Deception, Exploitation and Lifespan of Buyer-Seller Relationship in Experience Goods Markets*

Selçuk Özyurt†

November 8, 2017

Abstract

A market where short-lived customers, i.e., one-time shoppers, interact with a long-lived expert is considered. The expert privately observes whether or not a particular treatment is necessary for his customers and has incentive to recommend the treatment even if it is unnecessary. When the expert is known to be opportunist, i.e., rational in the usual sense, his best equilibrium payoff can be achieved by being honest at all times. However, if the customers believe that the expert is a commitment type who recommends the treatment only when it is necessary, then the expert can build reputation as an honest seller and exploit his customers to achieve higher payoffs. Exploiting customers for a long period of time is extremely unlikely even if the seller's past actions are imperfectly observed. However, this behavior is a part of many equilibria when the expert's customer is also a long-lived agent. (JEL C72, C73, D82, D83)

*I thank Mehmet Ekmekçi, Drew Fudenberg, Andrew Postlewaite, Simon Board, and audiences at several conferences for helpful comments. I am grateful to the Scientific and Technological Research Council of Turkey (TUBITAK) grant 112K135 for financial support.

†Carnegie Mellon University Tepper School of Business (e-mail: sozyurt@andrew.cmu.edu).

1. Introduction

Many forms of consulting and advisory, medical, or repair services are prime examples of what is known as a credence or an experience good in the economics literature.[1] Generally speaking, these goods have the characteristics that customers can observe the utility they derive from the good *ex post*, i.e., upon consumption, but cannot be sure about the extent of the good they actually need *ex ante*.[2] Therefore, sellers act as experts who determine the customers' needs by performing a diagnosis. They can then provide the right quality and charge for it or exploit the information asymmetry by deceiving the customer.

Consider the following scenario to illustrate the motivation of the paper. Sam has recently bought an old house, requiring multiple projects that can be spread over time. Sam has neither the experience nor time for construction, home repair and maintenance, and so, she is very much dependent on a handyman or a professional contractor. In addition, her limited budget makes her extremely anxious about potential rip-off. Sam always had terrific experience shopping online. She has a dilemma whether she should follow the conventional wisdom, find a "reputable" handyman and develop a long-run relationship with him or go with her gut and go online (e.g., Yelps) to find a handyman every time she needs one.

Sam's concern is well-founded because deception and mistreatment of experts are confirmed by numerous empirical studies in experience and credence goods markets.[3] The

[1]Nelson (1970) first made the distinction between a search good for which quality is evident prior to purchase and an experience good for which quality is known only after consumption. The notion of a credence good for which quality may never be known is proposed by Darby and Karni (1973).

[2]There are two strands of literature on the credence goods (Dulleck, Kerschbamer, and Sutter, 2011). One strand, which is also the one that I use in this paper, takes the abovementioned characteristics—customers do not know what they need, but they observe the utility from what they consume. The other strand assumes that customers know what they need but observe neither what they consume nor the utility derived from what they consume (e.g., whether food has been produced organically or not).

[3]Emons (1997) cites a Swiss study reporting that an average person's probability of receiving one of seven major surgical interventions is one-third above that of a physician or a member of a physician's family. In the late 1970s, the Department of Transportation estimated that 53% of auto-repair charges represented unnecessary repairs (see Wolinsky, 1993 and 1995). More recently, a field experiment by Schneider (2012) shows that completely unnecessary repairs were present in 27% of the cases, and serious undertreatment occurred in 77% of the cases. He also estimates that agency problems in the U.S. auto-repair market generate a welfare loss of approximately $8.2 billion, or 22% of industry revenue. Levitt and Syverson (2012) claims that real estate agents have an incentive to convince clients to sell their houses too cheaply and too quickly. He reports that for two comparable houses, one owned by a real estate agent and the other owned by a client of the real estate agent, the home of the real estate agent will stay on the market for a longer period (an extra 9.5 days) and sell for a higher price (approximately 3.7%), and the greater the informational advantage of the real estate agent, the larger these two differences are. Considerable evidence that exists in the health-care industry also indicates that monetary incentives matter for the provision of credence goods. Gruber, Kim, and Mayzlin (1999), for example, show that the frequencies of cesarean deliveries compared with normal childbirths react to the fee differentials of

literature provides valuable insights regarding how market characteristics influence the level of misconduct in these markets.[4] However, it is quiet about the potential impacts of lifespan of buyer-seller relationships. The conventional wisdom is that long-run relationships are less prone to deception. One legitimate reason for this belief is possibly our general insight from the repeated games literature: that is, repeated interactions support cooperation, and so hinder deception. Bolton, Katok and Ockenfels (2004), Özalp, Chen and Zheng (2011), and Özalp, Zheng and Ren (2014) experimentally support this belief, though not for credence goods markets, by showing that repetitive trade improves efficiency and reduces deception relative to one-shot trade.[5] Although the comparative statics analysis of Bolton et al. (2004) is the closest to the one that I theoretically explore in this paper, their experimental design does not allow us to make a clear comparison.[6] Therefore, it is fair to say that we empirically and theoretically lack full-fledged understanding for the Sam's dilemma.

To provide some insights on how lifespan of buyer-seller relationship affects sellers' deception in credence goods markets, I consider two simple reduced-form models and compare the customers' worst equilibrium payoffs in each. A customer's trade decision relies solely on seller's advice and 'reputation', which is endogenously created either by (1) imperfect experiences of her fellow customers or (2) her own experience. More formally, I consider two long-lived (monopolist) experts, denoted by L and S.[7] Seller L has long-run relationship with his customer, and so repeatedly trades with the same customer indefinitely. On the other hand, Seller S has short-run relationship with his customers, and so trades with the same customer only once and draws a new one from a pool of

health insurance programs. See the surveys in McGuire (2000) and Gaynor and Vogt (2000).

[4]The literature on credence goods markets focuses on three types of fraudulent behavior on the sellers' side (Dulleck and Kerschbamer, 2006): *overtreatment* (providing an extensive, sophisticated treatment for a condition that requires only a limited treatment), *undertreatment* (providing a quality that is insufficient to satisfy the consumer's needs), and *overpricing* (providing an inexpensive treatment but charging for an expensive one). A seminal paper by Darby and Karni (1973) investigates how market conditions affect the equilibrium amount of mistreatment in credence goods markets. Wolinsky (1993) demonstrates how cheating can be eliminated when customers search for second opinions or experts have reputation concerns. Emons (1997, 2001) study how the price mechanism can discipline experts to practice honestly. Pesendorfer and Wolinsky (2003) study whether a competitive sampling of opinions makes it attractive for experts to provide costly but unobservable diagnostic effort. Alger and Salanie (2006) study under which conditions sellers deceive customers to keep them uninformed, as this deters them from seeking a better price elsewhere. Fong (2005) studies which customers the expert sellers deceive if the customers have heterogeneous and identifiable characteristics (e.g., valuations for treatments or costs of treatment). Dulleck and Kerschbamer (2006) provide an excellent literature survey on credence goods markets.

[5]Bolton et al. (2004) study repeated trust game and Özalp et al. (2011) and (2014) study forecast information sharing game in supply chain management.

[6]In their repetitive trading setup, the matching is fixed: two subjects are always paired together as buyer and seller. However, subjects always alternate roles from round to round. This experimental design provides subjects unrealistically strong incentives to cooperate, and hence a result that supports the conventional wisdom.

[7]The case with competitive experts is discussed in Section 2

infinitely many customers at each stage. Then I compare the customers' worst average equilibrium payoffs against seller S and L, which also correspond to payoffs of equilibria with the highest amount of exploitation.

The current model differs from the literature on credence goods in two important aspects. First, I model the buyer-seller interaction as a simple sender-receiver game, where the seller's payoff does not directly depend on his advice. This modeling choice eliminates all the complications that arise when we include, for example, the expert seller's quality and pricing decisions, which are not directly linked to the central inquiries of this paper. Second, I consider a repeated interaction between the seller and his customer(s), providing a fruitful platform to study the expert seller's long-run and short-run trade-offs for deception and trustworthiness. All results are presented for sufficiently patient players.

Results show that lifespan of a buyer-seller relationship does not necessarily make sellers more honest. In fact, contrary to the popular belief, Seller S whose customers are one-time shoppers can be much more honest. Clearly, the last argument is not always true. It is true when there is a (monitoring) technology where customers can transfer their experience to other customers very much in the spirit of the feedback technologies in online feedback/shopping websites, such as Yelp, TripAdvisor, Amazon, eBay, or Alibaba. Seller S can exploit his customers more as the quality of the monitoring technology worsens, yet he is less exploitative than Seller L in the sense that maximum occurrence of "exploitation" by Seller S in any equilibrium is less than that of Seller L. Therefore, on seller-optimal equilibria, where the sellers' payoffs are maximized, Seller S offers a greater surplus to his customers on average.

Furthermore, Seller S has no incentive to be dishonest if he is known to be a rational agent. The reason for this is that the short-lived customers are so alert against the expert's advice that they never trust him unless he tells the truth with a sufficiently high probability. Because the expert must play a mixed strategy, i.e., tell the truth with a positive probability, if he wants to deceive his customers, his payoff from deception can be no more than his payoff of telling the truth. On the other hand, when customers believe that S is a commitment type who recommends the treatment only when it is necessary, then Seller S can manipulate his short-lived customers for higher payoffs by building reputation as an honest seller and recommending the treatment–once his reputation exceeds some threshold level—when it is unnecessary.[8] Customers can detect the seller's deception with a higher probability as their monitoring technology gets "stronger," and so stronger monitoring is always better for the customers. Its effect on the seller's welfare is not that

[8] Brown and Minor (2012) empirically support this prediction and show that more experienced experts are significantly more likely to mislead their customers.

straight. There is a threshold level of reputation under which the expert seller favors a stronger monitoring technology because he can build his reputation faster. Once the seller's reputation is above this threshold, he prefers a weaker monitoring technology.

However, Seller L can deceive his long-lived customer for an unlimited period of time, and so offer her a payoff of zero in a seller-optimal equilibrium. This is true because the expert can credibly promise his customer that he will report truthfully for a sufficiently long period of time if she does not punish his occasional deceptions. The customer does not deviate from such an 'exploitative' contract as long as her continuation payoff is no less than her outside option, which is zero. The expert seller's long-term promises have no benefit to short-lived customers, and so the expert cannot exploit them as much as he does his long-lived customer.

The rest of the paper is organized as follows. Section 2 explains the details of the sender-receiver stage game and provides its equilibrium predictions. Section 3 discusses the repeated sender-receiver game with short-lived receivers and presents the main results. Section 4 considers the repeated sender-receiver game where both the sender and receiver are long-lived agents. Finally, Section 5 concludes and discusses the related literature.

2. The Sender-Receiver Stage Game

In this section, I introduce and examine the sender-receiver stage game where an expert (he) and a receiver (she) interact sequentially. Sections 3 and 4 investigate the infinitely repeated version of this game. The receiver (e.g., a house owner) seeks service of an expert (a handyman) who provides a particular treatment. The expert can correctly diagnose whether or not his treatment is necessary for the receiver, but the receiver is unsure about it and cannot verify the expert's recommendation.

Timing of the game is as follows: at the beginning, the nature determines whether the treatment is "necessary" (n) or "unnecessary" (u), and so the true state is $s \in S = \{n, u\}$. The treatment is unnecessary with probability $\pi \in (0, 1)$. The expert observes the true state and then sends an unverifiable message to the receiver $m \in M = S$.[9] After observing the expert's message, the receiver who does not know the true state either approves (a) or rejects (r) the treatment. Regardless of the true state, the expert's payoff is $v_e > 0$ if the receiver approves the treatment and 0 otherwise. The receiver's payoff of approving the treatment is positive if the treatment is necessary and negative if it is unnecessary.

[9]The main focus of the paper is to study how long/much the sender can exploit the receiver as a function of his relationship with her. Therefore, I do not explore how vague or coarse the expert's messages would get, and thus, a richer state or message spaces will complicate the setup with no apparent benefit.

Receiver's payoff of rejecting the treatment is normalized to 0. Assuming a negative payoff of rejection when the treatment is necessary would have no effect on our results.[10] The payoffs are summarized as follows:

a	r
$v_e, -v_u$	$0, 0$

$s = u$

a	r
v_e, v_n	$0, 0$

$s = n$

where the real numbers v_e, v_u, and v_n are all strictly positive. The case where expected return of the treatment is positive is uninteresting because the receiver would always approve the treatment even if the expert's message is not informative. For this reason, I suppose that the parameters satisfy

$$-\pi v_u + (1 - \pi)v_n < 0 \tag{1}$$

so that the receiver prefers to reject the treatment ex ante.

Furthermore, the expert is one of two types: *honest* or *opportunist*. The opportunist expert is a rational player in the usual sense. That is, he chooses his message, given his beliefs about the receiver's play, to maximize his expected payoff. However, the honest expert always tells the truth. The expert knows his type, and let $\mu \in [0, 1)$ denote the probability that the expert is honest, that is μ is the expert's initial reputation. Call this sender-receiver game where all parameters are common knowledge G.

For expositional simplicity and clarity I took competition out of the model and supposed that the expert seller is a monopolist. This assumption ensures that the customer's minimax value is exogenous, which is zero according to the stage game G. Therefore, our results throughout the paper would be consistent with a competitive market setup where the expert seller is competing with other sellers in the sense of Diamond (1971): that is, customers have very small but non-zero search costs, and thus conjecture to receive the same expected payoff from all the other sellers, ensuring that the customer's minimax payoff can be fixed to zero. There might exist other interesting equilibria in that setup, but a complete analysis deserves to be a separate paper.

Let $\sigma_e(s) \in [0, 1]$ denote the probability that the expert sends message u when he observes state s. Therefore, $1 - \sigma_e(s)$ is the probability that the expert sends message n in state s. Given the expert's strategies $\sigma_e(u)$ and $\sigma_e(n)$, the receiver updates her belief

[10] This is true, in particular for the repeated game, because I will suppose that the expert's recommendation may become a public information only if the receiver trades with him, and thus, the ability of learning the true state when the treatment is refused does not affect the public history and equilibrium.

about the true state according to the Bayes' rule. Let $P(s|m)$ indicate the receiver's posterior probability that the true state is s conditional on the event that the expert sends message $m \in \{n, u\}$. Thus we have, for example,

$$P(u|n) = \frac{\pi(1-\mu)[1-\sigma_e(u)]}{\pi(1-\mu)[1-\sigma_e(u)]+(1-\pi)\mu+(1-\pi)(1-\mu)[1-\sigma_e(n)]}.$$

The receiver's mixed strategy $\sigma_r(m)$ is a function of the message $m \in S$ she receives, indicating the probability of playing a when she observes message m. A strategy profile and a belief structure constitute a *Perfect Bayesian Equilibrium* (or simply *equilibrium*) if (1) each player's strategy specifies optimal actions, given his/her beliefs and the strategies of the other player, (2) given the strategy profile, the beliefs are consistent with Bayes' rule whenever possible.

Definition 1. An equilibrium is **persuasive** if the receiver approves the treatment with a positive probability. An equilibrium is **fully revealing** if the (opportunist) expert truthfully reports the state with certainty, that is, he recommends the treatment if and only if it is necessary. An equilibrium is **babbling** if the (opportunist) expert's strategy is independent of the true state and the receiver's strategy is independent of the expert's message.

Remark 1. By optimality, a fully-revealing equilibrium must be persuasive. However, a persuasive equilibrium is not necessarily fully revealing. In fact, as the next result shows, deception occurs only in equilibria that are persuasive but not fully revealing.

Proposition 1. *There exists a persuasive equilibrium of the sender-receiver game G, where the opportunist expert always recommends the treatment and the receiver approves only if the expert recommends it, if and only if* $\mu \geq \mu^* \equiv 1 - \frac{(1-\pi)v_n}{\pi v_u}$*. It is unique whenever it exits. Furthermore, there does not exist a fully-revealing equilibrium of the game G.*

There are two important implications of this result. First, in a persuasive equilibrium the expert never lies when the treatment is necessary, and so, the receiver will certainly reject the project whenever she observes message u. However, the expert has incentive to lie when the treatment is unnecessary. Therefore, by expert's ***deception*** I formally mean a situation where the expert recommends the treatment although it is unnecessary and the receiver approves it. Second, the threshold μ^* indicates the *minimum level of trust the expert must possess to deceive the receiver.* If the expert's reputation is below this level, then he cannot deceive the receiver unless he tells the truth, i.e., reveals the state, with a sufficiently high probability. However, in game G, revealing the true state with a positive probability when the treatment is unnecessary is never optimal for the expert, and so he will never play a mixed strategy in equilibrium.

It is rather easy to see why there is no fully-revealing equilibrium. Suppose there is. Then the receiver's best response is to approve the treatment only when the expert recommends it. However, given the receiver's best response strategy, the expert prefers to lie and recommend the treatment also when it is unnecessary. Finally, babbling equilibrium exists only when $\mu < \mu^*$. This is true because for higher values of μ, the receiver approves the treatment when the expert recommends it and rejects otherwise. Therefore, the receiver's equilibrium strategy depends on the message she observes. Thus, the persuasive equilibrium is the unique equilibrium of the stage game G when $\mu \geq \mu^*$.

Remark 2. A babbling equilibrium exists if and only if $\mu < \mu^*$.

3. The Repeated Sender-Receiver Game with Short-Lived Receivers

This section considers an infinitely repeated game where a long-lived expert seller repeatedly plays the stage game G against a succession of agents (short-lived customers), each of whom plays the game once. Before the game starts, nature moves first and determines the type of the sender, which is fixed throughout the game. Then the expert privately learns his type: he is honest with probability $\mu \in [0, 1)$ and opportunistic with probability $1-\mu$. At each stage $t = 0, 1, ...$, nature determines the true state $s \in S$, where $\pi \in (0, 1)$ is the probability of state being u. Then the expert privately learns the state and sends a message $m \in M$. After observing the expert's message, the receiver decides whether to approve or reject the treatment. At the end of the stage, the expert and the receiver obtain their stage game payoffs. The receiver does not learn the true state if she rejects the treatment. The payoff structure of the stage game is as given in Section 2.

The opportunist expert discounts future with $\delta < 1$ and his objective is to maximize his discounted lifetime payoffs. The receiver is an infinite sequence of different short-lived agents who play the stage game with the expert only once. Each short-lived agent's objective is to maximize her expected payoff in the stage game that she plays.

The expert can perfectly observe the entire history of the game. All short-lived agents observe the same public signal $y_t \in Y = \{\emptyset, u_r(a^t, s^t)\}$ at the end of stage t. By $y^t = \emptyset$, I mean that the receivers observe no information about stage t and $u_r(a^t, s^t)$ is the receiver's payoff at stage t, which is a function of her action, a^t, and the true state, s^t. One can interpret $y^t = \emptyset$ case as though short-lived agent who plays the game with the expert in stage t does not share her experience with the other agents. Note that the receivers do not observe the expert seller's previous advices, but can correctly infer them by the previous receivers' payoffs, except in cases when receivers reject the treatment. As a result of, the monitoring technology is such that receivers leave their feedbacks only when they trade with the expert as is the case in most of the online shopping websites.

The receivers' monitoring technology is imperfect in the sense that for any t, $y^t = \emptyset$ with probability $1 - \beta$ and

$$y^t = u_r(a^t, s^t) = \begin{cases} -v_u, & \text{if } (a^t, s^t) = (a, u) \\ v_n, & \text{if } (a^t, s^t) = (a, n) \\ 0, & \text{otherwise.} \end{cases}$$

with probability β. Thus, the term $\beta \in [0, 1]$ is the rate at which a receiver's experience with the expert becomes public information. Higher values of β ensures that the receivers will be better informed about the expert's past play. Call this repeated sender-receiver game where all parameters are common knowledge G^∞.[11]

Let $\mathcal{H}_e = \bigcup_{t=0}^{\infty} (A \times S \times M)^t$ denote the set of histories for the expert. Therefore, a behavioral strategy of the opportunist expert is $\sigma_e : \mathcal{H}_e \times S \to \Delta(M)$. Given any history $h^t = (y^0, ..., y^{t-1})$ (possibly a null history h^0), $\sigma_e(h^t, s)(m)$ denotes the probability that the opportunist expert sends message m given that he observes state s after history h^t. The strategy of the honest expert is simple: he reports the true state at any stage. A behavioral strategy for the receiver is $\sigma_r : \mathcal{H}_r \times M \to \Delta(A)$ where $\mathcal{H}_r = \bigcup_{t=0}^{\infty} Y^t$ denotes the set of public histories for the receivers. Let $\sigma_r(h^t, m)(a)$ denote the probability that the receiver approves the treatment given that she observes message m after history h^t. If $\{a^t\}_{t=0}^{\infty}$ is the sequence of actions taken by the receivers and if $\{s^t\}_{t=0}^{\infty}$ is the sequence of state realizations throughout the game, then the expert's average payoff is $(1-\delta)\Sigma_{t=0}^{\infty} \delta^t u_e(a^t, s^t)$, a short-lived receiver's payoff is $u_r(a^t, s^t)$ if she enters the game in stage t, and the receivers' average payoff is $(1-\delta)\Sigma_{t=0}^{\infty} \delta^t u_r(a^t, s^t)$.

Definition 2. An equilibrium of the repeated sender-receiver game G^∞ is *fully revealing* if, on the equilibrium path, the (opportunist) expert reports truthfully at all stages. It is *babbling* if, on the equilibrium path, the (opportunist) expert's strategies are independent of the true state and the receiver's strategies are independent of the expert's messages at all stages. Finally, an equilibrium of the repeated sender-receiver game G^∞ is *persuasive*

[11] This infinite-horizon game is similar to the one in Ely and Välimäki (2003). The fundamental differences of their model are that the expert seller can offer one of two possible treatments and consumers only observe which treatment were done in the past. Only the expert seller knows which treatment was necessary. Therefore, the receivers have no feedback mechanism, where their experiences become public information, and so the seller's reputation can be built only upon the frequency with which various repairs are performed.

if, on the equilibrium path, a receiver approves the project with a positive probability at some stage, and *fully persuasive* if each receiver approves the project with a positive probability at all stages.

The results of Fudenberg and Levine (1989) do not apply here because their focus was to find the long-run player's lowest equilibrium payoff. However, I am interested in short-run players' minimum average equilibrium payoff when the expert's exploitation (and payoff) is maximal.

Given the payoff structure of the repeated game G^∞, it is fairly easy to see that the expert's best equilibrium must be a persuasive equilibrium, whenever one exists. Fully-revealing equilibrium is fully-persuasive, and as the next result shows, it always exist for sufficiently patient expert. However, there might exist other fully-persuasive equilibria that yield higher payoff to the expert than the fully-revealing equilibrium. Existence of such equilibria automatically implies that the expert must be deceiving the receivers at some stages, i.e., the expert suggests the treatment even though it is not necessary and the receiver accepts it. It is clear that the receivers' average equilibrium payoff gets lower as the expert deceives more of his customers. Intuition also suggests that the more customers the expert deceives the higher payoff he will get. However, the next result states that this is not entirely true.

The Main Results

Suppose for now that $\beta = 0$, that is the short-lived receivers cannot observe the history of the repeated sender-receiver game G^∞. This case resembles situations where the short-lived agents can never learn the benefit of the treatment as is true for some credence goods. There cannot exist a fully-revealing equilibrium of the repeated sender-receiver game because the short-lived receivers cannot coordinate on the expert's past play. In addition, the expert cannot build up or loose his reputation in the game. Therefore, if the expert's initial reputation is small, i.e., $\mu < \mu^*$, then there exists no persuasive equilibrium, in which a receiver approves the treatment. On the other hand, if $\mu \geq \mu^*$, then there exists a unique equilibrium of the game G^∞ in which each receiver approves the treatment and the expert deceives the receivers and recommends the treatment at all times. Clearly, this is the best case scenario for the expert seller and the worst case scenario for the customers.

For the rest of this section, I will investigate the more interesting case where the receivers can (partially) observe the expert's past play.

Proposition 2. *For any $\mu \in [0, 1)$ and $\beta \in (0, 1]$ there is some $\delta_\beta \in (0, 1)$ such that for all $\delta > \delta_\beta$ there exists a fully-revealing and fully-persuasive equilibrium of the game G^∞. Furthermore, when the expert is known to be the opportunistic type, i.e., $\mu = 0$,*

1. *deception is consistent with equilibrium: there exists a persuasive equilibrium where the expert deceives the receivers and each receiver's (ex ante) payoff is zero,*

2. *but the fully-revealing equilibrium, yielding payoff $v_e^f \equiv \frac{(1-\pi)v_e}{(1-\delta)}$ to the expert and average payoff $(1 - \pi)v_n$ to the receivers, is the best equilibrium for both the expert and the receivers, and*

3. *the expert's worst equilibrium is a babbling equilibrium with payoff zero.*

For any positive β, the repetition of the babbling equilibrium of the stage game is an equilibrium of the repeated game G^∞ given that the expert's reputation is less than μ^*, and players' payoffs are zero in this equilibrium. However, there exist other equilibria and the expert can achieve higher payoffs. If the expert is sufficiently patient, then mutual trust between the short-lived receivers and the opportunist expert supports an equilibrium where the expert truthfully reports at all stages. A punishment strategy that supports this equilibrium is simple: if the expert deviates and deceives a receiver, i.e., suggests the treatment when it is unnecessary, and if this deviation is observed by the short-lived receivers, then the expert and all the subsequent receivers play their babbling equilibrium strategies for the rest of the game.

When the expert is known to be the opportunist type, deception is consistent with equilibrium, but in the form of mixed strategies: that is, the expert must tell the truth with a sufficiently high probability to convince the receivers that the treatment is necessary. [12] By deceiving the receivers in mixed strategies, the expert can reduce their expected payoffs all the way to zero, but this would not benefit the expert. This is true because a mixed equilibrium requires that the expert's continuation payoff of deceiving (the discounted sum of the payoffs of deception and subsequent play) and telling the truth must be the same, and so the expert cannot do better than reporting truthfully. Hence, both the expert and the receivers prefer the fully-revealing equilibrium to all other equilibria. One may naturally expect the fully-revealing equilibrium to be the one that should survive in a competitive market environment when the expert is known to be the opportunistic type.

[12] In equilibrium, the expert must guarantee at least zero (ex ante) payoff to the receivers. Otherwise, they can simply reject the treatment. Because the expert is known to be the opportunist type, the receivers' expected payoff of accepting the treatment can be greater than zero only if the expert tells the truth with a sufficiently high probability.

An important message we get from Proposition 2 is that the expert's deception will always hurt the receivers but not necessarily benefit himself. Therefore, it is called ***exploitation*** or ***deliberate deception*** if deception benefits the expert. The second important message is that the expert can exploit a receiver only if he plays a pure strategy where he recommends the treatment with certainty. However, it means the opportunistic expert's messages are uninformative when he exploits the receivers. In equilibrium, a receiver approves the treatment even though the opportunistic expert's messages are uninformative only if she believes that the recommendation is coming from the honest type rather than the opportunistic type. Therefore, the expert can exploit the receivers only if he has sufficiently high reputation of being the honest type.

Recall from the stage game that the threshold reputation level, i.e., $\mu^* = 1 - \frac{(1-\pi)v_n}{\pi v_u}$, makes the (short-lived) receiver indifferent between accepting and rejecting the treatment when the sender recommends the treatment, and so plays a significant role. The expert can deceive the receivers without building his reputation any further if $\mu \geq \mu^*$. Otherwise, the expert first needs to build up his reputation to be able to exploit the receivers.

If the receivers' conjecture is such that the expert strictly prefers to tell the truth at stage t, then observing the expert telling the truth at that stage does not change the receivers' belief about the expert's actual type; he simply does what he was expected to do. However, observing the expert telling the truth even though he strictly prefers to lie changes the receivers' belief about the type of the expert. But this observation also proves that the receivers' conjecture was wrong, and "equilibrium" dictates that the receivers must have right conjectures to begin with. Therefore, in equilibrium, the expert can build his reputation for honesty if the receivers have the right conjecture that the expert has incentives to lie and to tell the truth, i.e., he is indifferent between lying and telling the truth. Thus, the expert chooses to materialize his short-term incentives as expected if he lies. If the expert tells the truth instead, then he chooses to postpone his short-term gains for something higher in return, which is a higher reputation for honesty.

Now suppose that the expert's initial reputation is low, i.e., $\mu < \mu^*$. Our ultimate purpose in next result is to find an equilibrium in which the expert's payoff is the highest, i.e., the expert's exploitation is maximal. Thus, we must seek a fully-persuasive equilibrium where the expert tells the truth 1) with certainty if the treatment is necessary and 2) with a positive probability if the treatment is unnecessary, which is (i) low enough so that his reputation can be updated quickly, but (ii) high enough so that the receivers approve the treatment when the expert recommends it. The reason why the first condition must hold is obvious.

The second condition must hold because the expert discounts time, and so his payoff is higher if he builds his reputation faster and starts exploiting the receivers earlier. The expert can build his reputation in just one stage if his strategy dictates him to tell the truth with a very low probability and if that stage payoffs are publicly observed. More formally, in stage 1, if the expert's strategy dictates him to tell the truth when the treatment is unnecessary with probability σ_e, then his reputation at the end of stage 1 following the history h^1, where treatment was unnecessary and expert was truthful, is $\mu_1 = \frac{\mu}{\mu+(1-\mu)\sigma_e}$ by the Bayes' rule.[13] The expert can deceive the receiver in the second stage only if his updated reputation is higher than μ^*. Thus, the expert can update his reputation to the required level μ^* in only one stage if the following holds:

$$\sigma_e \leq \frac{(1-\mu^*)\mu}{\mu^*(1-\mu)}. \tag{2}$$

As for the third condition, the expert prefers to build his reputation gradually and I formally prove this claim in Appendix (proof of Proposition 3). To get the intuition for this consider the case where the expert's initial reputation, μ, is very low. Then the condition in (2) requires that the expert is expected to lie with a very high probability when the treatment is unnecessary, and so the receiver in the first stage will reject the treatment even though the treatment is necessary and the expert truthfully reports so. The expert's continuation payoff in this case will be at most $\delta V(\mu)$, where $V(\mu)$ is the expert's highest continuation payoff in a subgame where his reputation is μ.[14] However, if the expert gradually builds his reputation, i.e., his strategy dictates him to tell the truth in stage 1 with a sufficiently high probability, then his continuation payoff will be at least $(1-\pi)v_e + \delta V(\mu)$, which is strictly higher than the payoff he could achieve if he builds his reputation in one stage.[15]

[13] Note that if stage 1 payoffs become public information, then the receivers will only observe that the treatment was rejected. The receivers will reason, given the expert's and stage 1 receiver's strategies, that this outcome could occur only when the true state was u and the expert was truthful, and thus update their beliefs about the expert's type.

[14] With probability $(1-\pi)$ true state is n. The expert reports truthfully and the receiver rejects the treatment. The next stage is the same subgame where the seller's reputation is μ. Thus, the expert's payoff is $0+\delta V(\mu)$ when state is n. If the state is u, then the expert's continuation payoff of telling the truth and of lying must be the same, which is $\widehat{V} = 0+\delta[\beta V(0)+(1-\beta)V(\mu)]$ where $V(0)$ is the expert's continuation payoff when he get caught lying, i.e., his rationality is revealed. Because the expert can achieve at least as high as $V(0)$ when his reputation is $\mu > 0$, we have $V(0) \leq V(\mu)$, and thus $\widehat{V} \leq \delta V(\mu)$.

[15] If state is n, then the expert reports truthfully and the receiver approves the treatment. Because the next stage is the same subgame where the seller's reputation is μ, the expert's payoff is $v_e + \delta V(\mu)$ when state is n. If state is u, then the expert's continuation payoff is at least as high as his continuation payoff when he reports truthfully, which is $\widehat{V} = 0+\delta[\beta V(\mu^1)+(1-\beta)V(\mu)]$ where $\mu^1 > \mu$ is the expert's updated reputation. Therefore, $\delta V(\mu) \leq \widehat{V}$ because $V(\mu) \leq V(\mu^1)$. Hence, the expert's highest payoff is at least $(1-\pi)v_e + \delta V(\mu)$ under a strategy where he gradually builds his reputation.

Thus, the expert should be telling the truth with a sufficiently high probability to ensure that the receivers approve the treatment whenever he recommends it. When the expert recommends the treatment, the receiver's expected payoff of approving it is

$$EU_r(a,n) = -v_u \frac{\pi(1-\mu)(1-\sigma_e)}{\pi(1-\mu)(1-\sigma_e)+(1-\pi)} + v_n \frac{(1-\pi)}{\pi(1-\mu)(1-\sigma_e)+(1-\pi)}$$

and rejecting it is $EU_r(r,n) = 0$. Hence, the receiver approves the treatment if $EU_r(a,n) \geq 0$, or equivalently

$$\sigma_e \geq 1 - \frac{v_n(1-\pi)}{v_u\pi(1-\mu)} = \frac{\mu^* - \mu}{1-\mu} \tag{3}$$

holds.

For some values of the primitives, in particular $(\mu^*)^2 \leq \mu$, the inequalities (2) and (3) can hold simultaneously. In this case, reputation building would take only one stage and the expert's best equilibrium dictates that he must tell the truth when the treatment is unnecessary with probability $\sigma_e = \frac{\mu^*-\mu}{1-\mu}$ until his honesty is publicly observed. However, if $(\mu^*)^2 > \mu$, then the inequalities (2) and (3) do not hold simultaneously. In this case, the expert should build up his reputation gradually in more than one stage.

To calculate the shortest publicly observed time required for the expert to build his reputation gradually up to the critical level μ^*, I set $\mu_0 = \mu$ and for all $t \geq 0$ define $\sigma_e^t = \frac{\mu^*-\mu_t}{1-\mu_t}$ and $\mu_{t+1} = \frac{\mu_t}{\mu_t+(1-\mu_t)\sigma_e^t}$ recursively. The term σ_e^t represents the probability that the expert reports truthfully when treatment is unnecessary conditional on the event that he previously did that and publicly observed in t times, and μ_t represents the expert's updated reputation given that the receivers observed t-many previous rejections of the treatment because the treatment was unnecessary.

If the expert observes the state u for k times and tells the truth at all times according to σ_e^t's as given above, then his reputation reaches $\mu_k = \frac{\mu}{\mu+(1-\mu)\Pi_{t=1}^{k-1}\sigma_e^t}$. The expert will stop building up his reputation whenever $\mu_k \geq \mu^*$ holds, which is equivalent to $\Pi_{t=1}^{k-1}\sigma_e^t \leq \frac{(1-\mu^*)\mu}{\mu^*(1-\mu)}$. Hence, the shortest time required for the expert to build up his reputation to μ^*—while the receivers prefer to approve the treatment conditional on observing the message n—is defined by

$$K^* = \min\left\{ k \in \mathbb{Z}^+ \,\middle|\, \prod_{t=0}^{k-1} \sigma_e^t \leq \frac{(1-\mu^*)\mu}{\mu^*(1-\mu)} \right\}. \tag{4}$$

By using this definition of K^*, it is rather easier to show that K^* is the smallest of the natural numbers k, satisfying $(\mu^*)^{k+1} \leq \mu$, and I show this last step in the appendix.

Lemma 1. *Suppose that $\mu > 0$ and $\beta \in (0, 1]$. In a fully-persuasive equilibrium, the shortest publicly observable time period that is required for the expert to build up his reputation to μ^* is*

$$N_G = \begin{cases} 0 & \text{if } \mu \geq \mu^* \\ K^* & \text{otherwise,} \end{cases}$$

where K^* is the smallest positive integer satisfying $(\mu^*)^{K^*+1} \leq \mu$, that is

$$K^* = \min \left\{ k \in \mathbb{Z}^+ \middle| \frac{\ln \mu}{\ln \mu^*} - 1 \leq k \right\}.$$

Note that K^* increases with μ^* but decreases with μ. Therefore, if the expert's initial reputation, μ, is higher, then he needs less time to build his reputation as an honest seller. If the level of trust the expert needs to possess in order to deceive the receivers, i.e., μ^*, is higher, then the expert needs more time to build his reputation. Recall that μ^* is an increasing function of v_u and π and a decreasing function of v_n. Therefore, *if the expected return of the treatment is higher, i.e., closer to 0, then μ^* gets lower, and thus, the expert needs less time to build up reputation.* The last result in this section provides the expert's highest equilibrium payoff when he has the change of building reputation.

Proposition 3. *Suppose that $\beta \in (0, 1]$ and $\mu > 0$. For sufficiently high values of δ, the expert's best equilibrium payoff, where the receivers are exploited the most, is*

$$V_e = (1 - \alpha_\beta)\ v_e^f + \alpha_\beta\ v_e^d$$

where $v_e^d \equiv v_e \left(\frac{1+\frac{\pi\delta\beta(1-\pi)}{1-\delta}}{1-(1-\pi\beta)\delta} \right)$ and $\alpha_\beta \equiv \left[\frac{\pi\beta\delta}{1-(1-\pi\beta)\delta} \right]^{N_G}$. In this equilibrium, the receivers' average ex ante payoff is no less than $v_n \left(\frac{\pi\delta\beta(1-\pi)}{1-(1-\pi\beta)\delta} \right) \delta^{N_G}$.

It is important to note that in the expert seller's best equilibrium, where exploitation is the highest, the receivers get a positive expected payoff on average. As I will show in next section, this is not the case when the receiver is a long-run player.

The term α_β is always in $(0, 1]$, and so, the expert's best equilibrium payoff is a convex combination of two numbers; v_e^f and v_e^d. The term v_e^f is the expert's payoff in the fully-revealing equilibrium. We know from Proposition 2 that v_e^f is the expert's best equilibrium payoff when he is known to be opportunist. On the other hand, the term v_e^d is the expert's best equilibrium payoff when the expert's initial reputation is higher than the threshold level μ^*, in which case he does not need to build his reputation to exploit the receivers, i.e., N_G is 0.

The equilibrium strategies yielding the highest payoff have three possible layers. The expert and the receivers start the game in the reputation building phase: the expert recommends the treatment with certainty if it is necessary and tells the truth with a positive probability if it is unnecessary. This phase lasts N_G publicly observed stages. The players move to the deception phase immediately once the expert's reputation exceeds the threshold level μ^*. In the deception phase, the opportunist expert recommends the treatment regardless of the true state. Deception phase ends whenever the receivers observe the expert's lie. Then players move to the truthful reporting phase and stay there forever. In this phase, the expert and the receivers play their fully-revealing equilibrium strategies. Any publicly observed deviation from these strategies are punished by babbling equilibrium strategies. The next two results are corollaries to Proposition 3, and their proofs are deferred to appendix.

Corollary 1. *The function V_e, indicating the expert's best equilibrium payoff, is maximized when*

$$\beta = \frac{N_G(1-\delta)}{\delta\pi}.$$

This result implies several testable hypothesis. A patient expert prefers weaker monitoring because the profit maximizing monitoring technology, β, decreases with δ. If the likelihood of the unnecessary treatment is higher, then the expert prefers a weaker monitoring technology because that reduces the likelihood of getting caught when he deceits. If the receiver's expected return from the treatment increases, i.e., μ^* is lower, and thus N_G is lower, then the expert prefers a weaker monitoring technology. Finally, higher initial reputation, μ, reduces N_G, and thus β. Therefore, the expert with a higher initial reputation reaches the threshold level of reputation faster, and thus prefers a weaker monitoring technology. Put differently, the expert with a low initial reputation prefers stronger monitoring system.

In fact, if the expert's initial reputation is higher than the critical threshold level, μ^*, then the expert can achieve his highest payoff under no monitoring, i.e., $\beta = 0$. However, when the expert needs to build up his reputation to deceive the receivers, i.e., $\mu < \mu^*$, then no monitoring is not in his best interest. The expert prefers to be monitored perfectly while he builds up his reputation. Once he reaches the threshold level, the expert prefers not to be monitored by the receivers at all. Overall, there is an inverted U-shaped relationship between the expert's best equilibrium payoff and the strength of the monitoring technology.

Corollary 2. *In the expert's best equilibrium, the expected number of stages that the expert exploits the receivers is $1/\beta$.*

Expected length of exploitation depends only on the receivers' monitoring technology. Consistent with the intuition, the length of exploitation decreases with the strength of the monitoring technology. However, the relationship has a degree of -1. The next section shows that the length of exploitation would be significantly longer when both the receiver and the expert are long-lived agents.

4. A Benchmark Result with a Long-Lived Receiver

In this section, I provide a benchmark result, not a complete analysis, for the case where the receiver is also a long-lived agent, and show that the expert can exploit the receiver for an unlimited period of time, leaving zero payoff to the receiver. For this purpose, I suppose that both the expert and the receiver are long-lived agents with the common discount factor $\delta < 1$. I restrict my attention to the case where the players can perfectly monitor the history of the game and the expert is known to be the opportunist type. That is, the players' actions and the true state at all stages are observable by the players and $\mu = 0$. The reason for this restriction is the intuition that the expert can deceive the receiver for a longer period of time when the monitoring is not perfect or when the expert has reputation for honesty.

The repeated sender-receiver game with a long-lived receiver has multiple equilibria. For example, there exists a fully-revealing and fully-persuasive equilibrium with sufficiently patient players, which delivers the receiver's best equilibrium payoff. On the other extreme, the next result proves that the expert can exploit the receiver for an unlimited period of time even though he is known to be the opportunist type and his deception is perfectly observable by the receiver.

Proposition 4. *For sufficiently high values of $\delta < 1$, there exists an equilibrium of the repeated sender-receiver game with a long-lived receiver in which the expert exploits the receiver for an unbounded period of time and delivers zero payoff to the receiver.*

As is standard in repeated games, equilibrium with indefinite period of exploitation is a result of mutual trust. The expert deceives the receiver for one period, for instance, and then reports truthfully for a several periods. As long as the expert balances the ratio of deception and truthful reporting in a way that the receiver's continuation payoff is never negative, the receiver neither deviates from her strategies of approving the treatment nor punishes the expert's deceptions. Thus, the cycle of exploitation and truthful reporting, which relies on mutual trust, can be repeated indefinitely. Such a cycle of mutual trust is not sustainable between the expert and his short-lived receivers. This is true because the short-lived receivers do not care about any future "reward" the expert may offer to their successors.

The expert might do better than this kind of cycling strategies by, for example, exploiting the receiver at the very early periods of the game. The important take-away, however, is that the expert can deceive the receiver in ways that leave no (or very little) surplus to the receiver and benefit from that.

5. Concluding Remarks and Related Literature

Following the seminal study by Crawford and Sobel (1982), sender-receiver (or cheap talk) games have became a natural framework to study issues of information transmission between an informed sender (expert) and an uninformed decision maker (receiver). Sobel (2011) provides a very detailed survey of this literature. Unlike the usual treatment in the sender-receiver games literature, I consider an infinite-horizon repeated game. Aumann and Hart (2003) and Krishna and Morgan (2004), for example, consider dynamic sender-receiver games. However, only the talk is repeated in their settings. Golosov, Skreta, Tsyvinsky, and Wilson (2009) study strategic information transmission game in a finite-horizon, dynamic Crawford and Sobel setup. The main result of their study is that fully-revealing equilibrium exists when both the expert and the receiver are long-lived and fully patient players.

Sobel (1985) considers a reputational sender-receiver game where the talk between the expert and the receiver is repeated finitely many stages, and both players are fully patient and long lived. His study assumes that the receiver is uncertain about the bias of the expert—she is either the "friendly" type, whose preferences are perfectly aligned with the receiver, or the "enemy" type, who has completely opposed preferences to the receiver. The main result of his study is that deception is sustainable in equilibrium only if the expert has sufficiently high reputation of being the "friendly" type, and deception would occur only once. Ottaviani and Sorensen (2006, 2006b) and Morris (2001) also investigate reputational sender-receiver games where the expert's bias is unknown to the receiver. The main message of these studies is that truth telling is incompatible with equilibrium when the expert is sufficiently concerned about his reputation. Unlike these former works, the expert's bias is not fixed—but state contingent—in the current model.

Benabou and Laroque (1992) study an infinitely repeated sender-receiver game between an expert and multiple audiences (i.e., the public). However, the players in their model have significantly different incentives. The expert receives an informative signal about the true state of the world. Unlike the current model, the expert in Benabou and

Laroque (1992) plays a trading game with his audiences right after sending his public message and directly affects his own payoff. In particular, the expert is an insider who manipulates his audiences' opinion through his cheap talk messages and trades with them in a purely speculative market. Benabou and Laroque (1992) show that the expert will deceive his audiences (who are short-lived agents) during an unbounded length of time. They conclude, contrary to the results of the present study, that an expert with very low reputation for honesty will make no significant attempt to build his reputation, and the issue of whether intermediate reputations are worth improving by "investing in truth" remains unresolved.

The standard results in the repeated games literature do not apply to our case, and there are two main reasons for this. First, the prevalent objective in the literature is to characterize players' equilibrium payoffs without any focus on strategies that lead to these payoffs. However, my primary objective is to find the equilibrium strategies where the expert is most exploitative. That is, I search an equilibrium in simple strategy profiles, in the sense of Abreu (1988), where the total number of deliberate deception is highest. Naturally, there are multiple such simple strategy profiles and some of them automatically yield the players' best/worst equilibrium payoffs. Second, the sender-receiver stage game does not satisfy the full dimensionality condition, i.e., the set of feasible and individually rational payoff set has empty interior in $\mathbb{R}^2$, and thus the machineries developed by Abreu, Pearce and Stacchetti (1990) and Fudenberg, Levine and Maskin (1994) do not obtain in our repeated game with two long-lived agents. Furthermore, our repeated sender-receiver game with one long-lived and one short-lived player is not moral hazard mixing game, and thus the machineries developed by Fudenberg and Levine (1994) do not obtain either.

Nevertheless, our predictions are consistent with the well-known results of the infinitely repeated games literature. Folk theorems simply show that for sufficiently patient players, cooperation is consistent with rational behavior, but many other behavioral profiles are equally rational. For this reason, it is not surprising to support deception in equilibrium as long as the seller keeps the buyer's continuation payoff no less than her minmax payoff. Somewhat surprisingly, we can show that for any positive natural number $T \leq \infty$, there exists an equilibrium where the total number of deception is T. Fudenberg, Kreps and Maskin (1990) show that in an environment where all players are known to be rational, deceiving a short-lived buyer cannot be consistent with equilibrium because short-run players always play short-run best responses. Therefore, it may come as no big surprise to show that deceiving short-lived buyers is not rewarding for the seller regardless of how imperfect the feedback technology is. What is interesting is that the short-lived buyers' propensity to believe that the seller is a commitment type is more effective to give the seller strong incentive for deception than imperfect monitoring technology. Furthermore, it is interesting that we can provide a nice upper bound for the total (expected) number of exploitation in equilibrium as a function of the efficiency of the monitoring technology.

Appendix

Proof of Proposition 1. The proof for the non-existence of a fully-revealing equilibrium is provided in the main text. Next, I prove that if there exists a persuasive equilibrium, then we must have (1) $\mu \geq \mu^*$ and (2) the persuasive equilibrium has the following form (and so it is unique): the opportunist expert sends the message n irrespective of the true state and the receiver approves (respectively rejects) the treatment when she observes the message n (respectively u). Given the receiver's best response correspondences, it is fairly easy to verify that this strategy profile is a PBE of the game G when $\mu \geq \mu^*$, and so, I skip the proof of the "if" part.

Suppose that there exists a PBE strategy profile σ in which after some message realization the receiver plays a with positive probability. Given the expert's strategies $\sigma_e(u)$, $\sigma_e(n)$, the receiver's best response correspondences (conditional on observing message u) are calculated as follows. The receiver's expected payoff of playing a and r are given by

$$EU_r(a|u) = -v_u\left(\frac{\pi\mu+\pi(1-\mu)\sigma_e(u)}{\pi\mu+\pi(1-\mu)\sigma_e(u)+(1-\pi)(1-\mu)\sigma_e(n)}\right) + v_n\left(\frac{(1-\pi)(1-\mu)\sigma_e(n)}{\pi\mu+\pi(1-\mu)\sigma_e(u)+(1-\pi)(1-\mu)\sigma_e(n)}\right)$$
$$EU_r(r|u) = 0.$$

Therefore,

$$BR_r(\sigma_e|u) = \begin{cases} 1, & \text{if } \sigma_e(u) < \overbrace{\dfrac{v_n(1-\pi)(1-\mu)\sigma_e(n) - v_u\pi\mu}{v_u\pi(1-\mu)}}^{A} \\ [0,1], & \text{if } \sigma_e(u) = A \\ 0, & \text{otherwise.} \end{cases}$$

That is, the receiver plays a (conditional on observing message u) with certainty whenever $\sigma_e(u) < A$ and r with certainty if $\sigma_e(u) > A$. Now, suppose that the receiver observes the message n. Then,

$$EU_r(a|n) = \frac{-v_u\left[\pi(1-\mu)[1-\sigma_e(u)]\right] + v_n\left[(1-\pi)\mu+(1-\pi)(1-\mu)[1-\sigma_e(n)]\right]}{\pi(1-\mu)[1-\sigma_e(u)]+(1-\pi)\mu+(1-\pi)(1-\mu)[1-\sigma_e(n)]}$$
$$EU_r(r|n) = 0.$$

Therefore,

$$BR_r(\sigma_e|n) = \begin{cases} 1, & \text{if } \sigma_e(u) > \overbrace{\dfrac{v_u\pi(1-\mu) - v_n(1-\pi)\mu - v_n(1-\pi)(1-\mu)[1-\sigma_e(n)]}{v_u\pi(1-\mu)}}^{B} \\ [0,1], & \text{if } \sigma_e(u) = B \\ 0, & \text{otherwise.} \end{cases}$$

That is, the receiver plays a (conditional on observing message n) with certainty whenever $\sigma_e(u) > B$ and r if $\sigma_e(u) < B$. Note that we have $A < B$ for all $\sigma_e(n) \in [0,1]$ because the inequality (1) holds. Recall that σ is such that the receiver plays a after some message realization with positive probability. There are five exhaustive cases regarding the value of $\sigma_e(u)$ relative to A and B:

(*i*) Assume that $\sigma_e(u) < A < B$. Then the receiver plays a and r when she receives the messages u and n, respectively. Therefore, the expert's payoffs of sending messages u and n are v_e and 0, respectively. Then, the optimality of equilibrium implies that the expert will send message u regardless of the true state, that is $\sigma_e(u) = 1$ and $\sigma_e(n) = 1$. However, we have $B = 1 - \frac{v_n(1-\pi)\mu}{v_u\pi(1-\mu)}$ when $\sigma_e(n) = 1$, contradicting the initial assumption $\sigma_e(u) < B$.

(*ii*) Assume that $A = \sigma_e(u) < B$. Then the receiver is indifferent between a and r when the expert sends the message u. However, the receiver chooses r when the expert sends the message n. Because the receiver plays a with positive probability in σ, we must have that the receiver plays a with positive probability after observing message u. Therefore, the expert's payoffs of sending messages u is positive whereas his payoff of sending message n is 0. Then, once again, the optimality of equilibrium implies that the expert will send message u regardless of the true state, that is $\sigma_e(u) = 1$ and $\sigma_e(n) = 1$. Similar to the previous case, we reach a contradiction because $\sigma_e(u) = 1 < B$ and B is strictly less than 1 for $\sigma_e(n) = 1$.

(*iii*) Assume that $A < \sigma_e(u) < B$. Then the receiver plays r regardless of the expert's message, contradicting that the receiver plays a with positive probability in σ.

(*iv*) Assume that $A < \sigma_e(u) = B$. Then the receiver plays r when the expert sends the message u, and she is indifferent between a and r when the expert sends the message n. Because the receiver plays a with positive probability in σ, we must have that the receiver plays a with positive probability after observing message n. Therefore, the expert's payoffs of sending messages n is positive whereas his payoff of sending message u is 0. Then, the optimality of equilibrium implies that the expert will send message n regardless of the true state, that is $\sigma_e(u) = 0$ and $\sigma_e(n) = 0$. For these values of $\sigma_e(u)$ and $\sigma_e(n)$, we have $A < 0$ and $B = 1 - \frac{v_n(1-\pi)}{v_u\pi(1-\mu)} = 0$ if and only if $\mu = \mu^*$.

(v) Assume that $A < B < \sigma_e(u)$. Then the receiver plays r when the expert sends the message u and a otherwise. Therefore, the expert's payoffs of sending messages n and u are v_e and 0, respectively. Then, the optimality of equilibrium implies that the expert will send message n regardless of the true state, that is $\sigma_e(u) = 0$ and $\sigma_e(n) = 0$. For these values of $\sigma_e(u)$ and $\sigma_e(n)$, we have $A < 0$ and $B = 1 - \frac{v_n(1-\pi)}{v_u\pi(1-\mu)} < 0$ if and only if $\mu > \mu^*$.

These five cases prove that in a persuasive equilibrium, the expert must send message n with certainty regardless of the true state (i.e., $\sigma_e(u) = \sigma_e(n) = 0$ must hold) and the expert's reputation must be no less than μ^* (see (iv) and (v)).

Proof of Proposition 2. Consider the following strategy profile. There are two phases; the coordination and the punishment. In the coordination phase, the opportunist expert truthfully reports the true state and the receivers approve the treatment when they observe the message n and reject it if they observe u. In the punishment phase, the expert always sends message n and the receivers reject the treatment regardless of the message they observe. The repeated game starts with the coordination phase and the players stay in this phase unless the expert deviates. Once the expert deviates and the receivers observe this deviation, then the game moves to the punishment phase and stays there for the rest of the game.

Next, I show that this strategy is a PBE. In the punishment phase, it will be public knowledge that the expert is the opportunist type with certainty and the punishment phase strategies forms a babbling equilibrium: a short-lived receiver will never approve the treatment because the expert's messages are uninformative and the expert cannot improve his payoff either. In the coordination phase, the expert's continuation payoff is

$$\frac{(1-\pi)v_e\delta}{1-\delta} = 0 + (0\pi + (1-\pi)v_e)\delta + (0\pi + (1-\pi)v_e)\delta^2 + \ldots$$

if the expert observes state u, and is $v_e + \frac{(1-\pi)v_e\delta}{1-\delta}$ if the expert observes state n. Therefore, the expert's payoff in the coordination phase is

$$\left[(1-\pi)v_e + \frac{(1-\pi)v_e\delta}{1-\delta}\right] = \frac{(1-\pi)v_e}{(1-\delta)} \equiv v_e^f.$$

One can easily calculate that the receivers' average payoff is $(1-\pi)v_n$. The expert has no incentive to deviate when the state is n. However, if he deviates when the state is u, then his deviation will lead to the continuation payoff of $v_e + \delta(1-\beta)v_e^f$. Therefore, the expert does not deviate from his coordination phase strategies if and only if $\frac{(1-\pi)v_e\delta}{1-\delta} \geq v_e + \delta(1-\beta)v_e^f$, or equivalently $\delta \geq \delta_\beta \equiv \frac{1}{1+(1-\pi)\beta}$.

Finally, given the expert's strategy, a short-lived receiver's expected payoff of following her strategy is 0 when the true state is u and v_n when the true state is n. However, if she deviates, her payoff will be $-v_u$ when the true state is u and 0 when the true state is n. Thus, coordination phase strategies are also optimal for each short-lived receiver.

Next, I will show that deception does not increase the expert's payoff. Let h^t be a history (possibly the null history) in which the expert is known to be the opportunist type. First, I will show that there is no equilibrium of the continuation game following the history h^t in which the receiver approves the treatment with positive probabilities after observing both messages. Suppose for a contradiction that there exists an equilibrium in which $\sigma_r(h^t, m)(a) \in (0,1)$ for each $m \in M$. Recall the receiver's best response correspondences from the proof of Proposition 1. Conditional on observing the message u, the receiver approves the treatment if $\sigma_e(u) \leq \sigma_e(n)\frac{v_n(1-\pi)}{v_u\pi} := A$. Note that the strategy $\sigma_e(m)$ in the one stage game corresponds to $\sigma_e(h^t, m)(u)$ in the repeated sender-receiver game. However, when the receiver observes message n, she approves the treatment if $\sigma_e(u) \geq 1 - (1-\sigma_e(n))\frac{v_n(1-\pi)}{v_u\pi} := B$. Therefore, the receiver approves the treatment with a positive probability regardless of the message she observes if and only if $B \leq \sigma_e(u) \leq A$ holds. However, since the inequality (1) holds, we have $A < B$ for all values of $\sigma_e(n)$, that yields the desired contradiction.

Therefore, if the expert wants to make receiver approve the treatment even when the true state is u, he must send the message n regardless of the true state. However, full deception (sending message n with certainty when true state is u) is not consistent with equilibrium. If the receiver believes that the expert will deceive her with certainty at some stage after h^t, then the short-lived receiver prefers to reject the treatment at that stage. However, partial deception (sending message n with probability less than one when true state is u) would be consistent with equilibrium. In particular, we know from the receiver's best response correspondences that if the expert sends the message n with certainty when the true state is n and with a probability $\sigma_e(h^t, u)(n) \leq \frac{v_n(1-\pi)}{v_u\pi}$ when the true state is u, then the receiver prefers to approve the treatment if she observes the message n. Next, I will show that partial deception does not improve the expert's payoff.

Let V be the expert's highest continuation payoff in any equilibrium following a history h^t in which the expert is known to be the opportunist type. The expert's expected payoff if the true state is n in stage $t+1$ is at most $v_e + \delta V$. However, if the true state is u in stage $t+1$, then the expert's continuation payoff of telling the truth is no more than $0 + \delta V$. In equilibrium where the expert tells the truth with a positive probability when the true state is u, his continuation payoff of lying and telling the truth must be the same. Hence, his continuation payoff when the state is u should be no more than δV. Thus, V must be less than or equal to $(1-\pi)(v_e + \delta V) + \pi\delta V$, implying that $V \leq \frac{(1-\pi)v_e}{1-\delta}$. Hence, v_e^f is the upper bound for the expert's equilibrium payoff following the history h^t.

Finally, I would like to argue (rather informally) that (partial) deception is consistent with equilibrium. I will support my claim for $\beta = 1$, and similar but more detailed strategies will exist for positive values of β. Consider the following strategy profile. There are two modes of the game: deception and babbling. Game starts in the deception phase. In this phase, the expert sends message n as long as he observes n. If he observes u, then he sends n with probability $\sigma_r(h^t, u)(n) = \frac{v_n(1-\pi)}{v_u\pi}$ and u with the remaining probability. Receivers always accept the treatment when the expert sends message n and reject otherwise. In the babbling mode, everyone plays his/her babbling equilibrium strategies. As long as the expert reports truthfully, the game stays in the deception mode. If he lies when the true state is u, then the game stays in the deception mode with probability $p = 1 - \frac{(1-\delta)^2}{(1-\pi)\delta}$ and moves to the babbling mode with the remaining probability. If the players ever deviate from their strategies, the game also moves to the babbling mode. Once the game moves to the babbling mode, it stays there forever. It is rather easy to verify that this strategy forms an equilibrium when δ is selected carefully and delivers zero payoff to all receivers.

Proof of Lemma 1. Suppose that $\mu \in (0, \mu^*)$. We know that $\mu^* = 1 - \frac{v_n(1-\pi)}{v_u\pi}$, and $\sigma_e^t = 1 - \frac{v_n(1-\pi)}{v_u\pi(1-\mu_t)}$. Therefore, we can write

$$\sigma_e^t = \frac{\mu^* - \mu_t}{1 - \mu_t}. \tag{5}$$

Moreover, we know that $\mu_{t+1} = \frac{\mu_t}{\mu_t + (1-\mu_t)\sigma_e^t}$ with $\mu_0 = \mu$, and thus, the recursive structure implies that

$$\mu_t = \frac{\mu}{\mu + (1-\mu)(\sigma_e^0\sigma_e^1...\sigma_e^{t-1})}. \tag{6}$$

The equations in (5) and (6) imply that

$$\sigma_e^t = \mu^* - \frac{\mu(1-\mu^*)}{(1-\mu)(\sigma_e^0\sigma_e^1...\sigma_e^{t-1})}. \tag{7}$$

Thus, given the starting point $\sigma_e^0 = \frac{\mu^* - \mu}{1-\mu}$ and the equation (7) we can recursively calculate $\sigma_e^0\sigma_e^1...\sigma_e^{k-1}$ for any $k \geq 1$ as follows: First $\sigma_e^0\sigma_e^1 = \frac{(\mu^*)^2-\mu}{1-\mu}$. Using the last equation and (7) we find that $\sigma_e^0\sigma_e^1\sigma_e^2 = \frac{(\mu^*)^3-\mu}{1-\mu}$. Repeating this process yields

$$\Pi_{t=0}^{n-1}\sigma_e^t = \frac{(\mu^*)^k - \mu}{1-\mu}$$

By using the definition of K^* in (4), it is rather easier to show that it is the smallest of the natural numbers k satisfying $(\mu^*)^{k+1} \leq \mu$, which is equivalent to $k \geq \frac{\ln \mu}{\ln \mu^*} - 1$.

Proof of Proposition 3. In what follows, I will describe some strategies and show that they are equilibrium strategies of the repeated game G^∞ that yield the expert the highest possible payoff. While I describe some (parts) of these strategies, I use a public randomization device simply because the description of these strategies are much shorter and easier. Fudenberg and Maskin (1991) show that one can actually get rid of the public randomization device for sufficiently high δ, by appropriate choice of which periods to play each action profile involved. For the same reason, we can also get rid of the public randomization device and prove the next result. I will start with presenting two lemmas that will be helpful for the proof of Proposition 3.

Lemma 2. *Consider a history h^t of the repeated sender-receiver game G^∞ in which the expert is known to be the opportunist type. For any payoff v in the range $[0, v_e^f]$ and sufficiently large values of $\delta < 1$, there exists an equilibrium of the continuation game following the history h^t in which the expert's payoff is v.*

Proof. We know from Proposition 2 that a fully-revealing and fully-persuasive equilibrium exists for the continuation game following the history h^t, where the expert's payoff is v_e^f. Likewise, a babbling equilibrium also exists for this subgame, where the receivers reject the treatment regardless of the message and the expert's payoff is 0. Let $v = \tau v_e^f$ and $\tau \in (0,1)$. Consider the following strategy profile: At the beginning of each stage, before the expert observes the true state, both the expert and the receivers observe the outcome of a public randomization device that has two possible outcomes; A and B. Outcome A occurs with probability τ and outcome B occurs with probability $1-\tau$. When the outcome is A in stage t, both players play their fully-revealing equilibrium strategies in this stage, where the expert tells the truth and the receiver approves (or rejects) the treatment when she observes the message n (or u). However, when the outcome is B, then players play their babbling equilibrium strategies, where the expert always sends the message n and the receivers always reject the treatment. If one of the players deviates and if the receivers publicly observe this deviation, then both players move to the punishment phase, where the expert and the receiver play their babbling equilibrium strategies forever.

According to this strategy profile, the expert's payoff is $\frac{\tau(1-\pi)v_e}{1-\delta}$, which is equal to τv_e^f as required. Similar to the arguments in the proof of Proposition 2, the punishment phase strategies form a babbling equilibrium. Because the expert is known to be the opportunist type, the receiver has no incentive to approve the treatment and the expert has no incentive to tell the truth when the public randomization outcome is B.

However, if the outcome is A, then the expert has no incentive to deviate when the true state is n. If he deviates in state u and sends the message n, then his continuation payoff will be $v_e + \frac{\delta\tau(1-\beta)(1-\pi)v_e}{1-\delta}$. If he tells the truth, his continuation payoff is $\frac{\delta\tau(1-\pi)v_e}{1-\delta}$. Therefore, the expert does not deviate from his strategy if and only $\frac{\delta\tau(1-\pi)v_e}{1-\delta} \geq v_e + \frac{\delta(1-\beta)\tau(1-\pi)v_e}{1-\delta}$, or equivalently $\delta \geq \delta^* \equiv \frac{1}{1+\tau\beta(1-\pi)} \in (0,1)$. This completes the proof.

□

Lemma 3. *Suppose that $\mu \geq \mu^*$ and $\beta \in (0,1]$. The expert's best equilibrium payoff in the repeated sender-receiver game G^∞ is $v_e^d = v_e \left(\frac{1+\frac{\pi\delta\beta(1-\pi)}{1-\delta}}{1-(1-\pi\beta)\delta} \right)$.*

Proof. First, I will show that the payoff v_e^d can be supported in equilibrium. Then I will argue that it is the highest expected payoff that the expert can achieve in any PBE of the repeated sender-receiver game. There are three phases of the strategy profile; exploitation, truthful-reporting and punishment. Players start in the exploitation phase, where the expert sends the message n regardless of the true state until he gets caught lying. Once he gets caught, the players move to the truthful-reporting phase, where the expert always tells the truth, and stay there forever. In the exploitation and truthful-reporting phases, the receivers always approve the treatment when they observe the message n and reject the treatment otherwise. If the expert ever deviates in any phase and if his deviation is detected by the receivers, then the players move to the punishment phase and stay there forever, where the receivers always reject the treatment and the expert always sends the message n. If a receiver deviates in the exploitation phase, then the players will continue to stay in the exploitation phase. However, if a receiver deviates in any other phase, then the players move to the punishment phase and stay there forever.

The expert's payoff under this strategy profile is calculated by solving the following recursive equation

$$V = \pi\Big(v_e + \delta\left[(1-\beta)V + \beta v_e^f\right]\Big) + (1-\pi)\Big(v_e + \delta V\Big)$$

The first parenthesis is the expert's continuation payoff when the true state is u. The expert will deceive the receiver in the first stage. Following the exploitation, the expert will get caught lying with probability β, and so, will be truthful forever. In this case, his continuation payoff will be v_e^f. However, the expert will not get caught lying with probability $1-\beta$, in which case the continuation game is identical with the game itself. The second parenthesis is the expert's expected payoff when the true state is n. The expert does not deceive the receiver in the first stage, and so, the continuation game will be identical to the game itself. The solution of this equation for V yields the value for v_e^d.

As we already argued before, punishment phase is a babbling equilibrium of the game. In the truthful reporting phase, the receivers know that the expert is opportunist. Thus, as we argued in Proposition 2, the players' strategies in the truthful reporting phase are optimal as well. Finally, in the exploitation phase, if the expert deviates and sends the message u, then his payoff will simply be 0. We know from Proposition 1 that each short-lived receiver's expected payoff of approving the treatment is positive because $\mu \geq \mu^*$. Therefore, if a receiver deviates in the exploitation phase and rejects the treatment when she observes n, then she will get 0 payoff. Hence, exploitation phase strategies are also optimal for both players.

Finally, I will argue that v_e^d is the highest payoff the expert can achieve in any equilibrium of the repeated sender-receiver game G^∞. Recall the strategy that is described above. The receiver approves the treatment if the true state is n. Moreover, she also approves the treatment when the true state is u if the expert's type has not revealed to the receivers yet. When the expert is get caught lying, the expert's type will be revealed. Thus, his best equilibrium payoff will be v_e^f as we proved in Proposition 2. Therefore, a higher game payoff for the expert is possible only if the expert sends the message n in the exploitation phase with a probability less than 1 when the true state is u, and thus, reduces the chances that she gets caught lying. This case (i.e., partial deception) would not increase the expert's payoff as we previously argued. In particular, in case of partial deception, the expert randomizes over the messages in M at some stage where the true state is u if and only if he is indifferent between these messages at that stage. When the true state is u, the expert's continuation payoff is at most $v_e + \delta\left[(1-\beta)V + \beta\frac{v_e(1-\pi)}{1-\delta}\right]$ if he lies (i.e., sends the message n), and thus his continuation payoff must also be at most this much when he tells the truth.

Therefore, if V is the highest continuation payoff of the expert in any equilibrium following a history where the expert's reputation is higher than or equal to μ^*, then $V \leq (1-\pi)(v_e + \delta V) + \pi\left(v_e + \delta\left[(1-\beta)V + \beta\frac{v_e(1-\pi)}{1-\delta}\right]\right)$, implying that $V \leq v_e\left[\frac{1+\frac{\pi\delta\beta(1-\pi)}{1-\delta}}{1-(1-\pi\beta)\delta}\right]$. This completes the proof.

□

Now I will start the proof of Proposition 3. First, I will show that the payoff V_e can be supported in equilibrium. Then I will argue that it is the expert's highest equilibrium payoff in the game. Consider the following strategy profile σ:

(i) The expert always sends message n whenever the true state is n.

(ii) Reputation building phase: Let h^t be a history where (1) the expert's reputation at the beginning of stage $t+1$ (i.e., $\mu(h^t)$) is strictly positive, but strictly less than μ^*, and (2) none of the players' deviation from his/her prescribed strategies has publicly observed by the receivers. In stage $t+1$, following the history h^t, (1) the expert sends message u with probability $\sigma_e(h^t, u)(u) = 1 - \frac{v_n(1-\pi)}{v_u \pi(1-\mu(h^t))}$ when the true state is u, and (2) the receivers reject (approve) the treatment if they observe the message u (n).

(iii) Let h^t be a history where (1) none of the players' deviation from his/her prescribed strategies has publicly observed by the receivers, (2) the true state in stage t was u, but (3) the expert has lied in stage t and his true type has revealed to the receivers (i.e., $\mu(h^t) = 0$). In the continuation game following the history h^t, players move to public randomization phase and stay there forever. In this phase, the expert and the receivers play the strategies that are described in the proof of Lemma 2. The public randomization device produce the outcome A with probability $\tau_k \equiv \frac{(\delta\beta V(k) - v_e)(1-\delta)}{\delta\beta(1-\pi)v_e}$, where

$$V(k) = v_e^f \left(1 - \left[\frac{\pi\beta\delta}{1-(1-\pi\beta)\delta}\right]^{N_G - k}\right) + v_e^d \left[\frac{\pi\beta\delta}{1-(1-\pi\beta)\delta}\right]^{N_G - k}$$

and $k \leq N_G$ is the number of stages (including stage t) that the receivers publicly observed in the history h^t, where the true state was u. Thus, in the public randomization phase (1) the expert always tells the truth and the receivers always approve (reject) the treatment if the message is n (u) when the outcome is A, and (2) the expert and the receivers play their babbling equilibrium strategies when the outcome is B.

(iv) Let h^t be a history where (1) the expert's reputation $\mu(h^t)$ is (weakly) higher than μ^*, and (2) none of the players' deviation from his/her prescribed strategies has publicly observed by the receivers. In the continuation game following the history h^t, all the players move to exploitation phase and stay there until the expert gets caught lying. In the exploitation phase, (1) the expert always sends the message n (regardless of the true state) and (2) the receivers reject (approve) the treatment if they observe the message u (n). Once the expert gets caught lying, all the players move to truthful reporting phase and stay there forever. In the truthful reporting phase, the expert always tells the truth, the receivers reject the treatment if the message is u and approve it if the message is n.

(v) After any history h^t where the expert has deviated from his strategies and this deviation is publicly observed by the receivers, players move to the punishment phase and stay there forever.[16] In the punishment phase, the receivers always reject the treatment, and the expert always sends the message n. If a receiver deviates in the exploitation phase, then the players will continue to stay in the exploitation phase. However, if a receiver deviates in any other phase, and if the receivers publicly observe her deviation, then the players move to the punishment phase and stay there forever.

First, note that, for any $\beta \in (0,1]$, there exists some $\delta_\beta < 1$ such that $\tau_k \in (0,1)$ for all $\delta \geq \delta_\beta$. τ_k is positive because $v_e < \delta\beta v_e^f < \delta\beta V(k)$ for high values of δ. Likewise, $\tau_k < 1$ because $V(k) < v_e^d$ and $\delta\beta v_e^d < v_e + \frac{\delta\beta(1-\pi)v_e}{1-\delta}$, implying $(\delta\beta V(k) - v_e)(1-\delta) < \delta\beta(1-\pi)v_e$ and $\frac{(\delta\beta V(k)-v_e)(1-\delta)}{\delta\beta(1-\pi)v_e} < 1$. Now, I will show that $\delta\beta v_e^d < v_e + \frac{\delta\beta(1-\pi)v_e}{1-\delta}$ holds: Suppose for a contradiction that it does not. That is, $\delta\beta v_e^d \geq v_e + \frac{\delta\beta(1-\pi)v_e}{1-\delta}$. If we insert the value of $v_e^d = v_e\left(\frac{1-\delta+\pi\delta\beta(1-\pi)}{(1-\delta)(1-(1-\pi\beta)\delta)}\right)$ to this inequality and cancel the v_e's, then we get $\delta\beta\left(\frac{1-\delta+\pi\delta\beta(1-\pi)}{(1-\delta)(1-(1-\pi\beta)\delta)}\right) \geq 1 + \frac{\delta\beta(1-\pi)}{1-\delta}$. Multiply both sides of this inequality with $1-\delta$ and divide by $\delta\beta$, and then subtract $\frac{\delta\pi(1-\pi)}{1-(1-\pi)\delta}$ from both sides to get $\frac{1-\delta}{1-(1-\pi\beta)\delta} \geq \frac{(1-\pi)(1-\delta)}{1-(1-\pi\beta)\delta} + \frac{1-\delta}{\delta\beta}$. Dividing both sides by $1-\delta$ and rearranging the terms yield $\frac{1}{\delta\beta} \leq \frac{\pi}{1-(1-\pi\beta)\delta}$. The last inequality implies $1 - \delta + \delta\pi\beta \leq \delta\pi\beta$, which is equivalent to $1 \leq \delta$. The last inequality yields the desired contradiction.

Let the term $V(k)$, where $0 \leq k \leq N_G$, represents the expert's continuation payoff, following a history h^t where (1) none of the players' deviation from his/her prescribed strategies has publicly observed by the receivers, and (2) the true state was u in exactly k stages prior to time t (including stage t), and these stages were publicly observed by the receivers. Therefore, $V(k)$ is the expert's continuation payoff when his public reputation is $\mu_k = \frac{\mu_{k-1}}{\mu_{k-1}+(1-\mu_{k-1})\sigma_e(h^y,u)(u)}$, where $y \leq t$ is the latest publicly observed stage in which the true state was u. Thus,

$$
\begin{aligned}
V(0) &= (1-\pi)[v_e + \delta V(0)] + \pi\delta[\beta V(1) + (1-\beta)V(0)] \\
V(1) &= (1-\pi)[v_e + \delta V(1)] + \pi\delta[\beta V(2) + (1-\beta)V(1)] \\
&\dots \\
V(N_G - 1) &= (1-\pi)[v_e + \delta V(N_G - 1)] + \pi\delta[\beta V(N_G) + (1-\beta)V(N_G - 1)]
\end{aligned}
$$

[16] Note that if the receivers ever realize that the expert has sent message u during the reputation building or exploitation phases and then later deceived a receiver, which is an off-equilibrium path, then they will deduce that the expert mush have been be the opportunist type and deviated.

and $V(N_G) = v_e^d$.

Recall that following the history h^t, where the true state was u in exactly k stages that are publicly observed by the receivers, the expert's reputation is μ_k. The first part of the recursive equation of $V(k)$ (i.e., $[v_e + \delta V(k)]$) is the expert's continuation payoff conditional on the event that the true state is n in stage $t+1$. In this case, the expert tells the truth in stage $t+1$, and so his stage payoff is v_e. The continuation game following stage $t+1$ will be identical to the continuation game following the history h^t. Thus, the expert's reputation is still μ_k, and so, the expert's continuation payoff following stage $t+1$ is $V(k)$. The second part of the recursive equation $\delta[\beta V(k+1) + (1-\beta)V(k)]$ (or δv_e^d in case $k = N_G - 1$) indicates the expert's continuation payoff if the true state is u in stage $t+1$. The expert tells the truth and his stage game payoff is 0. With probability β the receivers observe stage $t+1$, and so the expert's reputation will be updated to μ_{k+1}. Therefore, the expert's continuation payoff will be $V(k+1)$. However, with probability $1-\beta$ the receivers do not observe stage $t+1$, and thus the expert's continuation game payoff will be $V(k)$. If $k = N_G - 1$, then the expert's reputation will reach a level above μ^* in stage $t+1$ if this stage is publicly observed by the receivers, and so the expert's continuation payoff, according to the strategies given in (iv), will be v_e^d (the expert's highest equilibrium payoff in the exploitation phase), as we prove in Lemma 3.

The expert's expected payoff in the repeated sender-receiver game is equal to $V(0)$. In order to find its value, we must solve these N_G equations recursively. First solve the last equation, which implies $V(N_G - 1) = \left[1 - \frac{\pi\beta\delta}{1-(1-\pi\beta)\delta}\right] v_e^f + \left(\frac{\pi\beta\delta}{1-(1-\pi\beta)\delta}\right) v_e^d$. At the end of this process we get

$$V(0) = \left[1 - \left(\frac{\pi\beta\delta}{1-(1-\pi\beta)\delta}\right)^{N_G}\right] v_e^f + \left(\frac{\pi\beta\delta}{1-(1-\pi\beta)\delta}\right)^{N_G} v_e^d,$$

which is equal to V_e as desired.

Next, I will show that the strategy profile σ forms a PBE of the repeated sender-receiver game. We already know that the punishment strategies in part (v) are optimal. Lemma 3 shows that the strategies in part (iv) are optimal as well. By Lemma 2, the strategies in part (iii) are optimal because the expert's continuation payoff after a history h^t that fits to the description in (iii) is $\tau_k \frac{(1-\pi)v_e}{1-\delta}$ (i.e., $\tau_k v_e^f$), and $\tau_k < 1$ as we proved above.

We have also shown that the strategies in (i) are optimal when the expert's reputation is zero or above the threshold μ^*. To check that they are optimal during the reputation building phase, consider a history h^t where no deviation has ever publicly observed, true state in stage t is n, and previously k publicly observed stages were u. If the expert tells the truth in stage t, then his continuation payoff is $v_e + \delta V(k)$. However, if the expert reports u in stage t, then the receiver will reject the treatment and the expert's reputation will be updated, and so his payoff will be $0 + \delta[\beta V(k+1) + (1-\beta)V(k)]$. Given that $V(k+1) = [V(k)(1-\delta(1-\pi\beta)) - (1-\pi)v_e]/\delta\pi\beta$, it is easy to check that the latter is strictly less than the former.

Therefore, all we need to show that the strategies in (ii) are optimal as well. Consider a history h^t where (1) the expert's reputation at the end of stage t (i.e., $\mu(h^t)$) is strictly positive and strictly lower than μ^*, and (2) none of the players' deviation has publicly observed by the receivers before. The expert has no incentive to deviate when the true state is n in stage $t+1$. Therefore, suppose that the true state in stage $t+1$ is u. Furthermore, suppose that $1 \leq k \leq N_G$ is the number of stages (including stage t) in the history h^t, where the true state was u and the experts publicly observed the previous k stage outcomes. The expert's continuation payoff if he tells the truth in stage $t+1$ is $\delta\beta V(k+1) + \delta(1-\beta)V(k)$. However, if he lies and sends the message n in stage $t+1$, then his continuation payoff is $v_e + \delta\beta\tau_{k+1}\frac{(1-\pi)v_e}{1-\delta} + \delta(1-\beta)V(k)$ because with probability β the receivers will catch the expert lying, and thus all will move to public randomization phase starting in stage $t+2$. Because $\tau_{k+1} = \frac{(\delta\beta V(k+1)-v_e)(1-\delta)}{\delta\beta(1-\pi)v_e}$, it is easy to check that these two continuation payoffs are the same. That is, the expert is indifferent between lying and telling the truth when the true state is u in stage $t+1$. Thus, sending the message u with probability $\sigma_e(h^t,u)(u)$ in stage $t+1$ is a best response for the expert.

Moreover, given the expert's strategies and the history h^t described in the previous paragraph, a short lived receiver's expected payoff of approving the treatment conditional on observing the message n is $EU_r(a|n) = -v_u P(u|n) + v_n\left[1 - P(u|n)\right]$, where $P(u|n) = \frac{\pi(1-\mu(h^t))[1-\sigma_e(h^t,u)(u)]}{\pi(1-\mu(h^t))[1-\sigma_e(h^t,u)(u)]+(1-\pi)}$. For the value of $\sigma_e(h^t,u)(u)$ that is given above, $EU_r(a|n) = 0$, and thus, the receiver is indifferent between approving and rejecting the treatment whenever she observes the message n. Thus, approving the treatment when she observes n and rejecting the treatment when she observes u is a best response strategy for the short-lived receivers. Hence, the strategies in (ii) (together with the strategies in (iii) and (iv)) form a PBE of the continuation game.

$V(0) = V_e$ is the highest PBE payoff that the expert can attain in the repeated sender-receiver game G^∞. We can prove this recursively: By Lemma 3, in any equilibrium of the subgame, following a history where the expert's reputation is higher than μ^*, the expert's highest payoff is v_e^d. Thus, we must have $V_{N_G} = v_e^d$, where V_{N_G} denotes the expert's highest equilibrium payoff in any continuation game following such a history. If $N_G = 0$, then we are done. If, however, it is positive, then the expert's highest equilibrium payoff in any subgame following a history where the expert's reputation is μ_{N_G-1}, call it V_{N_G-1}, must be less than $(1-\pi)[v_e+\delta V_{N_G-1}]+\pi\delta V_{N_G}$. This is true because in the next stage, either the true state will be n, and so the expert's highest continuation payoff will be $v_e+\delta V_{N_G-1}$, or the true state will be u and the expert will tell the truth, build his reputation to a level above μ^*, and receive at most $0+\delta V_{N_G}$. Thus, solving all these inequalities recursively will yield that the expert's highest equilibrium payoff in the game must be less than or equal to V_e.

In the strategy profile that is described to prove Proposition 3, the expert builds his reputation gradually, which delays the payoff v_e^d. However, the expert would play a strategy in which he builds his reputation in few stages (by telling the truth when the true state is u with a probability less than $\sigma_e(h^t,u)(u)=\frac{v_n(1-\pi)}{v_u\pi(1-\mu(h^t))}$). However, if the expert follows a strategy in which he lies with a probability greater than $\sigma_e(h^t,u)(u)$, then in equilibrium, the receiver certainly rejects the treatment regardless of the message (recall that $\sigma_e(h^t,u)(u)$ is the probability of truth telling that makes the receivers indifferent between approving and rejecting the treatment). Thus, building reputation faster than N_G stages implies that the expert should give up his positive stage game payoffs until his reputation reaches μ^*. However, for such a strategy to be a part of an equilibrium strategy, the expert's continuation payoff of telling the truth and lying (in case the true state is u) must be the same. Suppose that the expert lies with a very high probability so that he can build up his reputation just in 1 stage by telling the truth when the true state is u. Therefore, the expert's stage game payoff of sending message u (when the true state is u) is 0, and so, his continuation payoff is at most $0+\delta v_e^d$. However, if he lies and gets caught, his continuation payoff will be at most $0+\delta v_e^f$ (by Proposition 2), not $v_e+\delta v_e^f$. The δv_e^f is lower than δv_e^d for all values of δ. Thus, if $0<\mu<\mu^*$, then in equilibrium, the expert cannot get a game payoff higher than v_e^f with a strategy where he does not build his reputation gradually. Equivalently, the expert can attain his highest expected payoff in a strategy profile where he gradually builds his reputation. Since N_G is the shortest time that the expert needs to build up his reputation, V_e must be the highest payoff the expert can attain in any PBE of the repeated sender-receiver game.

Proof of Corollary 1. Proposition 3 indicates that for any $\beta_0\in(0,1)$, there exists some $\delta^*<1$ high enough such that for all $\beta\in[\beta_0,1)$ and $\delta\in[\delta^*,1)$ there exist an equilibrium in which the expert's expected payoff is V_e. Now, fix the value of δ. Let β_0 be small enough, and so, δ^* be high enough so that $\frac{N_G(1-\delta)}{\delta\pi}\in[\beta_0,1]$ and $\delta\in[\delta_\beta,1)$. Then, the expert's expected payoff V_e takes its maximum value over $[\beta_0,1]$ at $\beta=\frac{N_G(1-\delta)}{\delta\pi}$. Here is why: Given the value of V_e in Proposition 3, we have

$$\frac{\partial V_e}{\partial \beta} = \frac{\partial \alpha_\beta}{\partial \beta}\left(\frac{v_e \pi}{1-(1-\pi\beta)\delta}\right) + \frac{\partial v_e^d}{\partial \beta}\alpha_\beta.$$

Since $\frac{\partial v_e^d}{\partial \beta} = -\frac{\delta \pi^2 v_e}{(1-(1-\pi\beta)\delta)^2}$ and $\frac{\partial \alpha_\beta}{\partial \beta} = \frac{\alpha_\beta N_G(1-\delta)}{\beta(1-(1-\pi\beta)\delta)}$, we have

$$\frac{\partial V_e}{\partial \beta} = \frac{\alpha_\beta v_e \pi}{(1-(1-\pi\beta)\delta)^2}\left[\frac{N_G(1-\delta)}{\beta} - \delta\pi\right].$$

Equating the last equation to 0 gives the value of β that maximizes V_e.

Proof of Corollary 2. For $N_G = 1$, it is easy to calculate the expected number of stages that the expert should be truthful to build his reputation up to μ^*. If $N_G = 1$, then it requires only one stage, where the true state is u and the receivers observe the payoffs, to build reputation. Therefore, the expected number of stages that the expert should be truthful is

$$\sum_{i=0}^{\infty}(i+1)(1-\pi\beta)^i \pi\beta = \frac{1}{\pi\beta},$$

where the $(i+1)^{th}$ stage is the first stage in which the true state is u and the receiver observes the realized payoffs, and $(1-\pi\beta)^i\pi\beta$ is the probability of this event. For arbitrarily large but finite N_G, we inductively calculate the expected number of stages that the expert should be truthful.

Recall that the expert can exploit the receivers in equilibrium only if the expert's reputation is higher than μ^*, and the expert's equilibrium payoff is positively related with the number of exploitation. According to the equilibrium strategies presented in the proof of Proposition 3, where the expert's payoff is the highest, the expert can exploit the receiver at only one stage when $\beta = 1$. However, for smaller values of β, the expert can deceive the receiver as long as the receivers do not observe the previous deceptions. According to these equilibrium strategies, the probability that the expert deceives the receiver during the entire repeated sender-receiver game only once is

$$\sum_{i=0}^{\infty}\pi(1-\pi)^i\beta = \beta.$$

Similarly, the probability that the expert deceives the receiver during the entire repeated sender-receiver game only twice is

$$\sum_{i=0}^{\infty}(i+1)\pi^2(1-\pi)^i(1-\beta)\beta = (1-\beta)\beta.$$

Inductively, we can find that the probability that the expert deceives only n times is

$$\sum_{i=n}^{\infty}\binom{i-1}{n-1}\pi^n(1-\pi)^{i-n+1}(1-\beta)^{n-1}\beta=(1-\beta)^{n-1}\beta.$$

Hence, in equilibrium, the expected number stages that the expert deceives the receivers is at most

$$\sum_{i=0}^{\infty}(i+1)(1-\beta)^i\beta=\frac{1}{\beta}.$$

Proof of Proposition 4. The following strategies form equilibrium of the repeated sender-receiver game where the expert exploits the receiver for an unlimited periods. Let M be the smallest natural number satisfying

$$M\geq\frac{ln\left[1-\frac{(1-\delta)v_u}{\delta(1-\pi)v_n}\right]}{ln\ \delta}. \tag{8}$$

Note that M is well defined for all $\delta\geq\delta^*\equiv\frac{v_u}{v_u+(1-\pi)v_n}$. The equilibrium strategies consist of three phases; exploitation, rewarding and punishment phases. Players start in the exploitation phase, where the expert sends the message n regardless of the true state. Exploitation phase ends when the expert deceives the receiver for one period. Once the exploitation phase ends, the players move to the rewarding phase and stay there for M periods. In the rewarding phase the expert reports the state truthfully. Once the rewarding phase ends, both players move to the exploitation phase again. Therefore, the exploitation phase, followed by the rewarding phase, repeats indefinitely. In both exploitation and rewarding phases, the receiver approves (reject) the treatment whenever she observes the message n (u). If at least one of the players deviates from his/her strategies in any phase, then the players move to the punishment phase and stay there forever. In the punishment phase, the receiver and the expert play their babbling equilibrium strategies.

We already know that the punishment phase strategies are optimal. The receiver has no incentive to deviate in the rewarding phase. Furthermore, for sufficiently large values of δ, the expert's rewarding phase strategies are also optimal: If the expert deviates and sends the message n when the true state is u, then his continuation payoff is only v_e. However, if he does not deviate, then his continuation payoff is higher than his fully-revealing equilibrium payoff $\frac{(1-\pi)v_e}{1-\delta}$, which is higher than v_e for sufficiently high values of δ.

Therefore, all we need to show is that the exploitation phase strategies are also optimal. Given the strategies described above, the expert's continuation payoff in the exploitation phase is strictly positive. However, his best deviation payoff is 0. Therefore, deviation is never optimal for the expert in the exploitation phase. Define V_r to be the receiver's continuation payoff in any subgame that begins with the exploitation phase. Likewise, let W_r denote the receiver's continuation payoff in any subgame that begins with the beginning of the rewarding phase. Therefore,

$$V_r = \pi[-v_u + \delta W_r] + (1-\pi)[v_n + V_r],$$

and

$$W_r = \frac{(1-\pi)v_n(1-\delta^M)}{1-\delta} + \delta^M V_r.$$

Solving these two equalities for V_r yields

$$V_r[1-(1-\pi+\delta^M\pi)\delta] = (1-\pi)v_n - \pi v_u + \delta\pi\frac{(1-\pi)v_n(1-\delta^M)}{1-\delta}. \tag{9}$$

If the receiver deviates in the exploitation phase and rejects the treatment, then her continuation payoff is simply 0. However, if she approves the treatment, then her continuation payoff is $\pi(-v_u + \delta W_r) + (1-\pi)(v_n + \delta V_r)$, which is equal to V_r.

Note that the left hand side of the equality (9) is always positive. To show that V_r is positive, all we need to show that the right hand side of (9) is also positive. Because M satisfies the inequality (8), it is true that $\delta\frac{(1-\pi)v_n(1-\delta^M)}{1-\delta} - v_u \geq 0$. Therefore, the right hand side of (9) is positive, implying that $V_r \geq 0$. Hence, the receiver has no incentive to deviate in the exploitation phase either.

Finally, one can easily construct an equilibrium where $V_r = 0$ by extending the deception phase (possibly with public randomization) or by shortening the rewarding mode so that the discounted expected loss from exploitation mode is equal to the discounted expected payoff from the rewarding mode. This completes the proof.

References

[1] Abreu, Dilip. On the theory of infinitely repeated games with discounting. Econometrica, 56(2): 383-396, March 1988.

[2] Abreu, Dilip, David Pearce, and Ennio Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. Econometrica. 58(5): 1041-1063, September 1990.

[3] Aumann, Robert J., and Sergiu Hart. Long Cheap Talk. Econometrica, 71(6): 1619-1660, November 2003.

[4] Benabou, Roland, and Guy Laroque. Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility. The Quarterly Journal of Economics: 921-958, 1992.

[5] Bolton, Gary E., Elena Katok, and Axel Ockenfels. How effective are electronic reputation mechanisms? An experimental investigation. Management science, 50(11): 1587-1602, November 2004.

[6] Brown, Jennifer, and Dylan B. Minor. Misconduct in Credence Good Markets. NBER working paper, 2012.

[7] Crawford, Vincent P., and Joel Sobel. Strategic information transmission. Econometrica, 50(6):1431-1451, November 1982.

[8] Darby, Michael R., and Edi Karni. Free competition and the optimal amount of fraud. Journal of law and economics: 67-88, 1973.

[9] Diamond, Peter. A Model of Price Adjustment. Journal of Economic Theory, 3(2): 156-168, 1971

[10] Dulleck, Uwe, and Rudolf Kerschbamer. On doctors, mechanics, and computer specialists: The economics of credence goods. Journal of Economic literature: 5-42, 2006.

[11] Dulleck, Uwe, Rudolf Kerschbamer, and Matthias Sutter. The Economics of Credence Goods: An Experiment on the Role of Liability, Verifiability, Reputation, and Competition. American Economic Review, 101(2): 526-55, 2011.

[12] Ely, Jeffrey C., and Juuso Välimäki. Bad reputation. The Quarterly Journal of Economics: 785-814, 2003.

[13] Emons, Winand. Credence goods and fraudulent experts. The RAND Journal of Economics: 107-119, 1997.

[14] Emons, Winand. Credence goods monopolists. International Journal of Industrial Organization. 19(3):375-389, 2001.

[15] Fong, Yuk-fai. When do experts cheat and whom do they target? RAND Journal of Economics: 113-130, 2005.

[16] Fudenberg, Drew, and Eric Maskin. On the dispensability of public randomization in discounted repeated games. Journal of Economic Theory. 53(2): 428-438, 1991.

[17] Fudenberg, Drew and David K. Levine. Reputation and Equilibrium Selection in Games with a Patient Player. Econometrica. 57(4):759-778, July 1989.

[18] Fudenberg, Drew, David Levine, and Eric Maskin. The folk theorem with imperfect public information. Econometrica. 62(5): 997-1039, September 1994.

[19] Fudenberg, Drew, David M. Kreps, and Eric S. Maskin. Repeated games with long-run and short-run players. The Review of Economic Studies. 57(4): 555-573, 1990.

[20] Golosov, Mikhail, Vasiliki Skreta, Aleh Tsyvinski, and Andrea Wilson. Dynamic Strategic Information Transmission. Journal of Economic Theory, 151: 304-341, May 2014.

[21] Gaynor, Martin and William Vogt, Antitrust and Competition in Health Care Markets, in "Handbook of Health Economics, Vol. 1," Elsevier, 2000, chapter 20.

[22] Gruber, Jon, John Kim, and Dina Mayzlin. Physician fees and procedure intensity: the case of cesarean delivery. Journal of health economics. 18(4): 473-490, 1999.

[23] Krishna, Vijay, and John Morgan. The art of conversation: eliciting information from experts through multi-stage communication. Journal of Economic Theory, 117(2): 147-179, August 2004.

[24] Levitt, Steven D., and Chad Syverson. Market distortions when agents are better informed: The value of information in real estate transactions. The Review of Economics and Statistics 90(4): 599-611, 2008.

[25] McGuire, Thomas, Physician Agency, in “Handbook of Health Economics, Vol. 1,” Elsevier, 2000, chapter 9.

[26] Morris, Stephen. Political Correctness. Journal of Political Economy, 109 (2): 231-265, April 2001.

[27] Ottaviani, Marco, and Peter Norman Sorensen (2006). Professional Advice. Journal of Economic Theory, 126 (1): 120-142, January 2006

[28] Ottaviani, Marco, and Peter Norman Sorensen (2006b). Reputational Cheap Talk. RAND Journal of Economics, 37(1): 155-175, Spring 2006

[29] Özer, Özalp, Yanchong Zheng, and Yufei Ren. Trust, trustworthiness, and information sharing in supply chains bridging China and the United States. Management Science. 60(10): 2435-2460, October 2014.

[30] Özer, Özalp, Yanchong Zheng, and Kay-Yut Chen. Trust in forecast information sharing. Management Science, 57(6): 1111-1137, April 2011.

[31] Pesendorfer, Wolfgang, and Asher Wolinsky. Second opinions and price competition: Inefficiency in the market for expert advice. The Review of Economic Studies 70(2): 417-437, 2003.

[32] Schneider, Henry S. Agency problems and reputation in expert services: Evidence from auto repair. The Journal of Industrial Economics 60(3): 406-433, 2012.

[33] Sobel, Joel. A Theory of Credibility. The Review of Economic Studies, 52 (4): 557-573, October 1985

[34] Sobel, Joel. Giving and Receiving Advice. Working Paper, 2011.

[35] Wolinsky, Asher. Competition in a market for informed experts' services. The RAND Journal of Economics: 380-398, 1993.

[36] Wolinsky, Asher. Competition in markets for credence goods. Journal of Institutional and Theoretical Economics: 117-131, 1995.