C(\ell)\tilde{\alpha}_P(\ell') \\ [C(\ell) - C(\ell')][\tilde{\alpha}_P(\ell) - \tilde{\alpha}_P(\ell')] &\leq 0. \end{aligned}$$

Clearly if $\ell > \ell'$, we have $C(\ell) > C(\ell')$. so $\tilde{\alpha}_P(\ell) \leq \tilde{\alpha}_P(\ell')$.

Next, to prove ii), suppose that it is false. That is, suppose that $P(\ell) < P(\ell')$ when $\ell > \ell'$. Then we have

$$[P(\ell) - C(\ell)]\tilde{\alpha}_P(\ell) < [P(\ell') - C(\ell)]\tilde{\alpha}_P(\ell) \leq [P(\ell') - C(\ell)]\tilde{\alpha}_P(\ell'),$$

where the first inequality follows from $\tilde{\alpha}_P(\ell) > 0$ and the second inequality follows from i). But this contradicts (2). We conclude that $P(\ell) \geq P(\ell')$, and $P(\ell)$ is weakly increasing.

Finally, we prove iii). Suppose $\ell > \ell'$. Then,

$$\Pi(\ell') \equiv [P(\ell') - C(\ell')]\tilde{\alpha}_P(\ell') \geq [P(\ell) - C(\ell')]\tilde{\alpha}_P(\ell) \geq [P(\ell) - C(\ell)]\tilde{\alpha}_P(\ell) \equiv \Pi(\ell).$$

The first inequality follows from (1) and the second inequality follows from $C(\ell') < C(\ell)$. If $\tilde{\alpha}_P(\ell) > 0$, the last weak inequality becomes a strict inequality.

Proof of Lemma 1: If equilibrium P is strictly increasing on (ℓ_1, ℓ_2) , then it is invertible, so from $P(\ell)$, the consumer infers that the loss is ℓ , which is her highest willingness to pay for repair. Hence, $P(\ell) \leq \ell$ if the offer is accepted with a positive probability. Suppose that $P(\ell) < \ell$ for some ℓ ; this implies that the offer $P(\ell)$ will be accepted with probability 1. Consider $\ell' < \ell$. At loss ℓ' , the equilibrium profit is $[P(\ell') - C(\ell')]\alpha(P(\ell'))$. Because P is strictly increasing, we have $P(\ell') < P(\ell)$. Let the expert deviate to price $P(\ell)$ at loss ℓ' , the profit now becomes $P(\ell) - C(\ell') > [P(\ell') - C(\ell')]\alpha(P(\ell'))$, which says that $P(\ell')$ fails to be an equilibrium price, a contradiction. We conclude that $P(\ell) = \ell$.

Proof of Proposition 3: We prove a more general result. Suppose that $P(\ell) = \ell$ for $\ell \in (\ell_1, \ell_2) \subseteq [\underline{\ell}, \hat{\ell}]$. Then the consumer will be indifferent between accepting and rejecting the price offer. We construct an acceptance probability function against which the expert's best response is indeed $P(\ell) = \ell$.

Let $\alpha(P(\ell)) = \alpha(\ell)$ be the equilibrium acceptance probability for $\ell \in (\ell_1, \ell_2)$. (We have dispensed with the $\tilde{\alpha}_P$ notation used in Proposition 1 because P is an identity map.) Define $\pi(\ell'; \ell) = [\ell' - C(\ell)] \alpha(\ell')$, any $\ell', \ell \in (\ell_1, \ell_2)$. For $P(\ell) = \ell$ to be an equilibrium price, we must have

$$[\ell' - C(\ell')] \alpha(\ell') \geq [\ell - C(\ell')] \alpha(\ell) \quad \text{and} \quad [\ell - C(\ell)] \alpha(\ell) \geq [\ell' - C(\ell)] \alpha(\ell'). \quad (9)$$

In other words, in equilibrium, we have $\pi(\ell; \ell) \geq \pi(\ell'; \ell)$ and $\pi(\ell'; \ell') \geq \pi(\ell; \ell')$.

Let $\Pi(\ell) \equiv \pi(\ell; \ell)$. Then, we have

$$\begin{aligned} \Pi(\ell) &\geq \pi(\ell'; \ell) = [\ell' - C(\ell)] \alpha(\ell') \\ &= [\ell' - C(\ell')] \alpha(\ell') - [C(\ell) - C(\ell')] \alpha(\ell') \\ &= \Pi(\ell') - [C(\ell) - C(\ell')] \alpha(\ell'). \end{aligned}$$

By symmetry, we have

$$\Pi(\ell') \geq \Pi(\ell) - [C(\ell') - C(\ell)] \alpha(\ell).$$

Combining and then dividing by $\ell - \ell'$, we have $\frac{-[C(\ell) - C(\ell')] \alpha(\ell)}{\ell - \ell'} \geq \frac{\Pi(\ell) - \Pi(\ell')}{\ell - \ell'} \geq \frac{-[C(\ell) - C(\ell')] \alpha(\ell')}{\ell - \ell'}$.

By taking limits, we have $\Pi'(\ell) = -C'(\ell) \alpha(\ell)$.

Next, we use the definition of $\Pi(\ell) = [\ell - C(\ell)] \alpha(\ell)$ to obtain

$$\Pi'(\ell) = [\ell - C(\ell)] \alpha'(\ell) + [1 - C'(\ell)] \alpha(\ell).$$

Substituting by $\Pi'(\ell) = -C'(\ell) \alpha(\ell)$, we have

$$[\ell - C(\ell)] \alpha'(\ell) + [1 - C'(\ell)] \alpha(\ell) = -C'(\ell) \alpha(\ell). \quad (10)$$

Hence, we get a differential equation for $\alpha(\ell)$

$$(\ell - C(\ell)) \alpha'(\ell) + \alpha(\ell) = 0 \quad \text{or} \quad \frac{d \ln \alpha(\ell)}{d \ell} = -\frac{1}{\ell - C(\ell)}. \quad (11)$$

We solve the differential equation for $\alpha(\ell)$:

$$\begin{aligned} \ln \alpha(\ell) &= - \int_{\ell_1}^{\ell} \frac{dx}{x - C(x)} + K \quad \text{some constant } K \\ \alpha(\ell) &= K' \exp \left\{ - \int_{\ell_1}^{\ell} \frac{dx}{x - C(x)} \right\} \quad \text{some constant } K' \\ \alpha(\ell) &= \alpha(\ell_1) \exp \left\{ - \int_{\ell_1}^{\ell} \frac{dx}{x - C(x)} \right\}, \end{aligned} \quad (12)$$

where we have relabelled K' as $\alpha(\ell_1)$, the acceptance probability at ℓ_1 .

The separating equilibrium is the case when $\ell_1 = \underline{\ell}$, and $\ell_2 = \widehat{\ell}$. The consumer must accept the lowest price with probability 1 because $\ell_1 > C(\ell_1)$, so $\alpha(\underline{\ell}) = 1$. We have obtained the equilibrium acceptance probability function in the Proposition.

For $\ell > \widehat{\ell}$, the consumer's maximum willingness to pay ℓ is less than the cost $C(\ell)$, but $P(\ell) \geq C(\ell)$ in an equilibrium. Hence, the consumer's strategy is optimal, and the expert's not offering a repair is also optimal.

Proof of Proposition 2: Given the expert's strategy, the consumer infers that her loss is drawn from $[\underline{\ell}, \ell^*]$ when the expert offers $P(\ell) = AL(\underline{\ell}, \ell^*)$, so her expected loss is $AL(\underline{\ell}, \ell^*)$. The consumer is indifferent between accepting and rejecting it, so it is a best response for her to accept the offer. If the consumer is recommended an off-the-equilibrium price, she continues to believe that her loss is drawn from $[\underline{\ell}, \ell^*]$ according to $F(\ell)$. Hence, the consumer will accept any price lower than $AL(\underline{\ell}, \ell^*)$ and reject any price higher than $AL(\underline{\ell}, \ell^*)$.

Next, given that the consumer accepts prices below $AL(\underline{\ell}, \ell^*)$ and rejects all higher prices, the expert offers price $AL(\underline{\ell}, \ell^*)$ if and only if $AL(\underline{\ell}, \ell^*) \geq C(\ell)$ for consumer with loss ℓ . From the definition of ℓ^* , we have $AL(\underline{\ell}, \ell^*) = C(\ell^*) < C(\ell)$ for any $\ell > \ell^*$. Hence offering $P(\ell) = AL(\underline{\ell}, \ell^*)$ is the expert's best response if and only if $\ell \leq \ell^*$. Given that the consumer rejects all prices higher than $AL(\underline{\ell}, \ell^*)$, it is the expert's best response to recommend no treatment for $\ell > \ell^*$.

Proof of Corollary 1: The expert's expected profit of the separating equilibrium, Π^S , is

$$\begin{aligned} \Pi^S &\equiv \int_{\underline{\ell}}^{\widehat{\ell}} [\ell - C(\ell)] \alpha(\ell) dF(\ell) \\ &= - \int_{\underline{\ell}}^{\widehat{\ell}} F(\ell) d[\ell - C(\ell)] \alpha(\ell) \\ &= \int_{\underline{\ell}}^{\widehat{\ell}} F(\ell) C'(\ell) \alpha(\ell) d\ell, \end{aligned}$$

where the first equality follows from integration by parts and the definition of $\widehat{\ell} = C(\widehat{\ell})$, and where the last equality follows from (10).

The expected profit of the pooling equilibrium, Π^P , is

$$\begin{aligned}
\Pi^P &\equiv \int_{\underline{\ell}}^{\ell^*} (\ell - C(\ell)) dF(\ell) \\
&= F(\ell^*)[\ell^* - C(\ell^*)] - \int_{\underline{\ell}}^{\ell^*} F(\ell)[1 - C'(\ell)] d\ell \\
&= F(\ell^*)[\ell^* - C(\ell^*)] - \int_{\underline{\ell}}^{\ell^*} F(\ell) d\ell + \int_{\underline{\ell}}^{\ell^*} F(\ell) C'(\ell) d\ell \\
&= \int_{\underline{\ell}}^{\ell^*} F(\ell) C'(\ell) d\ell,
\end{aligned}$$

where the first equality follows from integration by parts. The last equality is obtained by substitution:

$$\begin{aligned}
\int_{\underline{\ell}}^{\ell^*} F(\ell) d\ell &= F(\ell^*)\ell^* - \int_{\underline{\ell}}^{\ell^*} \ell dF(\ell) \\
&= F(\ell^*) \left(\ell^* - \frac{\int_{\underline{\ell}}^{\ell^*} \ell dF(\ell)}{F(\ell^*)} \right) \\
&= F(\ell^*)[\ell^* - C(\ell^*)],
\end{aligned}$$

where the first equality follows from integration by parts, and the last equality follows the definition of ℓ^* .

Proof of Proposition 5: We begin with the pooling equilibrium in Proposition 2. Consider a price deviation $p' > AL(\underline{\ell}, \ell^*)$. Define ℓ' by $p' = C(\ell')$. According to the Nonnegative Profit Principle, upon being offered price p' , the consumer believes that her loss is drawn from $[\underline{\ell}, \ell']$ according to $F(\ell)$. Because $C(\ell^*) = AL(\underline{\ell}, \ell^*) < p' = C(\ell')$, $\ell^* < \ell'$, and $AL(\underline{\ell}, \ell') < C(\ell') = p'$. Therefore, the consumer will reject p' . Now, consider a price deviation $p' < AL(\underline{\ell}, \ell^*)$. Such a price deviation is not profitable even if the consumer accepts p' with probability one because the transaction happens at a lower price. We conclude that the pooling equilibrium in Proposition 2 satisfies the Nonnegative Profit Principle.

Consider a pooling equilibrium in Corollary 2 indexed by $\tilde{\ell} < \ell^*$. Suppose the expert offers an off-equilibrium price $p' = AL(\underline{\ell}, \ell^*) - \varepsilon$, $\varepsilon > 0$. For a small enough ε , we apply the same argument in the first paragraph of this proof, so the consumer believes that her loss is drawn from $[\underline{\ell}, \ell']$, where $p' = C(\ell')$. Because $C(\ell') = p' = AL(\underline{\ell}, \ell^*) - \varepsilon < AL(\underline{\ell}, \ell^*) = C(\ell^*)$, $\ell' < \ell^*$, and ℓ' is arbitrarily close to ℓ^* when ε is sufficiently small. By the definition of ℓ^* , $AL(\underline{\ell}, \ell)$ crosses $C(\ell)$ from above at ℓ^* . Therefore, $AL(\underline{\ell}, \ell') > C(\ell') = p'$. So, the consumer must accept p' . The expert now has a profitable deviation from the lower price $AL(\underline{\ell}, \tilde{\ell})$ to the

higher price $p' = AL(\underline{\ell}, \ell^*) - \varepsilon$ for all $\ell \in [\underline{\ell}, \tilde{\ell}]$. We conclude that pooling equilibria in Corollary 2 do not satisfy the Nonnegative Profit Principle.

Proof of Proposition 6: Given the expert's strategy, the consumer is always indifferent between accepting and rejecting the on-equilibrium price, so her strategy against the price offer at $AL(\underline{\ell}, \tilde{\ell})$ and in $[\tilde{\ell}, \hat{\ell}]$ is optimal. At any off-equilibrium price between $AL(\underline{\ell}, \tilde{\ell})$ and $\tilde{\ell}$, the consumer continues to believe that the loss is drawn from $[\underline{\ell}, \tilde{\ell}]$, so strictly prefers to reject the (off-equilibrium) price. Upon a price offer above $\hat{\ell}$, the consumer believes that her loss is $\hat{\ell}$ and hence will reject the offer.

We next consider the expert's optimal choices at each ℓ , given the consumer's strategy in the Proposition. First, suppose that the expert has a consumer with loss $\ell < \tilde{\ell}$. According to the strategy in the Proposition, his payoff is $AL(\underline{\ell}, \tilde{\ell}) - C(\ell) > 0$. He earns less profit by lowering the price, and makes zero profit by offering a price between $AL(\underline{\ell}, \tilde{\ell})$ and $\tilde{\ell}$, according to the consumer's strategy in the Proposition. Suppose that the expert offers a price $\ell' > \tilde{\ell}$. In this case, his profit becomes $[\ell' - C(\ell)] \alpha(\ell') \equiv \pi(\ell'; \ell)$. Take the derivative

$$\begin{aligned} \frac{\partial \pi(\ell'; \ell)}{\partial \ell'} &= \alpha(\ell') + [\ell' - C(\ell)] \frac{d\alpha(\ell')}{d\ell'} \\ &= \alpha(\ell') \left[1 + \left[\ell' - C(\ell) \right] \frac{\alpha'(\ell')}{\alpha(\ell')} \right] \\ &= \alpha(\ell') \left[1 - \frac{\ell' - C(\ell)}{\ell' - C(\ell')} \right] \\ &< 0, \end{aligned}$$

where the third equality follows from (11) and the inequality follows from $C(\ell) < C(\ell')$. It follows that the expert's best deviation is at $\tilde{\ell}$. If the expert does deviate to price $\tilde{\ell}$, his expected profit is $[\tilde{\ell} - C(\ell)] \alpha(\tilde{\ell})$. We compare this candidate deviation profit with $AL(\underline{\ell}, \tilde{\ell}) - C(\ell)$ as follows:

$$\begin{aligned} &AL(\underline{\ell}, \tilde{\ell}) - C(\ell) - \left[\tilde{\ell} - C(\ell) \right] \alpha(\tilde{\ell}) \\ &= AL(\underline{\ell}, \tilde{\ell}) - \tilde{\ell} \alpha(\tilde{\ell}) - C(\ell) [1 - \alpha(\tilde{\ell})] \\ &= C(\tilde{\ell}) [1 - \alpha(\tilde{\ell})] - C(\ell) [1 - \alpha(\tilde{\ell})] \\ &= [C(\tilde{\ell}) - C(\ell)] [1 - \alpha(\tilde{\ell})] \geq 0, \end{aligned}$$

where the second equality follows from the definition of $\alpha(\tilde{\ell})$, and the last inequality follows from $C(\tilde{\ell}) > C(\ell)$. Hence, the expert's maximum profit is from offering price $AL(\underline{\ell}, \tilde{\ell})$ when ℓ is between $\underline{\ell}$ and $\tilde{\ell}$.

Second, suppose that the expert has a consumer with loss $\ell \geq \tilde{\ell}$. According to the Proposition, he should obtain an expected profit $[\ell - C(\ell)]\alpha(\ell) > 0$. The proof of Proposition 3 shows that he cannot earn more profit by deviating to any price above $\tilde{\ell}$. Clearly, the expert cannot profit from deviating to a price between $AL(\underline{\ell}, \tilde{\ell})$ and $\tilde{\ell}$ (because it will be rejected). Because the consumer always accepts the price $AL(\underline{\ell}, \tilde{\ell})$, it remains to show that he cannot earn more by deviating to price $AL(\underline{\ell}, \tilde{\ell})$. The profit from this deviation is $AL(\underline{\ell}, \tilde{\ell}) - C(\ell)$. From the definition of $\alpha(\tilde{\ell})$, we have $AL(\underline{\ell}, \tilde{\ell}) - C(\tilde{\ell}) = \alpha(\tilde{\ell})[\tilde{\ell} - C(\tilde{\ell})]$, hence

$$\begin{aligned} AL(\underline{\ell}, \tilde{\ell}) - C(\ell) &< \alpha(\tilde{\ell})[\tilde{\ell} - C(\tilde{\ell})] \\ &\leq \alpha(\ell)[\ell - C(\ell)], \end{aligned}$$

where the first inequality is due to $C(\ell) > C(\tilde{\ell})$ and $\alpha(\tilde{\ell}) < 1$, and the second inequality is from the proof of Proposition 3. We have verified that the strategies form an equilibrium.

Proof of Proposition 8: Consider an equilibrium, let $P(\ell, c)$ be the equilibrium price when the expert has a type (ℓ, c) consumer. Let $\alpha(P(\ell, c)) > 0$ be the corresponding acceptance probability when the consumer is offered the price $P(\ell, c)$. In an equilibrium an expert with a type (ℓ, c) consumer chooses a price to obtain equilibrium profit $\max_{\ell', c'} \{\alpha(P(\ell', c'))[P(\ell', c') - c]\} \equiv \Pi(c)$, which is independent of loss ℓ .

Consider equilibrium prices $P(\ell_1, c_1)$ and $P(\ell_2, c_2)$ for any (ℓ_1, c_1) and $(\ell_2, c_2) \in [\underline{\ell}, \bar{\ell}] \times [\underline{c}, \bar{c}]$. Hence, $\Pi(c_1) = \alpha(P(\ell_1, c_1))[P(\ell_1, c_1) - c_1]$ and $\Pi(c_2) = \alpha(P(\ell_2, c_2))[P(\ell_2, c_2) - c_2]$. Let $c_2 > c_1$

$$\begin{aligned} \Pi(c_2) &\geq \alpha(P(\ell_1, c_1))[P(\ell_1, c_1) - c_2] \\ &= \alpha(P(\ell_1, c_1))[P(\ell_1, c_1) - c_1 + c_1 - c_2] = \Pi(c_1) - (c_2 - c_1)\alpha(P(\ell_1, c_1)). \end{aligned}$$

Hence we have

$$\frac{\Pi(c_2) - \Pi(c_1)}{(c_2 - c_1)} \geq -\alpha(P(\ell_1, c_1)).$$

Analogously,

$$\begin{aligned} \Pi(c_1) &\geq \alpha(P(\ell_2, c_2))[P(\ell_2, c_2) - c_1] \\ &= \alpha(P(\ell_2, c_2))[P(\ell_2, c_2) - c_2 + c_2 - c_1] = \Pi(c_2) + (c_2 - c_1)\alpha(P(\ell_2, c_2)). \end{aligned}$$

Hence we have

$$-\alpha(P(\ell_2, c_2)) \geq \frac{\Pi(c_2) - \Pi(c_1)}{(c_2 - c_1)}.$$

Combining we have

$$-\alpha(P(\ell_2, c_2)) \geq \frac{\Pi(c_2) - \Pi(c_1)}{(c_2 - c_1)} \geq -\alpha(P(\ell_1, c_1)), \quad \text{for } c_1 < c_2, \text{ and for any } \ell_1 \text{ and } \ell_2. \quad (13)$$

Now consider $c_1 < c < c_2$. We now have

$$-\alpha(P(\ell_2, c_2)) \geq \frac{\Pi(c_2) - \Pi(c)}{(c_2 - c)} \geq -\alpha(P(\ell, c)), \quad \text{for } c < c_2, \text{ and for any } \ell \text{ and } \ell_2$$

where we have let c take the role of c_1 in (13). Let (ℓ_2, c_2) converge to (ℓ, c) , and we have the right-hand derivative satisfying

$$\left(\frac{d\Pi(c)}{dc} \right)^+ \geq -\alpha(P(\ell, c)).$$

Next, we have

$$\begin{aligned} -\alpha(P(\ell, c)) &\geq \frac{\Pi(c) - \Pi(c_1)}{(c - c_1)} \geq -\alpha(P(\ell_1, c_1)), \\ -\alpha(P(\ell, c)) &\geq \frac{\Pi(c_1) - \Pi(c)}{(c_1 - c)} \geq -\alpha(P(\ell_1, c_1)), \quad \text{for } c_1 < c, \text{ and for any } \ell \text{ and } \ell_1, \end{aligned}$$

where we have let c take the role of c_2 in (13). Let (ℓ_1, c_1) converge to (ℓ, c) , and we have the left-hand derivative satisfying

$$-\alpha(P(\ell, c)) \geq \left(\frac{d\Pi(c)}{dc} \right)^-.$$

Now $\Pi(c)$ is the maximum of an affine linear function in c , so it is convex, and almost everywhere differentiable. The left-hand and right-hand derivatives are equal, so we have

$$\left(\frac{d\Pi(c)}{dc} \right)^- = \left(\frac{d\Pi(c)}{dc} \right)^+ \geq -\alpha(P(\ell, c)) \geq \left(\frac{d\Pi(c)}{dc} \right)^- \quad \text{for any } \ell.$$

The above inequalities show that $\alpha(P(\ell, c))$ is independent of ℓ . Suppose the equilibrium price $P(\ell, c)$ varies in ℓ for a fixed c . Let $\tilde{\ell}$ denote the loss that leads to the highest equilibrium price $P(\tilde{\ell}, c)$. Because $\alpha(P(\ell, c))$ is constant in ℓ , the expert will deviate to recommending $P(\tilde{\ell}, c)$ for all $\ell \neq \tilde{\ell}$. We conclude that P must be independent of ℓ .

References

- [1] Akerlof, G.A. (1970), “The market for lemons: quality uncertainty and the market mechanism.” *Quarterly Journal of Economics*, 84, 488-500.
- [2] Alger, I., and Salanie, F. (2006), “A theory of fraud and over-treatment in experts markets.” *Journal of Economics and Management Strategy*, 15, 853-881.
- [3] Bagwell, K., and Riordan, M.H. (1991), “High and declining prices signal product quality.” *American Economic Review*, 81, 224-239.
- [4] Balafoutas, L., Beck, A., Kerschbamer, R., and Sutter, M. (2013), “What drives taxi drivers? A field experiment on fraud in a market for credence goods.” *Review of Economic Studies*, 80, 876-891.
- [5] Balafoutas, L., and Kerschbamer, R. (2020), “Credence goods in the literature: what the past fifteen years have taught us about fraud, incentives, and the role of institutions.” *Journal of Behavioral and Experimental Finance*, 26, 100285.
- [6] Balafoutas, L., Fornwagner, H., Kerschbamer, R., Sutter, M., and Tverdostup, M. (2021), “Diagnostic uncertainty and insurance coverage in credence goods markets.” Working paper.
- [7] Bester, H., and Dahm, M. (2018), “Credence goods, costly diagnosis, and subjective evaluation.” *Economic Journal*, 128, 1367-1394.
- [8] Chen Y., Li, J., and Zhang, J. (2021), “Efficient Liability in Expert Markets.” Working paper
- [9] Darby, R.M., and Karni E. (1973), “Free competition and the optimal amount of fraud.” *Journal of Law and Economics*, 16, 67-88.
- [10] Diamond, P. (1971), “A Model of Price Adjustment,” *Journal of Economic Theory*, III, 156-168.
- [11] Dulleck, U., and Kerschbamer, R. (2009), “Experts vs. discounters: consumer free-riding and experts withholding advice in markets for credence goods.” *International Journal of Industrial Organization*, 27, 15-23.

- [12] Dulleck, U., and Kerschbamer, R. (2006), “On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods.” *Journal of Economic Literature*, 44, 5–42.
- [13] Dulleck, U., Kerschbamer, R., and Sutter, M. (2011), “The economics of credence goods: on the role of liability, verifiability, reputation and competition.” *American Economic Review*, 101, 526–55.
- [14] Emons, W. (1997), “Credence goods and fraudulent experts.” *RAND Journal of Economics*, 2, 107–119.
- [15] Fong, Y. F. (2005), “When do experts cheat and whom do they target.” *RAND Journal of Economics*, 36, 113-130.
- [16] Fong, Y. F., and Liu, T. (2018), “Liability and reputation in credence goods markets.” *Economics Letters*, 166, 35-39.
- [17] Fong Y. F., Liu, T. and Wright, D. (2014), “On the role of verifiability and commitment in credence goods markets.” *International Journal of Industrial Organization*, 37, 118-129.
- [18] Fong Y. F., Liu, T., and Meng, X. (2022), “Trust building in credence goods markets.” *American Economic Journal–Microeconomics*, 14, 490-528.
- [19] Gottschalk, F., Mimra, W., and Waibel, C. (2020), “Health services as credence goods: a field experiment.” *Economic Journal*, 130, 1346-1383.
- [20] Kerschbamer, R., Sutter, M, and Dulleck, U. (2017), “How social preference shape incentives in (experimental) markets for credence goods.” *Economic Journal*, 127, 271-570.
- [21] Kim, K. (2012), “Endogenous market segmentation for lemons.” *Rand Journal of Economics*, 43, 562-576.
- [22] Liu, T. (2011), “Credence goods markets with conscientious and selfish experts.” *International Economic Review*, 52, 227-244.
- [23] Liu, T., Zhou, Y., Vergara Cobos, E. (2019), “Pricing schemes and seller fraud: evidence from New York City taxi rides.” *Journal of Industrial Economics*, 67, 56-90.

- [24] Mimra, W., Rasch, A., and Waibel, C. (2016), “Price competition and reputation in credence goods markets: experimental evidence.” *Games and Economic Behavior*, 100, 337-352.
- [25] Pesendorfer, W., and Wolinsky, A. (2003), “Second opinions and price competition: inefficiency in the market for expert advice.” *Review of Economic Studies*, 70, 417-437.
- [26] Pitchik, C., and Schotter, A. (1987), “Honesty in a model of strategic information transmission.” *American Economic Review*, 77, 1032-1036.
- [27] Taylor, C. (1995), “The economics of breakdowns, checkups, and cures.” *Journal of Political Economy*, 103, 53-74.
- [28] Schneider, H. (2012), “Agency problems and reputation in expert services: evidence from auto repair.” *Journal of Industrial Economics*, 60, 406–433.
- [29] Wolinsky, A. (1983), “Prices as signals of product quality.” *Review of Economic Studies*, 50, 647-658.
- [30] Wolinsky, A. (1993), “Competition in a market for informed experts’ services.” *RAND Journal of Economics*, 24: 380-398.
- [31] Wolinsky, A. (1995), “Competition in markets for credence goods.” *Journal of Institutional and Theoretical Economics*, 151, 117-131.