Games and Economic Behavior 145 (2024) 84–101

Contents lists available at ScienceDirect

Games and Economic Behavior

journal homepage: www.elsevier.com/locate/geb

Equilibrium information in credence goods ☆

Ting Liu [a,*], Ching-to Albert Ma [b]

[a] Department of Economics, Stony Brook University, 100 Nicolls Rd, Stony Brook, NY, 11794, USA
[b] Department of Economics, Boston University, USA

ARTICLE INFO

JEL classification:
D80
D82
D83

Keywords:
Credence goods
Experts
Separating equilibrium
Pooling equilibrium
Information acquisition

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.

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 may not 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. For example, a car's check-engine warning light can indicate more than ten problems, from a loose fuel cap to a defective catalytic converter (according to "11 Reasons for the Check Engine Light," in bumper.com). Different problems reduce a vehicle's performance to different degrees, so the consumer's value from repair depends on which fault the warning light indicates. Obviously this complexity arises in many other markets.

☆ **Acknowledgment:** For their comments and suggestions, we thank Jacobo Bizzotto, Lester Chan, Winand Emons, Sanghoon Kim, Henry Mak, Chiara Margaria, Juan Ortner, Edi Karni, and seminar participants at John Hopkins University, Boston University, Buffalo University, CREST in Paris, National Taiwan University, Oslo Metropolitan University, Queen's University in Kingston, Ontario, Stony Brook University, University of New South Wales, University of Queensland, Webinar Series on Credence Goods and Expert Markets 2021, Midwest Theory Conference 2021, International Industrial Organization Conference 2022, Econometric Society North American Summer Meetings 2022, and International Game Theory Festival 2022. We thank an editor and three reviewers for their comments and suggestions.

* Corresponding author.
E-mail addresses: ting.liu@stonybrook.edu (T. Liu), ma@bu.edu (C.A. Ma).

1 Pesendorfer 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. Chiu and Karni (2021) study a competitive market where consumers may have N, $N \geq 2$, finite losses.

https://doi.org/10.1016/j.geb.2024.03.002
Received 19 December 2022

Available online 15 March 2024
0899-8256/

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 model with finite 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 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. It is efficient to repair the problem if the loss is not too high. 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.

Our paper offers three new insights. First, we identify a novel pooling equilibrium absent in the existing literature. Second, we show that the expert can make more profit by revealing little information than full information. This finding contrasts sharply with the binary model's well-known result that the expert makes the highest profit by reporting honestly consumers' problems (Fong (2005)). Third, if the expert can choose diagnostic information endogenously, he optimally will acquire coarse information even if information acquisition is costless. We explain the three main results in turn.

In the pooling equilibrium, the expert recommends a price if the consumer's loss is below an endogenously determined 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 a 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 inefficiently.

The pooling equilibrium is absent in the binary model. If there are two losses, the only candidate pooling equilibrium is one in which the expert recommends the same price for both losses. This pooling equilibrium trivially solves the information asymmetry problem, so is dismissed by the common assumption that the average loss is lower than the higher repair cost.[2] Whereas we adopt the same assumption to rule out the "one-price-fix-all" pooling equilibrium, we endogenize the pooling threshold for an inefficient equilibrium.

We also derive a unique separating equilibrium, in which the expert charges the consumer her loss from the problem, so the consumer can infer losses from prices. The separating equilibrium is reminiscent of the honest equilibrium in the binary model. The consumer rejects each price with a positive probability that 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.

Profit comparison between pooling and separating equilibria coincides with welfare comparison because consumers always have zero (ex ante or ex post) equilibrium surplus. Both separating and pooling equilibria 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 equilibrium 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; the condition depends on the loss distribution and the cost function. The profit comparison indicates that the expert can make more profits from revealing coarse information than full information. Again, this is in stark contrast with the binary model, which has implicitly led one to believe that the highest profit results from an honest (separating) equilibrium.

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 pooling equilibrium is robust against a perturbation in the cost function while the separating equilibrium is sensitive to it (see Subsection 6.3).

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 are 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.

[2] See Pitchik and Schotter (1987), Wolinsky (1993), Fong (2005), Fong and Liu (2018), Fong et al. (2022).

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 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 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.[3] 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 for all equilibria. Then we characterize pooling and separating equilibria in closed forms. Next, Section 4 is about expert acquiring information. Section 5 discusses hybrid equilibria, which amalgamate pooling and separating equilibria, and the role of price commitment. In Section 6 we consider three extensions: a two-dimensional model in loss and cost, multiple experts, and an alternative class of cost functions. We offer some concluding remarks in Section 7. The Appendix contains proofs that are omitted in the main text. An Online Appendix contains a model with finite number of losses and its analysis, a noisy information structure, examples with specific loss distributions, and formal details about hybrid equilibria.

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 and Liu (2018), Fong et al. (2022), Chiu and Karni (2021), Karni (2023). The second line assumes that only the expert's repair 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. (2022).

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 while inflating 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 is 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 hidden actions; these papers include Pesendorfer and Wolinsky (2003), Dulleck and Kerschbamer (2009), Bester and Dahm (2018), and Chen et al. (2022). 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: 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 is unprofitable.

[3] For 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.

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. (2017), 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 many 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). Our work allows rich strategy sets, so 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 use prices to signal their goods' qualities. In each paper, there exists a separating equilibrium in which high prices signal high qualities. A common driving force for the separating equilibrium stems from some consumers being informed about qualities before purchase. We also identify a separating equilibrium in which high prices signal high losses. In our model, consumers are never informed about their losses before or after purchase. The separating equilibrium emerges due to the increasing probability of consumer rejection for higher prices. Kim (2012) studies endogenous market segmentation to alleviate information asymmetry, where 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 (details in the Online Appendix).

Our paper is also related to the literature on strategic information revelation by information intermediaries. Lizzeri (1999) and Mahenc (2017) study the role of certificate agency in mitigating inefficiency caused by asymmetric information. Whereas Lizzeri (1999) and Mahenc (2017) show that certificate agency may improve welfare by providing information to consumers, we demonstrate an alternative perspective: that the welfare can be improved by limiting the information available to the expert.

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 a random variable ℓ with distribution function F on the strictly positive support $[\underline{\ell}, \overline{\ell}]$; the distribution F is assumed to be absolutely continuous, so 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 $C : [\underline{\ell}, \overline{\ell}] \to \Re_+$ 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 $\ell < \hat{\ell}$, with $\underline{\ell} < \hat{\ell} \leq \overline{\ell}$, which says that it is efficient to treat only problems with losses lower than $\hat{\ell}$. 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, either before or after repair. The assumption that ℓ is unobservable to the consumer after a repair rules out contracting based on loss. 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.[4] The only verifiable event is that a repair avoids the loss, so the expert must perform the repair once the price offer is accepted.[5]

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 $[\underline{\ell}, \overline{\ell}]$, we can let prices belong to $[\underline{\ell}, \overline{\ell}]$ without 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 $\ell \in [\underline{\ell}, \overline{\ell}]$, drawn from distribution F. The expert offers to repair at a price $p \in [\underline{\ell}, \overline{\ell}] \cup \{+\infty\}$ (again, a price at $+\infty$ means the expert refusing to repair).

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: $P : [\underline{\ell}, \overline{\ell}] \to [\underline{\ell}, \overline{\ell}] \cup \{+\infty\}$, with $P(\ell) \in [\underline{\ell}, \overline{\ell}]$ 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: $\alpha : [\underline{\ell}, \overline{\ell}] \to [0, 1]$, with $\alpha(p)$ denoting the consumer's probability of accepting a repair offer $p \in [\underline{\ell}, \overline{\ell}]$.

Definition 1. A perfect-Bayesian equilibrium consists of a strategy profile (P, α) and a belief system. For any loss ℓ, the expert's pricing strategy $P(\ell)$ maximizes his profit given the consumer's strategy $\alpha(p)$; the consumer's strategy $\alpha(p)$ maximizes her expected

[4] In Subsection 6.1, we assume that costs and losses are random, but positively correlated. This relaxes the cost monotonicity assumption.

[5] Our analysis is not contingent on the expert's knowledge of B. The consumer's willingness to pay for a repair is determined by her loss and independent of the baseline utility. As a result, the expert's price does not depend on B.

payoff given $P(\ell)$ and her belief about the loss. The consumer's belief about the loss is updated according to the expert's strategy and Bayes rule whenever possible.

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 $P([\underline{\ell},\widetilde{\ell})) = p_1$ and $P([\widetilde{\ell},\overline{\ell}]) = p_2$, so the consumer responds to prices p_1 and p_2 on the equilibrium path. Bayes rule dictates that on observing p_1, the consumer believes that the loss is drawn from $[\underline{\ell},\widetilde{\ell})$ according to the prior F restricted to $[\underline{\ell},\widetilde{\ell})$, and that on observing p_2, it is from $[\widetilde{\ell},\overline{\ell}]$. Prices other than p_1 and p_2 never get offered in equilibrium. Yet, the consumer's equilibrium strategy must specify her acceptance probability at these prices.

Clearly, the consumer's responses at off-equilibrium price depend on the consumer's belief about the loss. We are free to specify these off-equilibrium beliefs to support a perfect-Bayesian equilibrium outcome. Not all off-equilibrium beliefs are reasonable or plausible. We will introduce a refinement, what we will call the Nonnegative Profit Principle, to be defined in Section 3; essentially it says that if a consumer receives a price offer, the consumer should believe that the expert's repair cost must be lower than that price—otherwise the expert will make a loss with that offer.

We do not include a price-posting stage in the extensive form. The literature usually assumes that the expert first posts a price list before diagnosis and then commits to recommending a price from the list if he is willing to treat the consumer. In our model, the expert does not post a set of prices before learning the loss. In complex repairs, it is uncommon that experts would make binding estimates before diagnosis. For example, few hospitals in the U.S. post prices of procedures online, despite such a federal requirement since 2021. More important, a combination of procedures may be needed for a repair, so posting prices of each procedure may not 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. We discuss the implications of including a price-posting stage in Subsection 5.2.

3. Equilibria

The continuous loss model admits many equilibria. We begin by deriving necessary conditions on players' equilibrium strategies and the expert's equilibrium profit. Guided by the equilibrium conditions, we then derive pooling and separating equilibria in closed forms.

3.1. *Perfect-Bayesian equilibrium strategies*

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

Next, for any given equilibrium strategy profile (P,α), define a function $\alpha_P : [\underline{\ell},\overline{\ell}] \to [0,1]$ by $\alpha_P(\ell) \equiv \alpha(P(\ell))$. The function α_P tracks the acceptance probability on the equilibrium path according to strategy profile (P,α), over all loss values; it does so by composing the equilibrium acceptance probability and the equilibrium price functions. Let $\Pi(\ell) \equiv \alpha_P(\ell) \times \{P(\ell) - C(\ell)\}$ denote the equilibrium profit from repairing ℓ, where, for brevity, we have omitted the equilibrium profile (P,α) in the argument of Π. The expert's expected equilibrium profit is $\int \Pi(\ell)\, \mathrm{d}\, F(\ell)$. We now state properties valid for all equilibria.

Proposition 1. *Any perfect-Bayesian equilibrium strategy profile (P,α) has the following properties:*
i) The equilibrium acceptance probability $\alpha_P(\ell)$ is weakly decreasing.
ii) The price $P(\ell)$ is weakly increasing if $\alpha_P(\ell) > 0$, $\ell \in [\underline{\ell},\overline{\ell}]$.
iii) Equilibrium profit $\Pi(\ell)$ is weakly decreasing in $\ell \in [\underline{\ell},\overline{\ell}]$ and is strictly decreasing in $\ell \in [\ell_1,\ell_2] \subset [\underline{\ell},\overline{\ell}]$ if $\alpha_P(\ell) > 0$ for $\ell \in [\ell_1,\ell_2]$.

Proof of Proposition 1. Let (P,α) be a perfect-Bayesian equilibrium. Define $\pi(\ell';\ell) \equiv \left[P(\ell') - C(\ell)\right]\alpha_p(\ell')$, for $\ell',\ell \in [\underline{\ell},\overline{\ell}]$. Hence, $\pi(\ell';\ell)$ is the expert's profit from recommending $P(\ell')$, the price meant for loss ℓ', when the consumer's loss is ℓ. Any perfect-Bayesian equilibrium (P,α) must satisfy the following no-deviation conditions: at two losses ℓ and ℓ'

$$\pi(\ell;\ell) = [P(\ell) - C(\ell)]\,\alpha_P(\ell) \geq \left[P(\ell') - C(\ell)\right]\alpha_P(\ell') = \pi(\ell';\ell) \tag{1}$$

$$\pi(\ell';\ell') = \left[P(\ell') - C(\ell')\right]\alpha_P(\ell') \geq \left[P(\ell) - C(\ell')\right]\alpha_P(\ell) = \pi(\ell;\ell'). \tag{2}$$

Adding the no-deviation constraints in (1) and (2) we have

$$-C(\ell)\alpha_P(\ell) - C(\ell')\alpha_P(\ell') \geq -C(\ell)\alpha_P(\ell') - C(\ell')\alpha_P(\ell)$$
$$[C(\ell) - C(\ell')][\alpha_P(\ell) - \alpha_P(\ell')] \leq 0.$$

Clearly if $\ell > \ell'$, we have $C(\ell) > C(\ell')$, so $\alpha_P(\ell) \leq \alpha_P(\ell')$. We have proven i).

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

$$[P(\ell) - C(\ell)]\alpha_P(\ell) < [P(\ell') - C(\ell)]\alpha_P(\ell) \leq [P(\ell') - C(\ell)]\alpha_P(\ell'),$$

where the first inequality follows from $\alpha_P(\ell) > 0$ and the second inequality follows from i). But this contradicts (1). We conclude that $P(\ell) \geq P(\ell')$, and $P(\ell)$ is weakly increasing.

Finally, suppose $\ell > \ell'$. Then

$$\Pi(\ell') \equiv [P(\ell') - C(\ell')]\alpha_P(\ell') \geq [P(\ell) - C(\ell')]\alpha_P(\ell) \geq [P(\ell) - C(\ell)]\alpha_P(\ell) \equiv \Pi(\ell).$$

The first inequality follows because it is optimal to recommend $P(\ell')$ and induce the acceptance rate $\alpha_P(\ell')$ for ℓ', and the second inequality follows from $C(\ell') < C(\ell)$. If $\alpha_P(\ell) > 0$ in some interval $[\ell_1, \ell_2]$, the last weak inequality becomes a strict inequality. This proves iii). ■

Part i) of Proposition 1 says that the consumer's acceptance probability must not increase with loss. Part ii) says that the expert must not decrease price when the repair cost goes up. These are intuitive equilibrium properties. Reducing prices when costs increase would just be against profit maximization; a higher acceptance probability at a higher price would encourage the expert to recommend high prices at low losses. Parts i) and ii) together say that as the loss and cost increase, the price increases but the acceptance rate decreases. Part iii) further says that the equilibrium profit must decrease with losses; again, this follows from profits varying inversely with costs. Proposition 1 derives properties on any equilibrium path, so off-path beliefs are irrelevant.

By 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) but can perfectly infer her losses from the price in case (ii). The last case (iii) is a combination of cases (i) and (ii). We characterize different types of equilibria, and begin with definitions.

Definition 2 *(Pooling equilibrium).* An equilibrium (P, α) is said to be pooling if for any ℓ and ℓ' in $[\underline{\ell}, \overline{\ell}]$ for which $\alpha(P(\ell)) > 0$ and $\alpha(P(\ell')) > 0$, then $P(\ell) = P(\ell')$.

Definition 3 *(Separating equilibrium).* An equilibrium (P, α) is said to be separating if for any ℓ and ℓ' in $[\underline{\ell}, \overline{\ell}]$, $\ell \neq \ell'$, 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 definitions are for *all* losses in $[\underline{\ell}, \overline{\ell}]$ with positive acceptance probabilities. Hence, these two definitions are not exhaustive. Continuum of equilibria that are neither separating nor pooling do exist. In the Online Appendix, these hybrid equilibria will be defined and analyzed.

3.2. Pooling equilibria

We begin with some notation and definitions.

Definition 4 *(Conditional average loss).* The expected loss in the interval $[\ell_1, \ell_2]$ is given by the function $AL : [\underline{\ell}, \overline{\ell}] \times [\underline{\ell}, \overline{\ell}] \to [\underline{\ell}, \overline{\ell}]$:

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

Of particular relevance is the expected loss conditional on the loss below ℓ, $AL(\underline{\ell}, \ell)$. Clearly, $AL(\underline{\ell}, \ell) < \ell$ for $\ell > \underline{\ell}$, and $AL(\underline{\ell}, \overline{\ell}) = \int \ell \ \mathrm{d}\ F(\ell) \equiv \mu$. Next we define a special set of losses.

Definition 5 *(Pooling equilibria indexed by losses).*

$$LP \equiv \{\ell : AL(\underline{\ell}, \ell) \geq C(\ell)\} \quad \text{and} \quad \ell^* \equiv \max\{\ell : AL(\underline{\ell}, \ell) \geq C(\ell)\}$$

The set LP denotes all losses each of which having a conditional expected repair value $AL(\underline{\ell}, \ell)$ larger than cost $C(\ell)$. The maximal element of LP is denoted as ℓ^*. We construct pooling equilibria indexed by elements of LP. The maximal element ℓ^* will index a special pooling equilibrium.

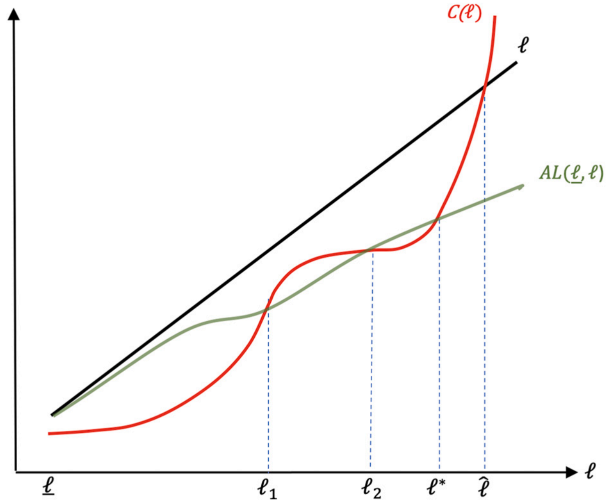

Fig. 1. A pooling equilibrium.

Now, both $AL(\underline{\ell}, \ell)$ and $C(\ell)$ are increasing, so may intersect multiple times. Hence, the set LP may consist of disjoint loss intervals. Recall, by the definition of $\hat{\ell}$, $C(\ell) < \ell$ if and only if $\ell < \hat{\ell}$. Obviously $AL(\underline{\ell}, \ell^*) = C(\ell^*)$: beyond ℓ^* cost will never be less than conditional average loss. Moreover, $\ell^* < \hat{\ell}$ because $C(\ell^*) = AL(\underline{\ell}, \ell^*) < \ell^*$ and $C(\ell) < \ell$ if and only if $\ell < \hat{\ell}$.

An example of various functions is in Fig. 1; we have drawn losses on the 45-degree line, the conditional average loss function AL, and the cost function C. The cost function $C(\ell)$ is higher than the conditional expected loss function $AL(\underline{\ell}, \ell)$ between ℓ_1 and ℓ_2. The set LP is $[\underline{\ell}, \ell^*] \backslash (\ell_1, \ell_2)$.

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

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

The consumer's strategy is

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

If the consumer is recommended $AL(\underline{\ell}, \ell^*)$, she believes that the loss is drawn from $[\underline{\ell}, \ell^*]$ according to $F(\ell)$ restricted to $[\underline{\ell}, \ell^*]$. The consumer holds the same belief when she receives any price offer different from $AL(\underline{\ell}, \ell^*)$.

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-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 $AL(\underline{\ell}, \ell^*) \ge C(\ell)$ 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 $\ell \le \ell^*$. 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 $\ell > \ell^*$. ■

The expert offers a single equilibrium price $AL(\underline{\ell}, \ell^*)$ that results in transactions. Upon receiving the price $AL(\underline{\ell}, \ell^*)$, the consumer believes that losses are in $[\underline{\ell}, \ell^*]$. The consumer accepts all prices below $AL(\underline{\ell}, \ell^*)$. The consumer rejects higher prices because she simply sticks with the equilibrium belief at price $AL(\underline{\ell}, \ell^*)$. 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^*. In the pooling equilibrium, losses below ℓ^* are repaired whereas losses between ℓ^* and $\hat{\ell}$ are not; these nonrepairs are inefficient. The consumer can only infer from the expert's strategy whether her problem is below or above ℓ^*. The pooling equilibrium is illustrated in Fig. 1.

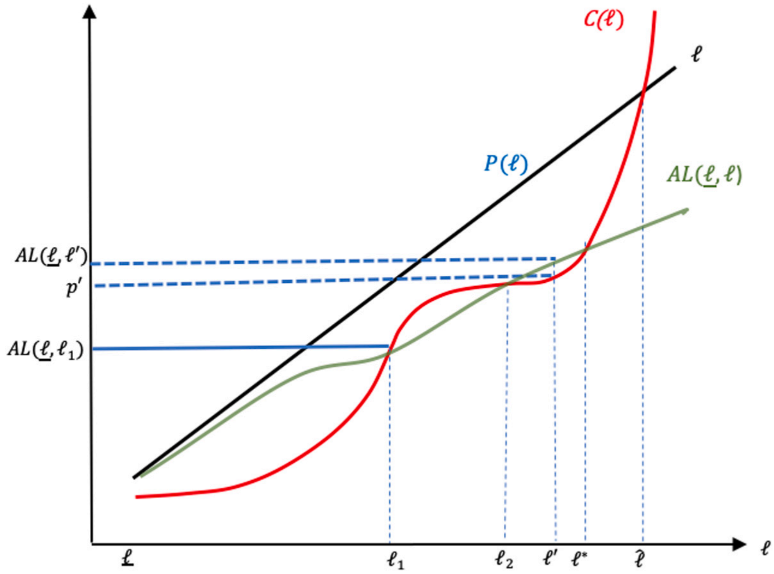

Fig. 2. The Nonnegative Profit Principle.

The off-equilibrium belief in Proposition 2 resembles *passive belief* due to McAfee and Schwartz (1994); the consumer never changes beliefs about losses even at off-equilibrium prices. Other beliefs that result in rejection of higher prices can support the equilibrium outcome in Proposition 2. For example, upon being offered an off-equilibrium, higher price, the consumer could believe that the loss was actually lower, say from an internal $[\underline{\ell}, \ell^{**}]$ with $\ell^{**} < \ell^*$. Indeed, such "optimistic" off-equilibrium belief can support other pooling equilibria.

Next, we show that there are other pooling equilibria which can be supported by arbitrary off-equilibrium beliefs. Then, we introduce the Nonnegative Profit Principle belief refinement to reject all these other pooling equilibria except the one in Proposition 2.

Corollary 1. *Each $\widetilde{\ell} \in LP \equiv \{\ell : AL(\underline{\ell}, \ell) \geq C(\ell)\}$ indexes a pooling equilibrium. The expert's strategy is to offer price $\widetilde{p} = C(\widetilde{\ell})$ if and only if $\ell \in [\underline{\ell}, \widetilde{\ell}]$. The consumer accepts $\widetilde{p}$ and rejects all others. The consumer believes that the loss is $\underline{\ell}$ if the expert offers any other price, and rejects it.*

The proof is straightforward, and omitted. The expert's strategy is already stated in the Corollary. The consumer's accepts $\widetilde{p}$ because, by definition of LP, we have $\widetilde{p} = C(\widetilde{\ell}) \leq AL(\underline{\ell}, \widetilde{\ell})$. To support the pooling equilibrium, we only need to specify off-equilibrium belief to be sufficiently optimistic, so that any higher price would be rejected. For example, upon receiving a price offer higher than $\widetilde{p}$, the consumer would become extremely optimistic and believe that the loss was at the minimal level, $\underline{\ell}$.

Optimistic and passive beliefs, as in the statements of Proposition 3 and Corollary 1, are unsophisticated, and fail to account for any notion of *forward induction*. We now introduce a refinement.

Definition 6 *(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 find it not credible 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 the truncated prior distribution of losses.

Proposition 3. *The pooling equilibrium in Proposition 2 satisfies the Nonnegative Profit Principle. A pooling equilibrium indexed by $\widetilde{\ell} < \ell^*$ in Corollary 1 does not satisfy the Nonnegative Profit Principle.*

The proof of the proposition is in the Appendix. (Proofs of other results are also in the Appendix unless they are already laid out in the text.) Fig. 2 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.

3.3. Separating equilibrium

To begin, we present the following key lemma.

Lemma 1. *Suppose that the equilibrium price function P is strictly increasing over an interval of losses, 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 be equal to 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 $P(\ell) = \ell$.

For the binary model, in the separating (honest) equilibrium, the consumer accepts the minor repair with probability one and rejects the major repair with a positive probability. The consumer's rejection probability makes the expert just indifferent between recommending the minor and the major repairs when the consumer's problem is minor. With a continuum of losses, it is impossible to find an acceptance function that makes the expert just indifferent between all prices for any given loss. Can we find the acceptance probability function to support a separating equilibrium?

Proposition 4. *There is a unique separating equilibrium which involves the threshold $\hat{\ell}$, defined by $C(\hat{\ell}) = \hat{\ell}$. The expert's strategy is a price function*

$$P(\ell) = \begin{cases} \ell & \text{if } \underline{\ell} \le \ell \le \hat{\ell} \\ +\infty & \text{if } \hat{\ell} < \ell \le \overline{\ell} \end{cases},$$

and the consumer's strategy is an acceptance probability function:

$$\alpha(\ell) = \begin{cases} \exp\left\{ -\int_{\underline{\ell}}^{\ell} \dfrac{dx}{x - C(x)} \right\} & \text{if } \underline{\ell} \le \ell \le \hat{\ell} \\ 0 & \text{if } \hat{\ell} < \ell \le \overline{\ell} \end{cases}.$$

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 4 does not have any unreached information set in the relevant price range $[\underline{\ell}, \hat{\ell}]$. 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 4 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.

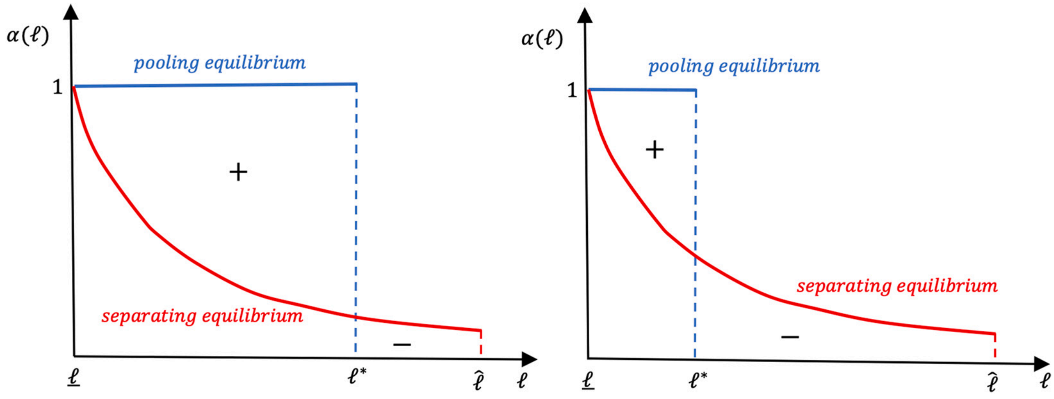

Fig. 3. Two examples comparing pooling and separating equilibria.

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 $\ell^* < \overline{\ell}$. 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. 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 equilibrium 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 2. *The pooling equilibrium in Proposition 2 yields more social surplus (and profit) than the separating equilibrium in Proposition 4 if and only if*

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

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

Fig. 3 shows the probabilities that loss ℓ is resolved in the pooling and separating equilibria. In the pooling equilibrium, losses below ℓ^* are always repaired, whereas in the separating equilibrium, such losses are sometimes left untreated. Hence, the pooling equilibrium yields a higher social surplus than the separating equilibrium for losses lower than ℓ^*. On the other hand, losses between ℓ^* and $\hat{\ell}$ are unrepaired with probability one in the pooling equilibrium but repaired with a positive probability in the separating equilibrium. Hence, in these cases, the pooling equilibrium results in a smaller social surplus than the separating equilibrium. The left panel of Fig. 3 illustrates the case where the pooling equilibrium dominates the separating equilibrium ex ante. The right panel depicts another scenario in which the separating equilibrium dominates the pooling equilibrium.

Corollary 2 gives the condition under which the pooling equilibrium is more efficient and profitable than the separating equilibrium. For a given distribution function F, condition (4) is more likely to hold if the cost function is more convex, which makes ℓ^* closer to $\hat{\ell}$. The impact of the distribution function F on the profit comparison is more subtle. In the Online Appendix, we present an example in which F follows a triangular distribution, and investigate how the shape of F affects condition (4). There, the pooling equilibrium always dominates the separating equilibrium. Additionally, the profit difference between the pooling equilibrium and the separating equilibrium becomes smaller as the density function f becomes flatter.

We work out two examples that verify Corollary 2:

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 $[\underline{\ell}, \overline{\ell}]$, $\underline{\ell} > 0$, and $C(\ell) = b\ell - k$, with $b > 1$ and $k \in ((b-1)\underline{\ell}, (b-1)\overline{\ell})$. In this case, $\hat{\ell} = \frac{k}{b-1}$, $\ell^* = \frac{l + 2k}{2b-1}$ and $\alpha(\ell) = \left[\frac{k-(b-1)\ell}{k-(b-1)\underline{\ell}}\right]^{-\frac{1}{1-b}}$. The difference between the left-hand side and the right-hand side of (4) is $\frac{[k-(b-1)\underline{\ell}]^2}{(\overline{\ell}-\underline{\ell})(2b-1)^2} > 0$. 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 2 and the above two examples show results that would be unavailable in a binary model. Fong (2005) shows that the expert's profit and efficiency is maximized by the separating equilibrium. In the separating equilibrium of the binary model, the minor loss is efficiently avoided but the major loss is sometimes left untreated. If there are many losses, all but the lowest loss may be left untreated. So, when problems are complex, the separating equilibrium may yield a lower profit and efficiency than the pooling equilibrium.

Propositions 1 to 2 together show that *all* equilibria are inefficient. Efficiency requires all losses $\ell \in [\underline{\ell}, \hat{\ell}]$ 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 $\hat{\ell}$, 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 making a price offer. In this section, we study the expert's information acquisition decision. 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 $S \subset \Re$ and a joint distribution, $J: S \times [\underline{\ell}, \overline{\ell}] \to [0,1]$, 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) $\int_S \mathrm{d}\, J(s,\ell) = J(\sup(S),\ell) = F(\ell)$ and ii) $J(s,\overline{\ell}) \equiv G(s)$ yields a probability distribution over S. The requirement i) implies that $\int_{\underline{\ell}}^{\overline{\ell}} \int_S \ell \,\mathrm{d}\, J(s,\ell) = \int_{\underline{\ell}}^{\overline{\ell}} \ell \,\mathrm{d}\, F(\ell) = \mu$. 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.[6] 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.

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 $\hat{\ell}$ (at which point, $\hat{\ell} = C(\hat{\ell})$). 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 4 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 $\hat{\ell}$, 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	$j(s_L,\ell)$	$j(s_H,\ell)$
$\underline{\ell} \le \ell \le \hat{\ell}$	$f(\ell)$	0
$\hat{\ell} \le \ell \le \overline{\ell}$	0	$f(\ell)$
signal probabilities	$F(\hat{\ell}) = \Pr(S = s_L)$	$1 - F(\hat{\ell}) = \Pr(S = s_H)$

(5)

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

$$k(\ell \mid S = s_L) = \begin{cases} \dfrac{f(\ell)}{F(\hat{\ell})} & \text{for } \ell \le \hat{\ell} \\ 0 \text{ for } & \ell > \hat{\ell} \end{cases} \quad \text{and} \quad k(\ell \mid S = s_H) = \begin{cases} 0 \text{ for } \ell \le \hat{\ell} \\ \dfrac{f(\ell)}{1-F(\hat{\ell})} & \text{for } \ell > \hat{\ell} \end{cases}. \tag{6}$$

[6] That is, there is an increasing, bijective function $\phi : [\underline{\ell}, \overline{\ell}] \to [\underline{s}, \overline{s}]$ such that $F(\ell) = G(s)$ where $s = \phi(\ell)$ and $G(s) \equiv \int_{\underline{\ell}}^{\overline{\ell}} \mathrm{d}\, J(s,\ell)$. The simplest case is where ϕ is the identity map.

Upon observing signal s_L, the expert's expected repair cost is $\dfrac{\int_{\underline{\ell}}^{\hat{\ell}} C(\ell)\mathrm{d}F(\ell)}{F(\hat{\ell})} \equiv AC(\underline{\ell},\hat{\ell})$, the average cost of repairing losses below $\hat{\ell}$. By the definition of $\hat{\ell}$, $\ell > C(\ell)$ for all $\ell < \hat{\ell}$. Hence, $AL(\underline{\ell},\hat{\ell}) > AC(\underline{\ell},\hat{\ell})$. 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 $AL(\underline{\ell},\hat{\ell})$.

Proposition 5. *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*

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

and the consumer's strategy is

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

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 $AL(\underline{\ell},\hat{\ell})$ if and only if the signal is s_L; the consumer always accepts an offer if it does not exceed $AL(\underline{\ell},\hat{\ell})$.

The proof of this Proposition is obvious, and omitted. The expert's equilibrium expected profit is $\left[AL(\underline{\ell},\hat{\ell}) - AC(\underline{\ell},\hat{\ell})\right] F(\hat{\ell}) = \int_{\underline{\ell}}^{\hat{\ell}} [\ell - C(\ell)] \,\mathrm{d}\, F(\ell)$, which is the first-best surplus. However, ex post, the repair cost for some losses may be above the equilibrium price. Indeed, $C(\hat{\ell}) = \hat{\ell} > AL(\underline{\ell},\hat{\ell})$, so for losses close to $\hat{\ell}$, 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 5, 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 5 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 $AL(\underline{\ell},\hat{\ell})$. 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 5 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.[7] However, in practice, consumers may not understand how a diagnostic test actually works, or how a signal is to be interpreted. Proposition 5 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.

The information structure described in Proposition 5 precisely reveals whether the consumer's loss is above or below $\hat{\ell}$. Often, tests are imprecise, so false positives and negatives are possible. In the Online Appendix, we generalize the information structure by allowing a small noise. With imperfect tests, both excessive and insufficient repair can occur in equilibrium. As the noise vanishes, the equilibrium outcome converges to the first best. Proposition 5 is robust against slight perturbation in information precision.

5. Discussions

5.1. Hybrid equilibria

As we have mentioned, pooling and separating equilibria do not exhaust all equilibrium classes. We can "cut and paste" pooling and separating equilibria to construct what we call hybrid equilibria. The definition of hybrid equilibria is provided in the last section of the Online Appendix.

[7] 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.

In that last section, we explicitly construct an equilibrium of the following sort. For low loss values, the equilibrium is pooling; the expert offers a single price. Then the equilibrium price jumps upwards when the loss level reaches a threshold; thereafter, the equilibrium is separating, the expert offering a price equal the loss value. The consumer accepts the pooling price with probability one, but rejects the separating prices with positive probabilities to support the equilibrium. A hybrid equilibrium combines pooling and separation. Other more complicated hybrid equilibria, with multiple pooling and separating intervals, can be constructed. We show that the Nonnegative Profit Principle refinement can eliminate some, but not all, hybrid equilibria.

5.2. Price commitment

The existing literature with binary losses includes a price-posting stage at the beginning of the game. After a price list has been posted, the expert learns the consumer's potential loss, and either recommends a price from the list or refuses to provide treatment. The list typically has two prices, one for each loss. Each price list defines a subgame which has a unique equilibrium. Hence, the price posting stage allows the expert to select the most profitable equilibrium. Two difficulties arise with price commitment when losses are continuous.

First, as we show in the Online Appendix, there is a continuum of subgames with hybrid equilibria in each. Although we can characterize the property of the hybrid equilibria, we are unable to rank their profits. As a result, we cannot pin down the price menu which yields the expert the highest profit.

Second, price commitment may not be useful for selecting the most profitable equilibrium because there are generally multiple equilibria following a price menu. For example, suppose the separating equilibrium yields the expert the highest profit under some parameter configurations. The posted price menu for the separating equilibrium is $[\underline{\ell}, \widehat{\ell}]$. However, $[\underline{\ell}, \widehat{\ell}]$ contains all relevant prices in the game without price posting. It follows that any equilibrium without price commitment remains an equilibrium in the subgame following the price menu $[\underline{\ell}, \widehat{\ell}]$.

In the Online Appendix, we have explored price commitment when there are 3 possible losses. Either the pooling or the separating equilibrium may emerge as the equilibrium of the whole game.

6. Extensions

6.1. Random losses and costs

Now we let a consumer's loss and cost both be random. Denote the cost support by $[\underline{c}, \overline{c}]$, a positive interval. Then the random variable $(\ell, c) \in [\underline{\ell}, \overline{\ell}] \times [\underline{c}, \overline{c}]$ has a joint distribution $F : [\underline{\ell}, \overline{\ell}] \times [\underline{c}, \overline{c}] \to [0, 1]$ and a density function $f(\ell, c) > 0$, $(\ell, c) \in [\underline{\ell}, \overline{\ell}] \times [\underline{c}, \overline{c}]$. (Do note that we abuse notation by using the same F to denote a random vector of losses and costs.) 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 $\ell \in [\underline{\ell}, \overline{\ell}]$ and repair cost $c \in [\underline{c}, \overline{c}]$, drawn with the joint distribution F. The expert chooses between offering to repair at a price $p \in [\underline{\ell}, \overline{\ell}] \cup \{+\infty\}$.
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: $P : [\underline{\ell}, \overline{\ell}] \times [\underline{c}, \overline{c}] \to [\underline{\ell}, \overline{\ell}]$, 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 6. *In any perfect-Bayesian equilibrium, the expert's strategy $P : [\underline{\ell}, \overline{\ell}] \times [\underline{c}, \overline{c}] \to [\underline{\ell}, \overline{\ell}]$ 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 6 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.[8]

[8] Suppose 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 $P : [\underline{c}, \overline{c}] \to [\underline{\ell}, \overline{\ell}]$, with $\underline{c} = C(\underline{\ell})$ and $\overline{c} = C(\overline{\ell})$. 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

Now Proposition 6 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:

$$L(c) \equiv \frac{\int_{\underline{\ell}}^{\overline{\ell}} \ell \mathrm{d}F(\ell, c)}{\int_{\underline{\ell}}^{\overline{\ell}} \mathrm{d}F(\ell, c)}.$$

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. 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 there are $N \geq 2$ identical experts and many consumers, with each consumer having a loss independently drawn from F. A 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 4 and the pooling equilibrium in Proposition 2 remain equilibria in the market with multiple experts. Precisely, it is an equilibrium for each expert and the consumer to play their respective strategies in Proposition 4; likewise, this is 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 4. 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 4. 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 4.

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 $\ell < \hat{\ell}$. A different scenario may have costs higher than losses at low and high loss values; that is $C(\ell) < \ell$ if and only if $\ell \in [\ell_0, \hat{\ell}]$ where $\underline{\ell} < \ell_0 < \hat{\ell} < \overline{\ell}$, but we maintain the assumption that C is an increasing function.[9] 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 made negative profits. 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. In fact, the separating equilibrium requires any loss below $\hat{\ell}$ to be resolved with a strictly positive probability. Therefore, the separating equilibrium will fail to exist if $C(\ell) > \ell$ for some but not all losses lower than ℓ_0.

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.

[9] For 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.

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. Moreover, the pooling equilibrium is the unique equilibrium that satisfies the Nonnegative Profit Principle. This observation highlights that the pooling equilibrium is more robust than other types of equilibrium under this alternative cost function.

7. Conclusion

We study complex credence-good problems. The consumer's problem is a loss among 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.

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. 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.

Declaration of competing interest

None.

Data availability

No data was used for the research described in the article.

Appendix A

Proof of Proposition 3. 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 1 indexed by $\widetilde{\ell} < \ell^*$. 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 $AL(\underline{\ell}, \widetilde{\ell})$ to the higher price $p' = AL(\underline{\ell}, \ell^*) - \varepsilon$ for all $\ell \in [\underline{\ell}, \widetilde{\ell}]$. We conclude that pooling equilibria in Corollary 1 do not satisfy the Nonnegative Profit Principle.

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, $P(\ell) \le \ell$ 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 4. We prove a more general result. Suppose that $P(\ell) = \ell$ for $\ell \in (\ell_1, \ell_2) \subseteq [\underline{\ell}, \widehat{\ell}]$. 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 α_P notation used in Proposition 1 because P is an identity map.) Conditions (1) and (2) must be satisfied in any equilibrium. Applying $P(\ell) = \ell$, conditions (1) and (2) become

$$[\ell' - C(\ell')]\alpha(\ell') \geq [\ell - C(\ell')]\alpha(\ell) \quad \text{and} \quad [\ell - C(\ell)]\alpha(\ell) \geq [\ell' - C(\ell)]\alpha(\ell'). \tag{7}$$

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

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

By symmetry, we have

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

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

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

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

$$[\ell - C(\ell)]\,\alpha'(\ell) + \left[1 - C'(\ell)\right]\alpha(\ell) = -C'(\ell)\alpha(\ell). \tag{8}$$

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

$$(\ell - C(\ell))\alpha'(\ell) + \alpha(\ell) = 0 \quad \text{or} \quad \frac{\mathrm{d}\ln\alpha(\ell)}{\mathrm{d}\ell} = -\frac{1}{\ell - C(\ell)}. \tag{9}$$

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

$$\ln\alpha(\ell) = -\int_{\ell_1}^{\ell} \frac{\mathrm{d}x}{x - C(x)} + K \quad \text{some constant } K$$

$$\alpha(\ell) = K' \exp\left\{-\int_{\ell_1}^{\ell} \frac{\mathrm{d}x}{x - C(x)}\right\} \quad \text{some constant } K'$$

$$\alpha(\ell) = \alpha(\ell_1)\exp\left\{-\int_{\ell_1}^{\ell} \frac{\mathrm{d}x}{x - C(x)}\right\}, \tag{10}$$

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

The separating equilibrium is the case when $\ell_1 = \underline{\ell}$, and $\ell_2 = \hat{\ell}$. 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 $\ell > \hat{\ell}$, 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 Corollary 2. The expert's expected profit of the separating equilibrium, Π^S, is

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

where the first equality follows from integration by parts and the definition of $\hat{\ell} = C(\hat{\ell})$, and where the last equality follows from (8).

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

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

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

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

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

Proof of Proposition 6. 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 $(\ell_2, c_2) \in [\underline{\ell}, \overline{\ell}] \times [\underline{c}, \overline{c}]$. 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

$$-\alpha(P(\ell_2, c_2)) \geq \frac{\Pi(c_2) - \Pi(c_1)}{(c_2 - c_1)} \geq -\alpha(P(\ell_1, c_1)), \quad \text{for } c_1 < c_2\text{, and for any } \ell_1 \text{ and } \ell_2. \tag{11}$$

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 (11). Let (ℓ_2, c_2) converge to (ℓ, c), and we have the right-hand derivative satisfying

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

Next, we have

$$-\alpha(P(\ell, c)) \geq \frac{\Pi(c) - \Pi(c_1)}{(c - c_1)} \geq -\alpha(P(\ell_1, c_1)),$$

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

where we have let c take the role of c_2 in (11). Let (ℓ_1, c_1) converge to (ℓ, c), and we have the left-hand derivative satisfying

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

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

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

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 $\widetilde{\ell}$ denote the loss that leads to the highest equilibrium price $P(\widetilde{\ell}, c)$. Because $\alpha(P(\ell, c))$ is constant in ℓ, the expert will deviate to recommending $P(\widetilde{\ell}, c)$ for all $\ell \neq \widetilde{\ell}$. We conclude that P must be independent of ℓ.

Appendix B. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.geb.2024.03.002.

References

Akerlof, G.A., 1970. The market for lemons: quality uncertainty and the market mechanism. Q. J. Econ. 84, 488–500.
Alger, I., Salanie, F., 2006. A theory of fraud and over-treatment in experts markets. J. Econ. Manag. Strategy 15, 853–881.
Bagwell, K., Riordan, M.H., 1991. High and declining prices signal product quality. Am. Econ. Rev. 81, 224–239.
Balafoutas, L., Beck, A., Kerschbamer, R., Sutter, M., 2013. What drives taxi drivers? A field experiment on fraud in a market for credence goods. Rev. Econ. Stud. 80, 876–891.
Balafoutas, L., Kerschbamer, R., 2020. Credence goods in the literature: what the past fifteen years have taught us about fraud, incentives, and the role of institutions. J. Behav. Exp. Finance 26, 100285.
Balafoutas, L., Fornwagner, H., Kerschbamer, R., Sutter, M., Tverdostup, M., 2021. Diagnostic uncertainty and insurance coverage in credence goods markets. Working paper.
Bester, H., Dahm, M., 2018. Credence goods, costly diagnosis, and subjective evaluation. Econ. J. 128, 1367–1394.
Chen, Y., Li, J., Zhang, J., 2022. Efficient liability in expert markets. Int. Econ. Rev. 63, 1717–1744.
Chiu, Y., Karni, E., 2021. Competitive equilibrium fraud in markets for credence-goods. J. Math. Econ. 96.
Darby, R.M., Karni, E., 1973. Free competition and the optimal amount of fraud. J. Law Econ. 16, 67–88.
Diamond, P., 1971. A model of price adjustment. J. Econ. Theory III, 156–168.
Dulleck, U., Kerschbamer, R., 2009. Experts vs. discounters: consumer free-riding and experts withholding advice in markets for credence goods. Int. J. Ind. Organ. 27, 15–23.
Dulleck, U., Kerschbamer, R., 2006. On doctors, mechanics, and computer specialists: the economics of credence goods. J. Econ. Lit. 44, 5–42.
Dulleck, U., Kerschbamer, R., Sutter, M., 2011. The economics of credence goods: on the role of liability, verifiability, reputation and competition. Am. Econ. Rev. 101, 526–555.
Emons, W., 1997. Credence goods and fraudulent experts. Rand J. Econ. 2, 107–119.
Fong, Y.F., 2005. When do experts cheat and whom do they target. Rand J. Econ. 36, 113–130.
Fong, Y.F., Liu, T., 2018. Liability and reputation in credence goods markets. Econ. Lett. 166, 35–39.
Fong, Y.F., Liu, T., Wright, D., 2014. On the role of verifiability and commitment in credence goods markets. Int. J. Ind. Organ. 37, 118–129.
Fong, Y.F., Liu, T., Meng, X., 2022. Trust building in credence goods markets. Am. Econ. J. Microecon. 14, 490–528.
Gottschalk, F., Mimra, W., Waibel, C., 2020. Health services as credence goods: a field experiment. Econ. J. 130, 1346–1383.
Karni, E., 2023. Preventive-service fraud in credence good markets. Econ. Theory. https://doi.org/10.1007/s00199-023-01540-x.
Kerschbamer, R., Sutter, M., Dulleck, U., 2017. How social preference shape incentives in (experimental) markets for credence goods. Econ. J. 127, 271–570.
Kim, K., 2012. Endogenous market segmentation for lemons. Rand J. Econ. 43, 562–576.
Liu, T., 2011. Credence goods markets with conscientious and selfish experts. Int. Econ. Rev. 52, 227–244.
Liu, T., Zhou, Y., Vergara Cobos, E., 2019. Pricing schemes and seller fraud: evidence from New York city taxi rides. J. Ind. Econ. 67, 56–90.
Lizzeri, A., 1999. Information revelation and certification intermediaries. Rand J. Econ. 30, 214–231.
Mahenc, P., 2017. Honest versus misleading certification. J. Econ. Manag. Strategy 26, 454–483.
McAfee, P., Schwartz, M., 1994. Opportunism in multilateral vertical contracting: nondiscrimination, exclusivity, and uniformity. Am. Econ. Rev. 84, 210–230.
Mimra, W., Rasch, A., Waibel, C., 2016. Price competition and reputation in credence goods markets: experimental evidence. Games Econ. Behav. 100, 337–352.
Pesendorfer, W., Wolinsky, A., 2003. Second opinions and price competition: inefficiency in the market for expert advice. Rev. Econ. Stud. 70, 417–437.
Pitchik, C., Schotter, A., 1987. Honesty in a model of strategic information transmission. Am. Econ. Rev. 77, 1032–1036.
Taylor, C., 1995. The economics of breakdowns, checkups, and cures. J. Polit. Econ. 103, 53–74.
Schneider, H., 2012. Agency problems and reputation in expert services: evidence from auto repair. J. Ind. Econ. 60, 406–433.
Wolinsky, A., 1983. Prices as signals of product quality. Rev. Econ. Stud. 50, 647–658.
Wolinsky, A., 1993. Competition in a market for informed experts' services. Rand J. Econ. 24, 380–398.
Wolinsky, A., 1995. Competition in markets for credence goods. J. Inst. Theor. Econ. 151, 117–131.