Title
Essay on Strategic Information Transmission in Trading Environments

Author
Ding, Xuan

Publication Date
2017

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Essay on Strategic Information Transmission in Trading Environments

A dissertation submitted in partial satisfaction of the
requirements for the degree of Doctor of Philosophy

in

Economics

by

Xuan Ding

Committee in charge:

Professor Simone Galperti, Co-Chair
Professor Joel Watson, Co-Chair
Professor Garey Ramey
Professor Yuval Rottenstreich
Professor Marta Serra-Garcia

2017

The Dissertation of Xuan Ding is approved and is acceptable in quality and form for publication on microfilm and electronically:

Co-Chair

Co-Chair

University of California, San Diego

2017

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGEMENTS

I would like to acknowledge all of my committee members for their great support and valuable guidance throughout my studies as a Ph.D. student. Without them, I would not have completed this dissertation.

I would like to express my most sincere gratitude to my Ph.D. advisors, Professor Joel Watson and Professor Simone Galperti. They are not only my research mentors but also excellent teachers. Their passion, enthusiasm and extensive knowledge in economic studies have greatly inspired me. They have not only showed me the elegance of microeconomics, but have also taught me to develop my research ideas and to think like an economist. I enjoyed every meeting and conversation with them, and I am truly fortunate to have worked with them.

I also want to extend my appreciation to Professor Garey Ramey, Professor Snehal Banerjee, and Professor Joel Sobel for their support and guidance during my Ph.D. studies. Their valuable comments and encouragement have helped me to improve my research significantly.

My parents, Mingyuan Wang and Zishou Ding, have always given me their full support and care while I was working on my Ph.D. dissertation. The completion of my Ph.D. will mean a lot to them, as well as to me. Without their love and encouragement, this work would not have been possible.

Last but not least, I would like to say thank you to the faculty and staff members, as well as my Ph.D. student colleagues in the UC San Diego Department of Economics. You have made my time in San Diego pleasant and memorable.

VITA

2012	Bachelor of Arts, Wuhan University
2012	Bachelor of Science, Wuhan University
2017	Doctor of Philosophy, University of California, San Diego

ABSTRACT OF THE DISSERTATION

Essay on Strategic Information Transmission in Trading Environments

by

Xuan Ding

Doctor of Philosophy in Economics

University of California, San Diego, 2017

Professor Simone Galperti, Co-Chair
Professor Joel Watson, Co-Chair

In my dissertation, I explore questions about strategic decision making and interactions between market participants, with an emphasis on information transmission and information design. I apply game-theoretic models to study how strategic information transmission influences economic outcomes in trading environments.

In Chapter 1, I study the information provision role and the incentive provision role of certification in a moral hazard setting in a seller-induced certification model versus a buyer-induced certification model. My results explain empirical observations in the credit rating market and provide policy implications regarding the "issuer-pays" rating mode versus the

"investor-pays" rating mode in the financial market.

Chapter 2 examines the incentive of a long-run seller to disclose previous offers in a dynamic market for lemons and identifies the impact of allowing voluntary disclosure on the market information structure. Compared with the models of mandatory disclosure and mandatory non-disclosure, the optional disclosure model generates a novel set of equilibria by allowing flexibility in the disclosure option. My result also proves the efficiency of optional disclosure, which explains a change in eBay's disclosure scheme about previous transactions.

In Chapter 3, I consider a long-term contracting problem between a monopolistic seller and a present-biased buyer with asymmetric information in a Markov environment. The buyer and the seller are fully aware of the degree of inconsistent discounting. I characterize the optimal contract and identify the novel impact of time inconsistency on the optimal allocations.

Chapter 1

Strategic Certification under Information Asymmetry

When there are severe moral hazard problems in the production of goods and services, both sellers and buyers demand certification. In these markets, certification has two important roles. One role is to reveal unobservable information about product quality to the market participants, while the other is to incentivize the sellers in production investment. In addition, the tradeoffs between these two roles depend on whether it is the buyer or the seller who initiates certification; this in turn affects how the certifier designs the certification mechanism. I show that the certifier can strategically inflate the equilibrium information structures in both models. However, which model creates a more efficient economic outcome is indeterminate and depends on the elasticity of the marginal cost of production. Although the certifier obtains a higher payoff when serving the seller, policy regulation is necessary when the marginal cost of production is very elastic.

1.1 Introduction

Producers make investments in R&D, capital, and other related resources to improve their product quality; however, consumers often have inadequate information about either these investments or the product quality. Such a moral hazard problem leads producers to shirk in production, brings under-provision of quality, and raises severe problems in market transactions.

One way to solve this inefficiency is to introduce a third party market intermediary —i.e., the certifier —who fosters credible communications between these two sides by providing extra information about product quality. Such intermediaries are widely seen in real marketplaces. For example, credit rating agencies certify financial assets, automotive companies certify pre-owned vehicles, and laboratories examine industrial products.

These information intermediaries alleviate the existing asymmetric information problem by providing credential information to the market participants, which increases information transparency. They also influence the seller's incentive to invest in production by alternating the extent to which they disclose information to the uninformed party. A noisy revelation makes it hard to distinguish the low-quality product and the high-quality product, which has a negative impact on the buyer's willingness to pay for the good. Furthermore, the producer becomes less willing to invest in product quality. For a profit-driven certifier, the design of the information revelation scheme needs to take into consideration both aspects.

In reality, there is demand for certification from both sides of the market. Sellers hire the certifier in order to differentiate themselves from their competitors, e.g., the organic food certification; buyers use certification to alleviate information asymmetry, e.g., the auto inspection for second-hand cars. These two payment structures can be observed as the two business modes in the financial market. One is the "issuer-pays" mode in which the seller pays for certification. The other is the "investor-pays" mode in which the buyer pays for certification. Currently, the "issuer-pays" model is the primary business model. However, since the 2008 financial crisis, the credit rating agencies (CRAs) have been heavily criticized for assigning inflated ratings to financial products. For example, the United States Senate Permanent Subcommittee on Investigations (2011, 6) concluded that inaccurate AAA credit rating was a key cause of the 2008 financial crisis. Moreover, it has been argued that the "issuer-pays" mode is a major cause of rating inflation ([50]). As the CRAs are paid by the sellers of the underlying assets, they tend to issue inflated ratings to attract more customers. This could lead to social loss since low-quality assets are sold as high-quality ones. A frequently raised claim is that the CRAs should abandon the current “issuer-pays” mode and return to the “investor-pays” mode([1]). [1]

A natural approach for looking into this problem is to compare the two business models of certification and examine how the business setting affects trading efficiency and economic outcomes. In this paper, I study the role of a certifier in two different models: in one model the seller certifies and in the other model the buyer certifies, namely the seller-certification model and the buyer-certification model. This paper adopts an information design approach by allowing the certifier to choose the signal structure as well as the certification fee, where there is asymmetric information regarding the product quality. Formally, the game has one buyer, one seller, and a certifier. The seller chooses to invest unobservable effort in production, which determines the asset’s quality. This quality is privately known to the seller. Without the certifier, the seller faces a moral hazard problem, thereby exerting no effort in production. With extra information from the certifier, the asymmetric information problem can be alleviated. Within this setup, I fully characterize the equilibria under two models which differ by the party that certifies.

In these two models, certification has two important roles. First, it alleviates the asymmetric information problem by revealing more information to the buyer. Second, it incentivizes the seller to invest in production. Altogether, these roles determine the information structure in the seller-certification model and the buyer-certification model.

In the seller-certification game, there are two possible information schemes. One potential information scheme is to improve market transparency and enable the buyer to perfectly distinguish different types of goods, thereby raising the high seller type’s valuation for certification. This suggests that the signal scheme has to be extremely informative. As a consequence, only the high type takes advantage of it. The certifier also wants to improve the chance that the seller produces a high-quality product, which leads it to implement a more revealing signal scheme. Here, the two roles of certification work in the same direction, resulting in an extremely informative signal scheme. However, the fee that the certifier can charge entails some constraints.

[1] Proposals from some regulation agencies explicitly state potential changes to the business mode of the credit rating market. For instance, the European Commission of the European Union suggests consideration of an international switch from the “issuer-pays” mode to the “investor pays” mode ([12]).

The greater the fee, the less the seller earns from producing a high-quality product, and the less the effort she exerts. In other words, there is a tradeoff between charging a higher fee and incentivizing the seller in production.

The other possible information mechanism in the seller-certification game involves manipulating the signal quality so that certification becomes valuable and affordable for both seller types. In this situation, the signal scheme has to be less informative so that the low type can take advantage of it. Nevertheless, the seller's valuation of certification also depends on the buyer's willingness to pay upon receiving a signal. When a good signal arrives, the buyer is willing to pay more for the good. Therefore, the certifier wants to improve the buyer's expectation of product quality conditional on the good signal. Accordingly, the signal mechanism is implemented to be a more revealing one so as to encourage the seller to produce a high-quality product. The two roles of certification work in different directions in this mechanism. The resulting signal scheme becomes somewhat noisy. In the paper, I find that these two mechanisms exist in equilibrium.

In the buyer-certification model, for a given prior of the product type, the buyer values certification the most when it is fully revealing. This causes the mechanism to be more informative. However, among all possible priors, the buyer values certification the most when he is very uncertain of the product quality. In addition, the certifier can charge more when the buyer is willing to pay more for the signal. Thus, it wants to design the signal mechanism so that the induced effort level leads to a very uncertain product type. This leads the mechanism in a way that incentivizes more effort when the effort level is low and could cause less effort when the effort level is high. Here, the two roles interact in a complicated fashion, and the resulting signal scheme could be fully revealing or noisy to some extent.

In both models, the certifier could strategically produce noise in the equilibrium signal structure, which exerts a negative impact on market transparency. In terms of social welfare, it is directly related to the induced effort level in equilibrium. In the paper, I find that the welfare comparison depends on the elasticity of the seller's marginal cost of production. This is because the elasticity of the marginal cost of production determines how effectively the certifier can influence the seller's effort in production, thereby shaping the signal structure. When the elasticity of the marginal cost of production is small, the seller-certification model generates a higher effort level. In contrast, when the elasticity of the marginal cost of production is large, the buyer-certification model results in a higher equilibrium effort level.

When the elasticity of the marginal cost of production is large, the induced equilibrium effort level is very irresponsive to a change in either the fee or the signal structure. To illustrate, take the seller-certification model first and consider the separating equilibrium with a very informative signal scheme. In this setting, the certifier's payoff depends on both the probability of producing a high-quality asset and the fee. Under the assumption of a very elastic marginal cost function, the gain from charging a higher fee outweighs the loss from lowering the probability of producing a high-quality asset. Therefore, the certifier will charge an expensive fee, which results in a low effort level in equilibrium. However, in the buyer-certification model, the buyer's valuation for certification is essentially the expected payment (to a low-type seller) that can be avoided with certification. This amount depends both on the probability of a low quality good and the payment (to a low-type seller) avoided with certification. Here, the gain from increasing the likelihood of producing a low-quality good is less than the loss from reducing the payment avoided with certification. The signal structure will be slightly noisy (or fully revealing), which induces a high effort level in equilibrium. As a result, the buyer-certification model could lead to a more efficient outcome than the seller-certification model. The opposite case can also be shown using this intuition.

I also discover that the certifier's equilibrium profit is higher under the seller-certification model; thus, the certifier's incentive to make profit and the policy maker's interest in improving social welfare are not always aligned. My theoretical results are consistent with the observation that in the current financial market the "issuer-pays" mode is more widespread. However, it is optimal to have active market interventions into the current business model under some circumstances.

The remainder of the paper is presented as follows. Section 1.2 presents a discussion of the related literature. Section 1.3 contains the detailed model setup. In Section 1.4 I show the equilibrium result of the seller-certification model. In Section 1.5 I discuss the buyer-certification model. Section 1.6 is devoted to a comparison of the two models. Section 1.7 extends the model setting, and Section 1.8 concludes. The proofs are shown in Appendix A.

1.2 Literature Review

1.2.1 Theoretical Findings on Certification

Since the financial crisis in 2008, there has been an explosion of research papers studying the causes of the financial crisis and the role of the credit rating agencies in the subprime crisis. Within this literature, there are a few theoretical papers that explicitly compare these two models. [46] discuss the provision of certification to the seller versus to the buyer in an adverse selection model when the monopolistic certifier commits to truthful reporting. They conclude that seller-certification leads to more transparency compared with buyer-certification. It also generates higher profit for the certifier; thus, it is not necessary to regulate the current business model. Another paper is [17]. They consider a model where the certifier can serve both the seller and the buyers. The product quality becomes publicly known if the seller demands a rating; it is the buyer's private information if the buyer pays for it. Both of these assume that the certifier truthfully communicates the quality of the product. [26] analyze a model where the investors decide whether to finance a project with unknown quality based on the CRA's rating. The rating quality depends on unobservable effort exerted by the CRA. They find that the rating is obtained less frequently and less accurately when the issuer pays for it.

The main difference in my paper is to approach the problem from an information design perspective. Specifically, I allow the certifier to choose the signal structure freely. This not only considers the certifier's role of increasing market transparency but also takes into account the impact of the rating criteria on production-related investment. The paper identifies the distinct tradeoffs between information provision and incentive provision in the two business models.

There is also a growing set of research papers that focus on the certifier's incentive to manipulate the signal structure. [31] and [2] examine the certifier's strategic information revelation in a setting where sellers pay for certification, and they find that partial disclosure could be optimal for the certifier. Another strand of research (e.g. [45], [8]) investigates the rating distortion due to rate shopping. Several research papers study the situation where the certifier colludes with the seller and distorts the true rating (e.g. [47], [33], [41]). My paper also falls into this category. What I find is that, even without the issue of capture and the certification cost, the certifier could still strategically produce noise in signals.

1.2.2 Empirical Findings on Certification

In the earlier 1970s, the business mode of major credit rating agencies changed from the "investor-pays" mode to the "issuer-pays" mode ([50]). Currently, the three major CRAs (Moody's, Standard and Poor and Fitch Ratings) serve only the issuer of the assets. There are empirical findings that document the difference between the two payment modes in the financial rating market; many of them point out that the "issuer-pays" mode leads to rating inflation. [24] show that Standard and Poor's assigned higher ratings to bonds after it converted to the "issuer-pays" business mode. Also, researchers ([51] and [9]) compare the ratings provided by CRAs using the "issuer-pays" business mode to the ratings assigned by those adopting the "investor-pays" business mode. They state that the ratings from CRAs choosing the "issuer-pays" business mode are systematically higher. Specifically, [32] document that most of the inaccurate ratings happened in the structured product market were on the mortgage-backed securities (MBS) and the collateralized debt obligations (CDO). [3] find evidence that there was a progressive decline in rating standards of MBSs between early 2005 and mid-2007. [22] conclude that the ratings for the MBS market were likely to be inflated during the boom period, especially for large issuers.

My paper's result is consistent with the market observation that the "issuer-pays" mode dominates the financial rating market. It also provides theoretical support to the empirical findings. In the paper, I show that there exist strategic rating distortions in equilibrium in the "issuer-pays" model. It is due to the reason that the certifier chooses to serve both the high-type seller and the low-type seller. Nevertheless, from a social welfare perspective, the "issuer-pays" mode could still generate more accurate ratings and more efficient economic outcomes than the "investor-pays" model when the elasticity of the marginal cost of production is small.

1.2.3 Literature on Information Design in Moral Hazard Settings

This paper adopts an information design perspective in a moral hazard model, which is related to the Bayesian persuasion literature. Since the pioneering work of [25], this strand of research has been growing rapidly. There are a few papers that incorporate moral hazard into the Bayesian persuasion model. [7] study a three-player game. In the game, the principal first sends a signal to the decision maker, the agent exerts unobservable effort that determines the underlying state, and finally the decision maker takes an action that determines all players' payoff. Also, [43] and [44] study the interplay between information disclosure and incentives in principal-agent relationships both in a career concern setting and a grading scheme design setting.

My paper complements this strand of literature and explores how the principal's objective affects the signal structure. The equilibrium signal structures in both certification models can be noisy, yet they are designed for different reasons. In the seller-certification model, the certifier could inflate the signals to attract both seller types. In contrast, in the buyer-certification model, the signal structure may be noisy in order to increase the uncertainty of the good's quality.

1.2.4 Literature on Using Moral Hazard Models in the Financial Market

In my paper, the asymmetric information comes from the seller's unobservable effort in production. This fits well into the situation of MBSs. An MBS issuer needs to exert costly effort to screen the candidate borrowers; moreover, she can securitize the loans and sell them in a secondary market. The effort she has exerted is usually unobservable by the other side of the market.

My paper is also related to research papers that use hidden action models in the underwriting practices of assets. One representative paper is [21]. They study the optimal design of MBS in a dynamic setting with moral hazard. Moreover, some empirical studies also document this moral hazard problem. [34] study the effect of shifts in the supply of mortgage credit, and their result suggests that security writers' moral hazard in screening was a main contributing factor to the mortgage default crisis.

1.3 The Model Setup

The paper studies a game with three players: a seller (she), a buyer (he) and a certifier (it). The seller's investment in production is depicted by an effort level $e \in [0,1]$, and she produces one unit of good. The effort level is only observable to the seller. The quality of the good is $\theta \in \Theta = \{\underline{\theta}, \overline{\theta}\}$, which is stochastically determined by e according to $f(\theta|e) = e$.[2] For simplicity, $\underline{\theta}$ is normalized to 0. The cost of production is given by the function $c(e)$. I assume that $c'(e) \geq 0$, $c''(e) \geq 0$ ($c''(e) > 0$ for $e > 0$), and $c(0) = 0$. The buyer's willingness to pay for a type-θ good is θ. The seller observes the good's quality, but the buyer does not. This asymmetric information would give the seller an incentive to shirk in production if the certifier were not present.

In addition, I assume that the certifier can perfectly observe the good's quality at no cost if either the seller or the buyer pays for the service.[3] The certifier chooses a fixed fee F and a signal structure π. A signal structure π consists of a finite signal realization space S and a set of distributions $\{\pi(\cdot|\theta)\}$ over S. Both the seller and the buyer can observe this signal structure and

[2] Here, I normalize the probability of producing a high quality good to be e itself.

[3] In this paper, I focus on how the certifier designs the information structure. The certification cost can be seen as a fixed cost to gain the techniques and expertise to conduct certification, and it is not the focus of this paper. This assumption allows me to isolate the noise in signals arising from strategic manipulation from the noise caused by the capability to certify. In fact, the assumption that the sender perfectly knows the underlying state is a common assumption in the Bayesian persuasion literature.

the fixed fee. If the good is certified, a signal realization $s \in S$ becomes publicly observable. In addition, Π is the set of all possible signal structures. The seller sets a price p in the market, and the buyer decides whether to accept p based on the available information. The certifier's payoff is denoted by U_c.

My research question centers on the two certification business models, i.e., the seller-certification model and the buyer-certification model. Therefore, I have two different games that capture these two business patterns. The timelines of the games are presented below:

The seller-certification game:

Stage 1: The certifier commits publicly to a signal structure $\pi_s(\cdot)$ and a fixed fee F_s.

Stage 2: Having observed π_s and F_s, the seller chooses an effort level e.

Stage 3: After learning θ, the seller picks a price p.

Stage 4: The seller decides whether to certify.

Stage 5: The buyer decides whether to purchase the good.

The buyer-certification game:

Stage 1: The certifier commits publicly to a signal structure $\pi_b(\cdot)$ and a fixed fee F_b.

Stage 2: Having observed π_b and F_b, the seller chooses an effort level e.

Stage 3: After learning θ, the seller picks a price p.

Stage 4: The buyer decides whether to certify.

Stage 5: The buyer decides whether to purchase the good.

Notice that the models incorporate both a moral hazard problem and an adverse selection problem. The moral hazard problem appears in the production stage; the adverse selection problem appears in the price-quotation stage. The only difference between these two games is the party that initiates and pays for certification. In the seller-certification game, I assume the seller sets a price before the realization of the certification outcome. However, there are situations in real marketplace in which sellers charge prices contingent on the actual certification outcomes. In Section 1.7, I consider a model capturing this model variation, and I find that the variation does not change the effort level in equilibrium. Therefore, I will maintain this assumption here.

In both of these two games, the certifier's signal structure and the fixed fee initiate a proper subgame that involves the participation of only the seller and the buyer. The solution concept I use is perfect Bayesian equilibrium (PBE). The certifier's signal structure and fixed fee maximize its payoff given the buyer's and the seller's strategies in the proper subgame.

Since the buyer has binary actions, i.e. buying or not buying, the signal realization space essentially has binary elements. I denote $S = \{G, B\}$. Without loss of generality, G stands for the good signal, while B stands for the bad signal. In other words, $E(\theta|G) \geq E(\theta|B)$. For notation purpose, I denote $p_G = E(\theta|G)$ and $p_B = E(\theta|B)$.

1.4 The Seller-Certification Game

I start with analyzing the equilibrium of the seller-certification game. The certifier picks $\pi_s(\cdot)$ and F_s at the beginning of the game. Observing the signal structure, the seller's (pure) strategy involves an effort choice e, a price choice p, and a choice c_s of certifying (1) or not (0). Observing the signal structure, the seller's certification choice, possibly a signal from the certifier, and the good's price, the buyer forms a belief μ_s of the seller being a high-type and decides whether to purchase the good. A (pure) strategy of the buyer is a choice b of buying the good (1) or not (0).

1.4.1 The Baseline Model with a Fully Revealing Signal Structure

Before I present the result of the full model, let me consider a model where the certifier is restricted to a fully revealing signal scheme ($\pi_s\left(G|\overline{\theta}\right) = 1$, $\pi_s\left(G|\underline{\theta}\right) = 0$). This restriction helps us understand the dynamics in the model. The difference between the baseline model and the full model also illustrates how information design works here.

In this restricted model, the seller only has an incentive to employ the certifier when she is the high type; certification itself becomes a signal of high quality. In equilibrium, only the high type certifies, and she must do so. Accordingly, the buyer is willing to pay $\overline{\theta}$ for a good with the signal G. His valuation is $\underline{\theta}$ when there is no certification or the signal is B. In equilibrium, the seller's payoff is $\overline{\theta} - F_s$ if $\theta = \overline{\theta}$, and 0 if $\theta = \underline{\theta}$. Using these arguments, I can derive the equilibrium effort choice e, which satisfies $c'(e) = \overline{\theta} - F_s$. The certifier's expected revenue is eF_s, where F_s satisfies the seller's participation constraint that $e\left(\overline{\theta} - F_s\right) - c(e) \geq 0$. The certifier could charge a larger fee, but it is able to collect that amount only when the realized type is high. Its optimality condition is presented as follows.

$$\frac{dU_c}{dF_s} = e + F_s \frac{de}{dF_s} = 0 \tag{1.1}$$

Equation (1.1) captures two effects of changing F_s on the certifier's payoff. One is a positive direct effect. For a fixed effort level, charging a higher fee directly leads to a higher payoff. The other is a negative indirect effect, which indicates how F_s influences U_c through changing the effort level. In equilibrium, the amount of effort is determined by $c'(e) = \overline{\theta} - F_s$. When the certifier raises F_s, it reduces the seller's marginal revenue from exerting effort; thus, it leads to a lower effort level and a lower chance to produce a high quality good. $\frac{de}{dF_s}$ is quantified as $-\frac{1}{c''(e)}$. When F_s is large, the indirect effect dominates in the two; when F_s is small, the direct effect dominates. The two effects cancel out at the optimal F_s. This is summarized in Proposition 1.

Proposition 1 *In every equilibrium of the baseline model, the seller certifies only if she is a high type, and she trades at a price equal to $\overline{\theta}$. The effort level on the equilibrium path satisfies $c'(e_s^*) = \overline{\theta} - e_s^* c''(e_s^*)$. The certification fee is $e_s^* c''(e_s^*)$.*

Nevertheless, this equilibrium outcome is not efficient as the induced effort level is strictly below the first-best. This is because a high seller type has to pay a fee to establish her credibility. Therefore, the marginal revenue is decreased from $\overline{\theta}$ to $\overline{\theta} - F_s^*$, and the seller's optimality condition becomes $c'(e) = \overline{\theta} - F_s^*$. It is easy to check that the participation constraint holds, which is done in the proof of Proposition 1.

1.4.2 The Full Model

After analyzing the restricted model, let me go back to the full model where the signal structure is flexible. Here, the equilibrium must involve certification as well. Suppose not. The certifier must have a payoff of zero. Because of the moral hazard problem, the buyer forms a belief that the good has low quality, and his willingness to pay is $\underline{\theta}$. Clearly, the seller has no incentive to put in any effort, and she ends up with a surplus of $\underline{\theta}$, which is essentially 0 from the assumption made earlier. However, from the result in the baseline model, if the certifier adopts a fully revealing signal structure ($\pi_s\left(G|\overline{\theta}\right) = 1$, $\pi_s\left(G|\underline{\theta}\right) = 0$) with a small fee ε, the seller can receive a much higher expected surplus by investing in production and certifying. Therefore, I conclude that certification must be obtained in equilibrium.

Naturally, there are two possible equilibrium types.[4] One is separating. In this case, the signal structure is so precise and informative that the seller benefits from certification only when she is a high type. This essentially leads to the same equilibrium outcome as the one I characterized in Proposition 1. The separating property can be preserved even under a flexible signal structure.

The other equilibrium type is pooling so that the seller types are distinguishable by neither the certification choices nor the asking prices. There is some noise in the signal structure so that even the low seller type could benefit from certifying. The certifier charges a small fee that both seller types find affordable. In the subgame, the seller certifies and charges the price p_G regardless of her type. The buyer accepts p_G only when the signal is G. Here, the certifier's

[4] Some readers may wonder if there exists a semi-pooling equilibrium. The answer is no. Suppose that the low seller type hires the certifier with probability $\alpha < 1$ in a semi-pooling equilibrium; as a result, the certifier receives payment from the low seller type with probability α. However, the certifier could slightly modify the signal structure and the certification fee to guarantee that certification is always obtained, which would raise its revenue. Thus, there is no semi-pooling equilibrium. The complete proof of this statement can be found in the proof of Proposition 2 in Appendix A.

maximum possible fee is $\pi_s(G|\underline{\theta})\,p_G$, which leaves the low type zero surplus. Moreover, there is also the seller's participation constraint, i.e., $\sum_{\theta\in\{\underline{\theta},\overline{\theta}\}} p_G\pi_s(s|\theta)\,f(\theta|e)-c(e)-F_s\geq 0$, and her effort incentivizing constraint, i.e., $c'(e)=p_G\left(\pi_s\left(G|\overline{\theta}\right)-\pi_s(G|\underline{\theta})\right)$. The ideal signal structure maximizes the certifier's payoff under the seller's participation constraint and the effort incentivizing constraint. In Proposition 2, I characterize the equilibria of this model; these two categories both exist.

Proposition 2 *In the seller-certification game, the effort level on the equilibrium path satisfies $c'(e_s^*)=\overline{\theta}-e_s^*c''(e_s^*)$. There are two equilibrium types:*

(1) One equilibrium involves separation of types. The certifier adopts a fixed fee F_s^* and a signal scheme such that only the high type certifies ($\pi_s^*\left(G|\overline{\theta}\right)=1$ and $\pi_s^*(G|\underline{\theta})\,\overline{\theta}<F_s^*$). The high seller type trades at a price equal to $\overline{\theta}$.

(2) The other equilibrium involves pooling of types. The certifier adopts a fixed fee F_s^* and a signal scheme such that both types certify ($\pi_s^*\left(G|\overline{\theta}\right)=1$ and $\pi_s^*(G|\underline{\theta})\,p_G^*=F_s^*$). The good is traded only when the signal is G, and is traded at p_G^*.

In the separating equilibrium, $\overline{\theta}\pi_s^*(G|\underline{\theta})<F_s^*$ guarantees that only the high seller type certifies. The certification mechanism is so informative that the two types become naturally distinguishable by their certification actions. Here, the certifier faces the same tradeoff as in the baseline model. In equilibrium, the certifier charges F_s^*, which balances the direct effect and the indirect effect of changing F_s. Using equation (1.1), I still have $F_s^*=e_s^*c''(e_s^*)$. Accordingly, the effort level satisfies $c'(e_s^*)=\overline{\theta}-e_s^*c''(e_s^*)$.

In the pooling equilibrium, the high seller type and the low seller type are not completely distinguishable by their certification behaviors and the certification results. Here, even a low-quality good has a chance to be assigned a positive signal. The fee is equal to the gain that the low type can receive from certification; thus, $U_c=\pi_s(G|\underline{\theta})\,p_G$. The optimality conditions are given as follows.

$$\frac{dU_c}{d\pi_s\left(G|\overline{\theta}\right)}=\underbrace{\frac{\left(\overline{\theta}-p_G\right)e\pi_s\left(G|\underline{\theta}\right)}{e\pi_s\left(G|\overline{\theta}\right)+\left(1-e\right)\pi_s\left(G|\underline{\theta}\right)}}_{\text{the direct effect of }\pi_s\left(G|\overline{\theta}\right)\text{ on }U_c}+\underbrace{\pi_s\left(G|\underline{\theta}\right)\frac{dp_G}{de}\frac{de}{d\pi_s\left(G|\overline{\theta}\right)}}_{\text{the indirect effect of }\pi_s\left(G|\overline{\theta}\right)\text{ on }U_c} \tag{1.2}$$

$$\frac{dU_c}{d\pi_s\left(G|\underline{\theta}\right)}=\underbrace{\frac{e\pi_s\left(G|\overline{\theta}\right)p_G}{e\pi_s\left(G|\overline{\theta}\right)+\left(1-e\right)\pi_s\left(G|\underline{\theta}\right)}}_{\text{the direct effect of }\pi_s\left(G|\underline{\theta}\right)\text{ on }U_c}+\underbrace{\pi_s\left(G|\underline{\theta}\right)\frac{dp_G}{de}\frac{de}{d\pi_s\left(G|\underline{\theta}\right)}}_{\text{the indirect effect of }\pi_s\left(G|\underline{\theta}\right)\text{ on }U_c} \tag{1.3}$$

Equation (1.2) and (1.3) state the effects of changing $\pi_s\left(G|\overline{\theta}\right)$ and $\pi_s\left(G|\underline{\theta}\right)$ on U_c. On the one hand, in equation (1.2) and (1.3), the first arguments show how the change in $\pi_s\left(G|\underline{\theta}\right)$ and $\pi_s\left(G|\underline{\theta}\right)$ influence the certifier's payoff directly, when e is fixed. If $\pi_s\left(G|\underline{\theta}\right)$ increases, the low type can have a better chance to trade at p_G, which overshadows the drop in p_G. Therefore, the direct effect of $\pi_s\left(G|\underline{\theta}\right)$ is positive. Similarly, if $\pi_s\left(G|\overline{\theta}\right)$ increases, the expected valuation conditional on the signal G improves; the direct effect of $\pi_s\left(G|\overline{\theta}\right)$ is also positive.

On the other hand, the effort level is determined according to

$$c'\left(e\right)=p_G\left(\pi_s\left(G|\overline{\theta}\right)-\pi_s\left(G|\underline{\theta}\right)\right) \tag{1.4}$$

A higher effort level increases the probability of producing a type-$\overline{\theta}$ good, which improves the buyer's willingness to pay conditional on the good signal. Therefore $\frac{dp_G}{de}\geq 0$. The indirect effects also depend on $\frac{de}{d\pi_s\left(G|\overline{\theta}\right)}$ and $\frac{de}{d\pi_s\left(G|\underline{\theta}\right)}$, which are presented as follows.

$$\frac{de}{d\pi_s\left(G|\overline{\theta}\right)}=H^{-1}\frac{e\overline{\theta}\left(\pi_s\left(G|\overline{\theta}\right)^2-\left(1-e\right)\left(\pi_s\left(G|\overline{\theta}\right)-\pi_s\left(G|\underline{\theta}\right)\right)^2\right)}{\left(e\pi_s\left(G|\overline{\theta}\right)+\left(1-e\right)\pi_s\left(G|\underline{\theta}\right)\right)^2} \tag{1.5}$$

$$\frac{de}{d\pi_s(G|\underline{\theta})} = -H^{-1}\frac{e\overline{\theta}\pi_s\left(G|\overline{\theta}\right)^2}{\left(e\pi_s\left(G|\overline{\theta}\right)+(1-e)\pi_s(G|\underline{\theta})\right)^2} \tag{1.6}$$

Here, $H = c''(e) - \frac{\overline{\theta}\pi_s\left(G|\overline{\theta}\right)\pi_s(G|\underline{\theta})\left(\pi_s\left(G|\overline{\theta}\right)-\pi_s(G|\underline{\theta})\right)}{\left(e\pi_s\left(G|\overline{\theta}\right)+(1-e)\pi_s(G|\underline{\theta})\right)^2}$. A higher $\pi_s\left(G|\overline{\theta}\right)$ raises p_G. It also raises the marginal payoff of devoting effort. As a result, if $\pi_s\left(G|\overline{\theta}\right)$ increases, the seller will devote more effort when $H > 0$, which gives rise to a positive indirect effect. On the contrary, the indirect effect of $\pi_s(G|\underline{\theta})$ is negative from equation (1.6) under the same condition. In equilibrium, $H > 0$ must hold. Otherwise, I would have $\pi_s\left(G|\overline{\theta}\right) \leq \pi_s(G|\underline{\theta})$, which contradicts the assumption made earlier. Since both the direct effect and the indirect effect of $\pi_s\left(G|\overline{\theta}\right)$ go in the same way, $\pi_s\left(G|\overline{\theta}\right) = 1$. As for $\pi_s(G|\underline{\theta})$, the two effects work in different directions, and they cancel out at the optimality condition, which results in $0 < \pi_s(G|\underline{\theta}) < 1$.

The certifier strategically produces inflated ratings in equilibrium. This noisy signal scheme discourages the seller from exerting effort in the production stage; the marginal payoff of effort decreases to $p_G^*(1-\pi_s^*(G|\underline{\theta}))$. The resulting effort level is below the first-best. The next corollary summarizes the welfare result in the seller-certification game.

Corollary 1 *In any equilibrium of the seller-certification game, the effort level is strictly below the first-best.*

1.5 The Buyer-Certification Game

In this section, I characterize the equilibrium in the buyer-certification game. The certifier chooses $\pi_b(\cdot)$ and F_b upfront. Observing the signal structure, the seller picks an effort level and a price. Her (pure) strategy involves an effort choice e, and a price choice p. Observing the signal structure and the asking price, the buyer forms a belief μ_b of the good being a high-type and decides whether to certify and purchase the good. A pure strategy for the buyer involves a choice c_b of certifying the good (1) or not (0), and a choice b of purchasing the good (1) or not (0).

1.5.1 The Baseline Model with a Fully Revealing Signal Structure

Similar to the previous section, I start my analysis with the baseline model where the signal scheme is restricted to be fully revealing. The equilibrium can be solved backwards. The buyer believes that the good is of high quality with probability μ_b, before he certifies. Let $U_b(p|c_b=0)$ denote the buyer's payoff when there is no certification, and $U_b(p|c_b=1)$ denote his payoff when there is certification. Without certification, he purchases the good only when his expected valuation is greater than the asking price, i.e., $U_b(p|c_b=0)=\mu_b\overline{\theta}+(1-\mu_b)\underline{\theta}-p\geq 0$. With certification, the buyer acquires extra information. He purchases the good only when a positive signal is sent, and $U_b(p|c_b=1)\geq 0$. Conditional on the good signal, let μ_b^G denote the updated belief; specifically, $\mu_b^G=\frac{\mu_b\pi_b(G|\overline{\theta})}{\mu_b\pi_b(G|\overline{\theta})+(1-\mu_b)\pi_b(G|\underline{\theta})}$. In addition, $p_G=\mu_b^G\overline{\theta}+\left(1-\mu_b^G\right)\underline{\theta}$. Given $\pi_b(\cdot)$, F_b, and the asking price p, the buyer's expected surplus conditional on certifying is $U_b(p|c_b=1)=(p_G-p)\left(\mu_b\pi_b\left(G|\overline{\theta}\right)+(1-\mu_b)\pi_b(G|\underline{\theta})\right)-F_b$. He certifies only when $U_b(p|c_b=1)\geq U_b(p|c_b=0)$.

Whether certification is obtained depends on the belief μ_b. Accordingly, there are two possible equilibrium types. In one type of equilibrium, the buyer's belief is so pessimistic ($\mu_b=0$) that he believes he cannot get extra information from certifying. This occurs when he believes the asset is bad for sure. Certification is not obtained in equilibrium; therefore, the seller shirks and devotes no effort to production. This equilibrium outcome coincides with the market outcome without a certifier. Due to the buyer's negative belief, the certifier is not actively involved in the market. This equilibrium arises here since there is no uncertainty of the good's quality when $e=0$. However, if there is a slight change to the probability transition function of effort such that there is still uncertainty at the lowest effort level, this inefficient equilibrium will not exist. This is because even at the lowest effort level, the buyer still values the extra information provided by the certifier, and he is willing to acquire its service at an appropriate fee. The seller will be incentivized to invest positive effort level in production, which points to the equilibrium in which certification is obtained.

In the other type of equilibrium, the buyer pays for certification with positive probability. Here, the low seller type mimics the high type's price. Let p_H denote the high type's asking price. One possible situation is that the buyer certifies with probability β ($\beta < 1$) conditional on the price being p_H. Accordingly, the low type always wants to mimic the high type's price. The high type will definitely receive a good rating if the good is subject to inspection. Therefore, she wants to charge the maximum possible price that makes the buyer willing to purchase the good. This suggests that $\max\{U_b(p_H|c_b=1), U_b(p_H|c_b=0)\}=0$; moreover, the asking price p_H is $\max\{\mu_b\overline{\theta}+(1-\mu_b)\underline{\theta}, p_G-\frac{F_b}{\mu_b\pi_b(G|\overline{\theta})+(1-\mu_b)\pi_b(G|\underline{\theta})}\}$. As for the certifier, it wants to guarantee the purchase of certification so that $U_b(p_H|c_b=1)\geq U_b(p_H|c_b=0)$ and $U_b(p_H|c_b=1)\geq 0$. In equilibrium, it must be that $U_b(p_H|c_b=1)=U_b(p_H|c_b=0)=0$. Otherwise, there is space for the certifier or the seller to charge more. This indicates that $p_H=\mu_b\overline{\theta}$ on the equilibrium path. Since the seller's payoff is $ep_H+(1-e)(1-\beta)p_H-c(e)$, I further derive the optimal effort choice e by using $c'(e)=\beta p_H$. From the consistency of beliefs, $\mu_b=e$. The equilibrium effort level satisfies $c'(e)=\beta e\overline{\theta}$.

The other possible situation is that the buyer certifies with probability 1 conditional on the price being p_H. Here, the low seller type is indifferent between mimicking the high type's price p_H or picking a lower price p_L. p_L is accepted only when $p_L=\underline{\theta}$. Suppose that the low type chooses p_H with probability α ($\alpha\leq 1$). Let μ_{b_H} denote the buyer's belief of a high type when the asking price is p_H. Using the same argument as in the previous paragraph, $p_H=\mu_{b_H}\overline{\theta}$ on the equilibrium path. Since the seller's payoff is $ep_H-c(e)$, I further characterize the optimal choice e by using $c'(e)=p_H$. From the consistency of beliefs, $\mu_{b_H}=\frac{e}{e+\alpha(1-e)}$; the equilibrium effort level satisfies $c'(e)=\frac{e}{e+\alpha(1-e)}\overline{\theta}$. The following proposition summarizes the equilibria in the baseline model.

Proposition 3 *In the baseline model, there are two equilibrium types:*

(1) In the type-1 equilibrium, the high type charges p_H^*, and the buyer certifies with

probability β conditional on the price being p_H^*.

(i) If $\beta < 1$, the low type charges p_H^* with probability 1. The buyer accepts p_H^* if the signal is G when certifying. With probability $1-\beta$, he accepts p_H^* directly. The effort level on the equilibrium path satisfies $c'\left(e_b^*\right) = \beta e_b^* \overline{\theta}$; in addition, $p_H^* = e_b^* \overline{\theta}$.

(ii) If $\beta = 1$, the low type randomizes between p_H^* and a lowr offer p_L^* with probability $(\alpha, 1-\alpha)$. The buyer accepts p_H^* or p_L^* only when the signal is G or $p_L^* = 0$, respectively. The effort level on the equilibrium path satisfies $c'\left(e_b^*\right) = \frac{e_b^*}{e_b^* + \alpha\left(1-e_b^*\right)}\overline{\theta}$. In addition, $p_H^* = \frac{e_b^*}{e_b^* + \alpha\left(1-e_b^*\right)}\overline{\theta}$, and p_L^* can be any non-negative value.

(2) In the type-2 equilibrium, the seller exerts no effort and asks for $p^* \geq 0$. The buyer does not certify, and he purchases the good only when $p^* = 0$. The certification fee can be any non-negative value.

1.5.2 The Full Model

Now consider the full model where the signal structure is flexible. Some properties in the baseline model carry to the full model. In any equilibrium, the buyer certifies only when $U_b(p|c_b = 1) \geq U_b(p|c_b = 0)$. After certifying, he only accepts the good when the signal is G. Similar to what I just presented in the baseline model, there are still two equilibrium types. In one type of equilibrium, the buyer's belief is so pessimistic ($\mu_b = 0$) that he thinks the good certainly has bad quality. In this type of equilibrium, certification is not acquired.

In the other type of equilibrium, the buyer must certify with probability 1. This is different from the baseline model because the certifier can use the signal structure to incentivize the seller in production and affect the outcome of the game. In the baseline model, the certifier can only pick the fee. It cannot influence the equilibrium effort level; therefore, its revenue is fixed in some sense. However, if the certifier is able to design the information structure, the effort level is endogenous and is determined by the actual signal scheme. The certifier can pick a signal scheme that maximizes its revenue. As a result, certification is always acquired in equilibrium.

If not, the certifier could adjust the signaling scheme and the fee so that the buyer strictly prefers acquiring the signal. The details of the construction can be found in the proof of Proposition 4 in Appendix A.

I also find that the seller charges the same price regardless of her type. This suggests that the asking price does not serve as a signal. It generates uncertainty regarding the quality of the asset and creates a demand for certification since the buyer cannot identify the seller type purely from the asking price. Otherwise, suppose the high type and the low type ask for different prices. One possible case is a fully separating equilibrium where the seller types set completely different prices. This requires the signal scheme to be fully separating ($\pi_b\left(G|\overline{\theta}\right)=1$, $\pi_b\left(G|\underline{\theta}\right)=0$). However, this situation is not possible because the certifier would have no revenue. The certifier could distort the signal scheme so that $\pi_b\left(G|\underline{\theta}\right)>0$. Here, the low seller type would always mimic the high type's price as she could have a chance to receive the better price; the buyer would therefore have a demand for certification. Another possible situation is a semi-pooling equilibrium where the low type randomizes between a high price and a low price. This still suggests $\pi_b\left(G|\overline{\theta}\right)=1$ and $\pi_b\left(G|\underline{\theta}\right)=0$. In this situation, the buyer demands certification only when the low type picks a high price. However, the certifier is not always hired. Similar to the previous case, the certifier could slightly change $\pi_b\left(G|\underline{\theta}\right)$ to be a small positive number to reach a better payoff.

Let me go one step back and study the pricing strategies of the seller and the certifier. Denote $\overline{p}=p_G-\frac{F_b}{\mu_b\pi_b\left(G|\overline{\theta}\right)+(1-\mu_b)\pi_b\left(G|\underline{\theta}\right)}$. At price $\overline{p}$, $U_b\left(\overline{p}|c_b=1\right)=0$. If $U_b\left(\overline{p}|c_b=1\right)<U_b\left(\overline{p}|c_b=0\right)$, the seller would charge $p=\mu_b\overline{\theta}$ to achieve her maximum payoff, and the buyer would not certify. If $U_b\left(\overline{p}|c_b=1\right)\geq U_b\left(\overline{p}|c_b=0\right)$, the high type's expected payoff would be $\pi_b\left(G|\overline{\theta}\right)\overline{p}$ when she asked for $\overline{p}$, conditional on being certified. In addition, if $U_b\left(\pi_b\left(G|\overline{\theta}\right)\overline{p}|c_b=1\right)<U_b\left(\pi_b\left(G|\overline{\theta}\right)\overline{p}|c_b=0\right)$, the high seller type could charge a price p higher than $\pi_b\left(G|\overline{\theta}\right)\overline{p}$ that makes $U_b\left(p|c_b=1\right)<U_b\left(p|c_b=0\right)$. At this price, the seller would obtain a better payoff than $\pi_b\left(G|\overline{\theta}\right)\overline{p}$; the buyer would not certify. If $U_b\left(\pi_b\left(G|\overline{\theta}\right)\overline{p}|c_b=1\right)\geq U_b\left(\pi_b\left(G|\overline{\theta}\right)\overline{p}|c_b=0\right)$, the seller would charge $\overline{p}$, which could induce the buyer to certify. As

for the certifier, it wants to guarantee that certification is demanded. Therefore, I establish the condition that $U_b\left(\pi_b\left(G|\overline{\theta}\right)\overline{p}|c_b=1\right)=U_b\left(\pi_b\left(G|\overline{\theta}\right)\overline{p}|c_b=0\right)$. The fee has to satisfy the following expression.

$$F_b=\mu_b\pi_b\left(G|\overline{\theta}\right)\overline{\theta}-\frac{\mu_b\pi_b\left(G|\overline{\theta}\right)+(1-\mu_b)\,\pi_b\left(G|\underline{\theta}\right)}{\pi_b\left(G|\overline{\theta}\right)+\left(1-\pi_b\left(G|\overline{\theta}\right)\right)\left(\mu_b\pi_b\left(G|\overline{\theta}\right)+(1-\mu_b)\,\pi_b\left(G|\underline{\theta}\right)\right)}\mu_b\overline{\theta} \tag{1.7}$$

In addition, the seller's asking price is $\overline{p}$. The optimality condition of the seller's effort choice is characterized by the following equation.

$$c'(e)=\overline{p}\left(\pi_b\left(G|\overline{\theta}\right)-\pi_b\left(G|\underline{\theta}\right)\right) \tag{1.8}$$

Given the consistency of beliefs, $\mu_b=e$ and $p_G=\frac{e\overline{\theta}\pi_b\left(G|\overline{\theta}\right)+(1-e)\underline{\theta}\pi_b\left(G|\underline{\theta}\right)}{e\pi_b\left(G|\overline{\theta}\right)+(1-e)\pi_b\left(G|\underline{\theta}\right)}$. The seller's effort choice can be derived using equation (1.8) for a given fee and a signal scheme. Moreover, $U_c=F_b$. There are two optimality conditions.

$$\frac{dU_c}{d\pi_b\left(G|\overline{\theta}\right)}=\underbrace{e\overline{\theta}\left(1-\frac{\left(e\pi_b\left(G|\overline{\theta}\right)+(1-e)\,\pi_b\left(G|\underline{\theta}\right)\right)^2-(1-e)\,\pi_b\left(G|\underline{\theta}\right)}{\left(\pi_b\left(G|\overline{\theta}\right)+\left(1-\pi_b\left(G|\overline{\theta}\right)\right)\left(e\pi_b\left(G|\overline{\theta}\right)+(1-e)\,\pi_b\left(G|\underline{\theta}\right)\right)\right)^2}\right)}_{\text{the direct effect of }\pi_b\left(G|\overline{\theta}\right)\text{ on }U_c}+\underbrace{\frac{dU_c}{de}\frac{de}{d\pi_b\left(G|\overline{\theta}\right)}}_{\substack{\text{the indirect effect of}\\ \pi_b\left(G|\overline{\theta}\right)\text{ on }U_c}} \tag{1.9}$$

$$\frac{dU_c}{d\pi_b\left(G|\underline{\theta}\right)}=\underbrace{-\frac{(1-e)\,\pi_b\left(G|\overline{\theta}\right)e\overline{\theta}}{\left(\pi_b\left(G|\overline{\theta}\right)+\left(1-\pi_b\left(G|\overline{\theta}\right)\right)\left(e\pi_b\left(G|\overline{\theta}\right)+(1-e)\,\pi_b\left(G|\underline{\theta}\right)\right)\right)^2}}_{\text{the direct effect of }\pi_b\left(G|\underline{\theta}\right)\text{ on }U_c}+\underbrace{\frac{dU_c}{de}\frac{de}{d\pi_b\left(G|\underline{\theta}\right)}}_{\substack{\text{the indirect effect of}\\ \pi_b\left(G|\underline{\theta}\right)\text{ on }U_c}} \tag{1.10}$$

These two equations characterize the overall effects of $\pi_b\left(G|\overline{\theta}\right)$ and $\pi_b\left(G|\underline{\theta}\right)$ on the certifier's payoff. Given a fixed effort level, for a higher $\pi_b\left(G|\overline{\theta}\right)$ or a lower $\pi_b\left(G|\underline{\theta}\right)$, the signal structure becomes more informative; as a result, the buyer values certification more. In summary, the direct effect is positive for $\pi_b\left(G|\overline{\theta}\right)$ while it is negative for $\pi_b\left(G|\underline{\theta}\right)$.

The indirect effects depend on $\frac{dU_c}{de}$. If the effort level is pretty low, the seller is highly likely to produce $\underline{\theta}$; the buyer does not value certification much. Similarly, he does not value certification much when the effort level is pretty high. He values certification the most when he is very uncertain of the asset's quality; thus, the certifier's payoff is maximized when e is of moderate value. This can be shown mathematically as $\frac{dU_c}{de} = \overline{\theta}\left(1-2e\right)\left(\pi_b\left(G|\overline{\theta}\right) - \pi_b\left(G|\underline{\theta}\right)\right)$. $\frac{dU_c}{de}$ is positive when $e \geq \frac{1}{2}$ and is negative otherwise. The indirect effects also depend on $\frac{de}{d\pi_b\left(G|\overline{\theta}\right)}$ and $\frac{de}{d\pi_b\left(G|\underline{\theta}\right)}$, which are derived in Appendix A.

Whether the buyer-certification game has an interior solution or a corner solution depends on the cost function. Let $\tilde{e}$ denote the effort level at the corner solution ($\pi_b\left(G|\overline{\theta}\right) = 1$ and $\pi_b\left(G|\underline{\theta}\right) = 0$), i.e., $c'\left(\tilde{e}\right) = \tilde{e}\overline{\theta}$. The specific equilibrium characterization can be found in the following proposition.

Proposition 4 *In the buyer-certification, there are two equilibrium types:*

(1) In the type-1 equilibrium, the certifier adopts a fixed fee F_b^* and signal scheme such that $\pi_b^*\left(G|\overline{\theta}\right) = 1$ and $\pi_b^*\left(G|\underline{\theta}\right) < 1$. The seller sets $p^* = e_b^*\overline{\theta}$. The buyer always certifies; he accepts the asking price only if the signal is G.

(i) If $\varepsilon_{c'(e)} < \min\{\frac{e}{1-e}, 1\}$ or $\varepsilon_{c'(e)} > \max\{\frac{e}{1-e}, 1\}$ for all e, $\pi_b^*\left(G|\underline{\theta}\right) = 0$. The equilibrium effort level satisfies $c'\left(e_b^*\right) = e_b^*\overline{\theta}$.

(ii) If $\varepsilon_{c'(\tilde{e})} \in \left(\frac{\tilde{e}}{1-\tilde{e}}, 1\right)$ (or $\varepsilon_{c'(\tilde{e})} \in \left(1, \frac{\tilde{e}}{1-\tilde{e}}\right)$), $\pi_b^*\left(G|\underline{\theta}\right) = 1 - \frac{c''\left(e_b^*\right)\left(1-e_b^*\right)}{e_b^*\overline{\theta}}$. The equilibrium effort level satisfies $c'\left(e_b^*\right) = \left(1-e_b^*\right)c''\left(e_b^*\right)$.

(2) In the type-2 equilibrium, the seller exerts no effort. The buyer does not certify, and he purchases the good only when the asking price is $\underline{\theta}$. The certifier's signal scheme and fee can be any non-negative values.

These two equilibrium types generate very different equilibrium outcomes. The type-1 equilibrium induces a market failure;[5] the type-1 equilibrium does not. However, the outcome of the type-1 equilibrium is also not efficient. This is because the seller cannot set a price that extracts the entire payoff of producing $\overline{\theta}$. The buyer receives an information rent, which is extracted by the certifier. As a result, the induced effort level is strictly below the first-best. Corollary 2 summarizes the welfare result in the buyer-certification game.

Corollary 2 *In any equilibrium of the buyer-certification game, the effort level is strictly below the first-best.*

1.6 Comparison of the Two Business Models

This section presents a comparison of the equilibrium results in the seller-certification model versus the buyer-certification model. In the seller-certification model, the effort level e_s^* satisfies equation (1.11) derived from Proposition 2.

$$\varepsilon_{c'(e_s^*)} = \frac{\overline{\theta}}{c'(e_s^*)} - 1 \tag{1.11}$$

The certifier's equilibrium payoff satisfies $U_c^* = e_s^*\left(\overline{\theta} - c'(e_s^*)\right)$. In the buyer-certification model, I will focus on the type-1 equilibrium as the type-2 equilibrium is not robust to a slight variation in the probability transition function of effort. If $\varepsilon_{c'(e)} < \min\{\frac{e}{1-e}, 1\}$ or $\varepsilon_{c'(e)} > \max\{\frac{e}{1-e}, 1\}$ for all e, the type-1 equilibrium has a corner solution; the effort level e_b^* satisfies $c'(e_b^*) = e_b^*\overline{\theta}$, and $U_c^* = (1-e_b^*)\, c'(e_b^*)$. If $\varepsilon_{c'(\tilde{e})} \in \left(\frac{\tilde{e}}{1-\tilde{e}}, 1\right)$ (or $\varepsilon_{c'(\tilde{e})} \in \left(1, \frac{\tilde{e}}{1-\tilde{e}}\right)$), this equilibrium has an interior solution; the effort level e_b^* satisfies equation (1.12) derived from Proposition 4.

$$\varepsilon_{c'(e_b^*)} = \frac{e_b^*}{1-e_b^*} \tag{1.12}$$

In this case, $U_c^* = (1-e_b^*)\, c'(e_b^*)$. In these two models, the social welfare is $e\overline{\theta} - c(e)$. Since both e_s^* and e_b^* are lower than the first best effort level, the larger one of these two leads to

[5] For the same reason as in the baseline model, the type-2 equilibrium does not exist if there is a slight variation in the probability transition function of effort.

the model with a higher social welfare in equilibrium. Theorem 1 summarizes the welfare comparison.

Theorem 1 *Considering the equilibria where the certifier is employed, there exist $A > A' > 0$ such that:*

(1) If $\varepsilon_{c'(e)} \leq A'$ for all e, the seller-certification model yields a higher social welfare than the buyer-certification model.

(2) If $\varepsilon_{c'(e)} \geq A$ for all e, the buyer-certification model yields a higher social welfare than the seller-certification model.

(3) The certifier always earns a higher profit in the seller-certification model.

Regarding the intuition behind Theorem 1, let me explain it in the seller-certification game first. In the separating equilibrium, the certification fee F_s^* and the fully revealing signal structure induce a subgame where the effort level is generated by $c'(e) = \overline{\theta} - F_s$. Since it only collects payment from the high seller type, the certifier's expected payoff is eF_s, where e and F_s satisfy $c'(e) = \overline{\theta} - F_s$. The first component, e, is the probability of receiving the payment, which is essentially the probability of producing a high-quality asset. The second component represents the amount of the certification fee that induces e. The certifier's goal is to induce an effort level that maximizes $e\left(\overline{\theta} - c'(e)\right)$. A larger e induced in the game has two impacts. On the one hand, the seller is more likely to be a high type and therefore certify, which could raise the certifier's payoff. On the other hand, to generate a larger e, the certifier has to charge a lower fee. The tradeoff between these two impacts can be seen from the following condition.

$$\frac{dU_c}{de} = \left(\overline{\theta} - c'(e)\right) - c'(e)\,\varepsilon_{c'(e)} \tag{1.13}$$

Whether to charge an expensive fee and generate a low e or charge a cheap fee and generate a high e depends on $\varepsilon_{c'(e)}$. This is reasonable because the elasticity of $c'(e)$ reflects how effectively the effort level can be incentivized through changing the signal scheme or the fee. If

$\varepsilon_{c'(e)}$ is small, $c'(e)$ is not very responsive to a change in e. Only a slight change in F_s is needed to change the effort level by one percent; thus, the second effect is small, and the first effect is dominant between the two. This leads the certifier to charge a cheap fee in order to induce a high effort level. However, it is the opposite if $\varepsilon_{c'(e)}$ is large. In this scenario, the effort level is not responsive to a change in F_s, and the second effect is large. The certifier will charge an expensive fee, which generates a low effort level.

In the pooling equilibrium of the seller-certification model, the certifier is able to serve both seller types with an inflated signal scheme ($\pi_s\left(G|\overline{\theta}\right) = 1$ and $0 < \pi_s\left(G|\underline{\theta}\right) < 1$). Here, $U_c = \pi_s\left(G|\underline{\theta}\right) p_G$, which is equal to the gain of the low type from obtaining certification. The first-best effort level satisfies $c'(e) = \overline{\theta}$. However, here the noisy signal structure distorts the marginal benefit of exerting effort. The amount of distortion can be decomposed into two parts. One part is the additional gain of the low seller type, which is $\pi_s\left(G|\underline{\theta}\right) p_G$. This reduces the seller's incentive to invest in production and is equal to $\frac{e\pi_s(G|\underline{\theta})\overline{\theta}}{e+(1-e)\pi_s(G|\theta)}$. The other part is $\overline{\theta} - p_G$, which is the reduction in the gain of the high seller type. This further reduces the seller's production incentive and is equal to $\frac{(1-e)\pi_s(G|\underline{\theta})\overline{\theta}}{e+(1-e)\pi_s(G|\theta)}$. Combining all these together, a fraction, e, of the total distortion goes to the first part, and the remaining goes to the second part. The total distortion is $\overline{\theta} - c'(e)$; thus, the gain of the low type is $e\left(\overline{\theta} - c'(e)\right)$. From here, it is clear to see that the certifier has the same expected payoff as in the separating equilibrium. The intuition of how the induced effort level depends on $\varepsilon_{c'(e)}$ can be explained in a similar way.

In the buyer-certification model, the information design problem is very different. I will focus on the interior solution ($\pi_b\left(G|\overline{\theta}\right) = 1$ and $0 < \pi_b\left(G|\underline{\theta}\right) < 1$) of the type-1 equilibrium. Here, the certification fee is essentially equal to the buyer's valuation of certification. This is the expected amount of payment to a low-type seller that can be avoided with certification. In this equilibrium, both the high type and the low type charge $p = e\overline{\theta}$. Nevertheless, the low type is only able to receive the asking price with probability $\pi_b\left(G|\underline{\theta}\right)$; as a result, the amount of money avoided paying to the low type is $\left(1 - \pi_b\left(G|\underline{\theta}\right)\right) e\overline{\theta}$. The high type seller always trades at a price $e\overline{\theta}$. In other words, the amount of money which is not collected by the low type is essentially the additional gain of the high type. $(1-\pi_b(G|\underline{\theta}))\, e\overline{\theta}$ is also the marginal benefit of exerting effort, which should be equal to the marginal cost $c'(e)$ in equilibrium. Moreover, the low-type good is produced with probability $1-e$. Altogether, the expected amount of loss averted is $(1-e)(1-\pi_b(G|\underline{\theta}))\, e\overline{\theta} = (1-e)\, c'(e)$. This is also the information rent left to the buyer. The certifier will induce an effort level that maximizes $(1-e)\, c'(e)$. Here, there are two effects of changing e. On the one hand, a higher e reduces the chance of producing a low quality good, which could reduce the buyer's valuation for certification. On the other hand, a higher e suggests a higher marginal cost of production in equilibrium. This indicates that the additional gain of the high type would be improved in the subgame; the low type would obtain a lower payoff, and the buyer could save more from obtaining certification. Equation (1.14) summarizes these two effects.

$$\frac{dU_c}{de} = -c'(e) + \frac{1-e}{e} c'(e)\, \varepsilon_{c'(e)} \tag{1.14}$$

If $\varepsilon_{c'(e)}$ is small, the second effect changes very slowly, and it is small in magnitude for a large e. This suggests that there is not much change in the buyer's saving with certification, even if there is a large change in the effort level. The first effect is the dominant one between the two. Therefore, it is desirable for the certifier to inflate the rating scheme further, which results in a low e_b^*. However, it is the opposite if $\varepsilon_{c'(e)}$ is large. In this scenario, the second effect changes very rapidly, and it is large even at a high effort level. Here, the change to the buyer's saving with certification drops so drastically (as e decreases) that it is more desirable for the certifier to incentivize a high e_b^*. Consequently, when $\varepsilon_{c'(e)}$ is small, the equilibrium effort is less distorted.

In Figure 1.1, I present the equilibrium effort levels in both the seller-certification model and the buyer-certification model (interior solution) when the elasticity of $c'(e)$ is small. Here, I use $c(e) = \frac{2}{3}e^{\frac{3}{2}}$ and $\overline{\theta} = 1$. The elasticity of the marginal cost function is $\frac{1}{2}$. The intersection of the blue dotted line and the red line depicts the equilibrium effort level in the buyer-certification model. Moreover, the intersection of the green dashed line and the red line depicts the equilibrium effort level in the seller-certification model.

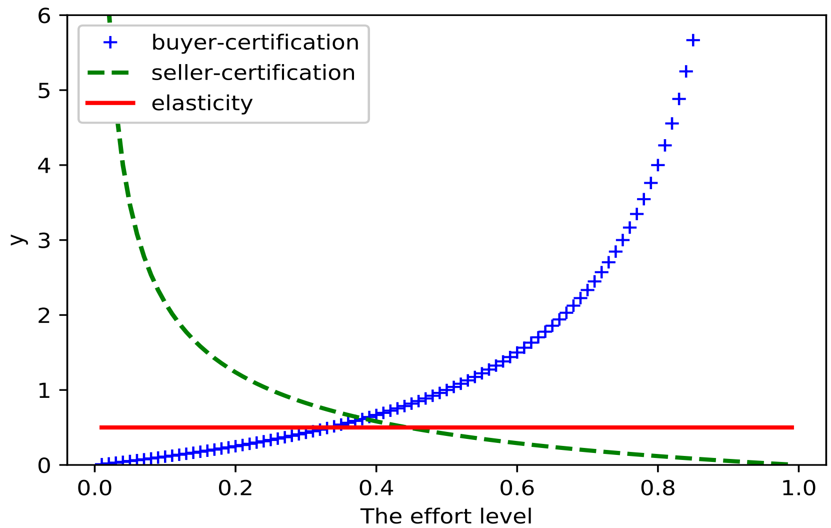

Figure 1.1. Comparison of the equilibrium effort levels when the elasticity of $c'(e)$ is small

Here, the seller-certification model generates a higher equilibrium effort level. When $\varepsilon_{c'(e)}$ is small, the seller-certification model leads to a higher social welfare level than the buyer-certification model.

On the contrary, in Figure 1.2, I show an example of an elastic $c'(e)$. Here, I use $c(e) = \frac{1}{3}e^3$ and $\overline{\theta} = 1$. The elasticity of the marginal cost function is 2. As in the previous graph, the intersection of the blue dotted line and the red line depicts the equilibrium effort level in the buyer-certification model. Also, the intersection of the green dashed line and the red line depicts the equilibrium effort level in the seller-certification model.

In Figure 1.2, we can see that the buyer-certification model creates a higher effort level. The buyer-certification model generates a better social surplus when $\varepsilon_{c'(e)}$ is large.

I can also compare the certifier's payoff in these two business models. In the type-1 equilibrium of the buyer-certification model, the certification fee is equal to the expected amount of payment (to the low seller type) avoided with certification. Specifically, $F_b^* = \left(1-e_b^*\right) c'\left(e_b^*\right) < e_b^*\left(\overline{\theta} - c'\left(e_b^*\right)\right)$. Nevertheless, in the seller-certification model, the certifier

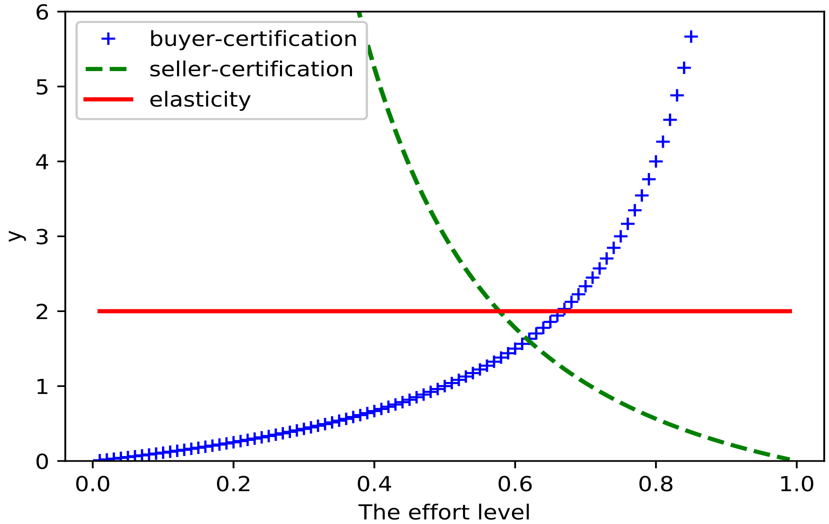

Figure 1.2. Comparison of the equilibrium effort levels when the elasticity of $c'(e)$ is large

ends up with payoff $e_s^*\left(\overline{\theta}-c'(e_s^*)\right)$ where $e_s^*c''(e_s^*)=\overline{\theta}-c'(e_s^*)$. As e_s^* maximizes the value of the function $e\left(\overline{\theta}-c'(e)\right)$, the certifier is always better off in the seller-certification game. This result is established because of information asymmetry. In the seller-certification model, the seller's certification choice is made conditional on the exact type, whereas in the buyer-certification model, the buyer's certification choice is based on the expectation of the seller type. As the buyer certifies to avoid purchasing the low quality good, the valuation he places on certification is less than the seller's valuation of certification.

In terms of other market participants, the buyer obtains a payoff of zero in both certification models. The seller is the residual claimant. In the seller-certification model, this amount is $e_s^*c'(e_s^*)-c(e_s^*)$; in the buyer-certification model, it is $e_b^*\overline{\theta}-c\left(e_b^*\right)-\left(1-e_b^*\right)c'\left(e_b^*\right)$. When $e_b^*\geq e_s^*$, it is clear that the seller is better off in the buyer-certification game. However, this comparison becomes ambiguous when $e_b^*<e_s^*$. If $e_s^*-e_b^*$ is relatively small, the seller could still be better off when the buyer initiates certification.

1.7 Extensions

In this section, I discuss some possible variations in the model setting and how they potentially influence the current results.

1.7.1 Variation in the Timeline

One possible variation is the timing of certification. In the seller-certification model, the seller may certify first, and then set the price contingent on the signal realization. This setting is practical in the real marketplace as the seller can adjust the asking price based on the actual rating outcome. Even after I incorporate this change in the model, my results remain robust. There are still two possible equilibria, a separating one where only the high type certifies, and a pooling one where both types certify. I summarize the results in the next proposition.

Proposition 5 *Suppose the seller can set the asking price conditional on the certification outcome. There are two equilibrium types, and both of them yield the same equilibrium outcome as the two equilibria characterized in Proposition 2.*

1.7.2 Choice of the Business Model

Another possible variation is that the certifier is able to choose the party to work with. To incorporate this issue, I allow the certifier to commit to work with the party it prefers at the beginning of the game. If it chooses to work with the seller, the game extends as the seller-certification game. If it chooses to work with the buyer, the game extends as the buyer-certification game. Here, the certifier compares its payoff in the two models and chooses the one with a better payoff. Therefore, I can directly apply the results from the previous sections. Specifically, the certifier always has a higher payoff in the seller-certification game; thus, it will commit to work with the seller. This is also consistent with the current financial market observation that the majority of business is done in the "issuer-pays" mode.

However, from Theorem 1, we know that the social welfare comparison between the seller-certification model and the buyer-certification model depends on $\varepsilon_{c'(e)}$. When $\varepsilon_{c'(e)}$ is small, the seller-certification model generates a higher social surplus. Here, the policy maker's interest of improving social welfare and the certifier's incentive of making profit are aligned. When $\varepsilon_{c'(e)}$ is large, the buyer-certification model induces a higher social surplus. However, the policy maker's interest and the certifier's incentive work in opposite directions. This suggests that it is not always optimal to let the certifier pick the business mode itself. Under some conditions, the policy maker's interest and the certifier's incentive are not aligned. The policy maker could improve social welfare by changing the current business mode to the "investor-pays" mode when the marginal cost of production is very elastic.

1.7.3 A Two-Part Tariff Contract

In the seller-certification game, another possibility is that the certifier uses a two-part tariff contract. Suppose that the certifier can charge an entry fee before the seller exerts any effort. After the seller knows the good's quality, there is an additional fee to have the good certified. The research question here is how the certifier designs the certification mechanism if a two-part tariff contract is allowed.

To answer this question, I can utilize the results in Section 1.4. Since the certifier can use the two-part tariff to extract all the surplus in the game, it will induce the mechanism that maximizes the total payoff, which is $e\overline{\theta} - c(e)$. Therefore, the certifier has the incentive to generate the first-best effort level e^*, which satisfies $c'(e^*) = \overline{\theta}$. This shows that the marginal benefit of exerting effort is $\overline{\theta}$. Accordingly, the signal mechanism must be fully revealing, and the additional fee to certify is set to be zero. The certifier uses the entry fee to extract all the surplus generated in the game, which is $e^*\overline{\theta} - c(e^*)$. Here, the seller and the buyer both have zero payoffs.

Surprisingly, allowing a two-part tariff contract can restore the effort level to the first-best. The intuition is that the certifier fully internalizes the gain and the cost of production so that the certifier's incentive is perfectly aligned with the general interest of improving social welfare.

However, this requires a two-part tariff to be implemented before the seller produces the good. If the fixed fee is not forced to be collected before the production process, the seller has an incentive to deviate. Suppose the fixed fee and the additional fees are F and 0 respectively. The seller's payoff is $e\overline{\theta} - F - c(e)$ for an effort level e. However, $e\overline{\theta} - F - c(e) \leq e\left(\overline{\theta} - F\right) - c(e)$. The seller has a strong incentive to produce the good first before she pays the fixed fee, as such deviation generates a strictly better payoff when $e < 1$. She pays the fixed fee only when the good is a high quality one. The corresponding effort level satisfies $c'(e^{*}) = \overline{\theta} - F$, which is the same as in the separating equilibrium of the seller-certification model.

The seller-certification model leads to a better social outcome if a two-part tariff contract is feasible. However, this kind of contract is very difficult to enforce since the production process is hard to monitor in reality. The two-part tariff contract itself cannot guarantee that the seller takes the initiative to pay the fixed fee before she produces the good. If the seller has the flexibility to choose the time to certify, the two-part tariff contract will not work in the way it is designed to be. The induced effort level will still be below the first-best.

1.8 Conclusion

In markets with severe moral hazard problems in the production of goods or services, both market participants demand certification. In this paper, I study the roles of certification —revealing information and incentivizing the seller —in two models: the seller-certification model and the buyer-certification model. The tradeoffs between these two roles are different in the two models, thereby leading to different market transparency levels and economic outcomes. The question of who pays for certification is crucial to the market outcome.

Specifically, the paper adopts an information design approach to capture the two roles of certification by allowing the certifier to choose the signal structure freely. In particular, I show that the certifier may strategically produce noise when serving the seller, which is consistent with the empirical observation that the "issuer-pays" business model tends to release inflated ratings.

However, which business model creates a more efficient outcome depends on the induced effort level, which is determined by the elasticity of the marginal cost of production. This function determines the effectiveness of incentivizing the seller to exert effort through altering the signal scheme. If the marginal cost of production is very inelastic, the seller-certification model creates a more efficient outcome, whereas the buyer-certification model does better when the marginal cost of production is very elastic.

Moreover, I also find that the certifier always earns a higher profit in the seller-certification model. This prediction is consistent with the observation that the "issuer-pays" model dominates the current financial rating market. The welfare analysis in this paper provides a novel policy implication concerning the business mode of certification. When the marginal cost of production appears to be very elastic, the buyer-certification mode does a better job of incentivizing effort. In this situation, policy interventions could be beneficial to society. These results complement the current discussion of which business mode should be employed in the certification market.

The models of the seller-certification and the buyer-certification are two representative examples of the certification mechanisms. It is a subclass of a more general class of mechanisms where the certifier can freely charge both sides of the market. The optimal certification mechanism design remains a significant and open future research question.

Chapter 2

Information Disclosure in the Dynamic Market for Lemons

This paper studies the incentive of a long-run seller to disclose previous offers in a dynamic market for lemons and identifies the impact of voluntary disclosure on the market information structure. Comparing the optional disclosure model with the models of mandatory disclosure and mandatory non-disclosure, I find that there is a novel set of equilibria generated by allowing flexibility in the disclosure option. In this new class of equilibria, the seller adopts a threshold rule and only discloses rejected price offers above the threshold to future buyers. In the two-period model, I show that the optional disclosure model could induce a strictly higher social surplus than models with mandatory restrictions on the disclosure policy. Policy makers may not necessarily enforce mandatory disclosure or concealment of every past offer in the real marketplace. Moreover, they can adopt a non-disclosure policy for lower prices and a voluntary (or mandatory) disclosure policy for higher prices to enhance trading efficiency.

2.1 Introduction

In real life trading, buyers submit offers based on how much they know about sellers' willingness to sell. One important source of such information is previous offers that were not accepted. When a seller has the option to disclose past buyers' offers, whether to disclose or not becomes a practical question to consider. In the real estate market, for instance, the listing agent has the option to disclose past offers that were rejected while conducting business with the current selling agent. In the labor market, an employee's past salary is always of great concern. Potential employees must choose carefully whether to disclose this piece of information during wage rate bargaining.[1] Moreover, even in the trading of professional soccer players, soccer club managers face the problem of what information they should reveal regarding rejected transactions. For example, they could simply reveal the final result, or they could discuss the exact terms of the contract they had rejected.[2] All of these practical examples have a common, central question: should a seller disclose past prices to current buyers when they have a chance to do so?

The above question is complicated but also meaningful to answer. The answer not only helps us to understand a seller's incentive to make the disclosure decision, but also has great significance for designing or regulating market structure regarding to what extent the observability of past offers is at the discretion of the seller. In the economics literature, much attention has been paid to games in which the observability of past offers is exogenously determined in the model. For example, [23] and [18] have studied the bargaining problem in a dynamic market for lemons, where past offers were either always observable or always unobservable to buyers. This model setting is appropriate when the observability of past prices is a natural element in the market structure. However, there are many situations, such as those described earlier, where the observability of past offers is up to the seller. A strategic seller can choose whether or not to disclose past offers to her buyers. On the one hand, disclosure may act as a good signal and reduce the amount of asymmetric information, which would facilitate trade. On the other

[1] In New York City, a new salary history law became effective on Oct. 31st, 2017. It is now illegal for public and private employers of any size to require information about job applicants' salary histories. Disclosure of salary history is not mandatory in wage bargaining. (https://www1.nyc.gov/site/cchr/media/salary-history.page)

[2] Transactions of soccer players are often released by Britain's free press. Here, sometimes soccer club managers choose only to reveal the failure of a particular transaction, as when Manchester United rejected a Wayne Rooney bid from Chelsea (See: http://www.telegraph.co.uk/sport/football/teams/manchester-united/10184651/Wayne-Rooney-bid-from-Chelsea-rejected-by-Premier-League-champions-Manchester-United.html). At other times, the details of the rejected offers are disclosed, as when Napoli rejected the offer for Gonzalo Higuain from Atletico Madrid (http://soccerlens.com/arsenal-news-atletico-madrid-see-e60m-higuain-bid-rejected-by-napoli/196937/).

hand, withholding information could preclude players from speculating in trades or from taking opportunistic actions.

This paper studies the seller's optional disclosure question by considering the interaction between a strategic long-run seller (she) and two short-run buyers (he) for a single item in a two-period game. In each period, a single buyer has the chance to trade with the seller, and there is asymmetric information about the quality of the item. The key feature of my paper is the seller's ability to decide whether or not to disclose a past offer to the current buyer. In the model, if there is no trade in period 1, the seller has the freedom to decide whether to reveal information about the unsuccessful offer in period 1 before the buyer in period 2 submits an offer. The observability of the previous offer is endogenous in the model, and it is a decision made by the seller.

In the paper, I characterize the equilibria in a two-period model and then extend the results to an infinite-horizon game. Buyers form beliefs about the quality of the good based on the available information and submit offers according to those beliefs. The available information to the buyers includes the disclosed offers and the total number of rejected offers. With the disclosure option, the seller has more flexibility in sending out information to the buyers; as a result, she is able to influence the buyers' beliefs through disclosing or withholding past offers. Therefore, the equilibrium beliefs in the optional disclosure game are quite different from both the mandatory disclosure model and the mandatory non-disclosure model in the literature, as are the equilibrium dynamics and social welfare. Comparing my model with the models of mandatory disclosure and mandatory non-disclosure ([23], [18]), I find that there is a new set of equilibria generated by allowing optional disclosure. In this new class of equilibria, the seller implements a threshold rule for disclosing past offers. Disclosing a high past offer can be seen as sending a positive signal to the current buyer while disclosing a low past offer can be seen as sending a negative signal. In the two-period game, the seller selects a disclosure threshold in period 2 and only discloses the rejected price offer if it is above this threshold. She adopts this threshold rule as the criterion to evaluate whether or not to disclose the previous offer and sends a positive signal to the current buyer when this criterion is fulfilled. Only an offer above the threshold is considered to be adequate for sending a positive signal. On the equilibrium path, there is no disclosure, and no signal is sent from the seller's side.

In the two-period model, I further show that the welfare-maximizing equilibrium of the optional disclosure model could generate a strictly higher social surplus than any equilibrium of the model where the mandatory non-disclosure restriction is imposed. In addition, any equilibrium of the optional disclosure model yields a higher social surplus than the equilibrium of the mandatory disclosure model. Policy makers may not necessarily enforce mandatory disclosure or concealment of every past offer in the real marketplace. In the mandatory disclosure or mandatory concealment model, past offers are naturally observable or unobservable; the seller mechanically reveals everything or nothing to the buyers. However, providing the seller the disclosure option gives her the freedom to take different disclosure actions for different price offers. The seller has perfect control over what messages are transmitted to the buyers in the optional disclosure model. In equilibrium, she only conveys positive messages to the buyers, and by doing so potentially improves the buyers' beliefs. Therefore, optional disclosure could lead to a higher trading price and a better economic outcome. In the optional disclosure model, finer information could be transmitted in the market, and trading efficiency could be improved.

The results presented in this paper are related to many observed behaviors in the real marketplace. For example, listing agents often do not disclose rejected offers when conducting business with current buyers. This is consistent with the model's predictions. Sometimes listing agents may selectively hint that they have rejected prices above a certain level, which is also relevant to the seller's strategy of only disclosing offers above a certain threshold in my model. According to some headhunters, it is not advisable to disclose one's salary history during wage bargaining.[3] My results indicate that the headhunter's advice is sometimes reasonable. Reporting one's previous salaries is definitely not a wise action to take if they are below potential employers'

[3] See http://www.vault.com/blog/salary-and-benefits/should-you-disclose-your-salary-in-an-interview and http://www.salary.com/disclosing-salary-history/

expectations.

When designing the information structure in a market, policy makers should contemplate the disclosure rule of past offers. The paper suggests that providing the disclosure option to the seller could potentially generate higher social welfare in the market. More importantly, policy makers can apply selective restrictions on the disclosure rule to achieve the welfare level in which they are interested. Adopting a non-disclosure policy for lower prices and voluntary (or mandatory) disclosure policy for higher prices could potentially enhance trading efficiency.

The rest of the paper is organized as follows. Section 2.2 briefly reviews the related literature. Section 2.3 introduces the model setting. Section 2.4 discusses a two-period game. Section 2.5 discusses generalizations of the model and related policy implications. Section 2.6 explores extensions of the model. Finally, in Section 2.7, I present my concluding remarks. The formal proofs are in Appendix B.

2.2 Literature Review

Starting in the 1980s, a great number of researchers have studied questions about information disclosure. Under costless revelation, [19] and [35] argue that an informed seller will always deliberately reveal her private information to buyers if the information is ex-post verifiable; therefore, the problem of asymmetric information is eventually solved. The main difference between my paper and the previous literature on this topic is the disclosure content. In the literature, the content usually refers to signals or evidence of a seller's type, which is exogenously determined, whereas the content I consider is the buyers' behaviors in past transactions, which are endogenous in the model. As historical prices are the same for every seller type, the disclosure content is not type-dependent. Disclosure itself is not a perfect signal of the seller's type, and disclosure is not a dominant strategy for the highest seller type. Therefore, it is not necessary to have the unraveling result in my model. Another difference is that I consider a dynamic model with a long-run seller and multiple short-run buyers. In this setting, the information regarding past offers becomes critical to the current buyer. The information disclosure literature considers only static or dynamic models with long-run players, where buyers are aware of previous offers.

There is a strand of the bargaining literature related to information transparency. Both [23] and [18] study the bargaining problem in a dynamic market for lemons, where past offers are either always observable or always unobservable to buyers. Their results suggest that observable offers tend to induce a market breakdown or a bargaining impasse, whereas trade is eventually reached with private offers. [28] compares three information structures under which sellers and buyers randomly and bilaterally match. He finds that market efficiency is not monotonic in the amount of information available to buyers. [27] consider a Coasian bargaining environment and show that the unobservability of past negotiations leads to lower prices and faster trading. For this strand of literature, the intuition is that the seller has a stronger incentive to reject high offers to signal a high-quality commodity when past offers are observable. This contributes to delay in trade and gives rises to trading inefficiency.

Similar to the settings of [23] and [18], my model also incorporates a dynamic market structure. However, in my work, the seller has the option to disclose past rejected prices. The disclosure option is endogenous, as opposed to exogenous, as found in the literature; this completely changes the information structure of the model. In the optional disclosure model, the seller has more flexibility in signaling her type, and she is able to influence the outcome of the game through the signaling effect on buyers' beliefs. Moreover, the questions I discuss in my paper are fundamentally different from the prior work in this area. The literature focuses more on comparing the two market structures, i.e., full information versus no information. In my work, I concentrate on the seller's choice of whether to disclose previous offers. However, my paper does relate to this strand of research in terms of its welfare analysis and policy implications. I find that the optional disclosure model can generate a higher social welfare level under some circumstances.

Besides the information disclosure literature, this paper also relates itself to the dynamic signaling literature in which the adverse selection problem appears ([37], [49], [29], [11]).

Rejecting and disclosing high offers can be seen as a behavior of signaling a high type.

2.3 The Model and Preliminaries

The model features a two-period bargaining game between a long-run seller (she) and a sequence of short-run buyers (he). Time is discrete and indexed by $t = 1, 2$. There is a single trading item, and the quality of the item, q, is the seller's private information. Assume that q is distributed uniformly on the interval $\left[\underline{q}, \overline{q}\right]$, and let $G_1(q)$ denote this distribution. $G_1(q)$ is common knowledge. Given q, the seller's reservation value of the good is αq with $\frac{1}{2} < \alpha < 1$. [4] The buyers are homogeneous in the sense that they share a common valuation the good, which is equal to q. I label buyer t as the buyer who trades in period t.

In period t, buyer t interacts with the seller and proposes an offer p_t for the good. The buyer can only make one offer in a period, and he is only able to trade in that period. If the offer is accepted by the seller in period 1, the game ends; otherwise, the game continues to period 2. In period 2, the seller can choose to disclose p_1 that she has rejected to buyer 2 before he submits his offer. The disclosure content can be any subset of the offer history. Disclosure is costless and verifiable, which is to say, "the talk is not cheap". If she discloses nothing, the only thing that the current buyer is aware of is that all past offers were rejected.

If there is no trade after the two periods, a type-q seller will receive the continuation value Aq ($\alpha \leq A \leq 1$) in period 2.[5] In period 2, a trade will occur if and only if $p_2 \geq Aq$. I maintain the assumption that Aq is no less than αq. Let Δq denote $\overline{q} - \underline{q}$. For the two-period model, I assume that Δq is large enough so that the highest type is not traded in these two periods. A general analysis of trading the highest type is discussed in the infinite-period model. The seller discounts the payoff across all periods according to a discount factor δ, and all players are risk

[4] In a one-period game, the marginal increase in average quality if trading with an extra high-type seller is $\frac{1}{2}$ while the cost is α. If $\alpha \leq \frac{1}{2}$, the marginal benefit outweighs the cost for any q, and I have $\frac{\underline{q}+q}{2} \geq \alpha q$. Here, buyer 1 will simply offer $\alpha\overline{q}$. As this result is relatively trivial, I exclude it in my model.

[5] In [18], they work with a two-period model under a similar setting when there is no trade. The authors discuss the value of A, which represents the efficiency loss at the trading deadline. Here, the value of A does not change the equilibrium characterization. A discussion of A is not a focus of this paper.

neutral. The model's timeline is presented below.

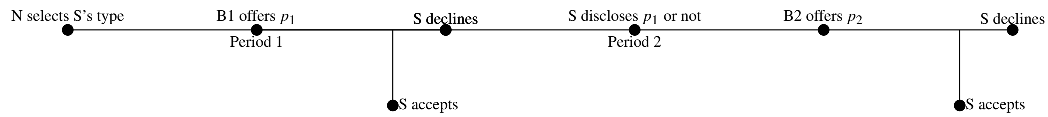

Figure 2.1. The timeline of the two-period game

Additionally, buyer 2 has a belief about the seller's type, and this is described by a cumulative distribution function $G_2(q)$. Suppose the seller's type is q, and trade happens in period t. The seller's payoff is $\delta^{t-1}(p_t - \alpha q)$, buyer t's payoff is $q - p_t$, and the other buyer ends up with a payoff of zero. Let p^1 be the price history in period 2, and P^1 be the set of all possible price histories in period 2.

The seller's strategy is $\left\{a^1(q,p_1), d^2(q,p^1), a^2(q,d^2,p_1,p_2)\right\}$. $a^1(q,p_1)$ captures the probability that the type-q seller accepts p_1; the disclosure rule, $d^2(q,p^1)$, maps the seller's type q and the price history p^1 into a subset of p^1; $a^2(q,d^2,p_1,p_2)$ maps the seller's type q, the disclosure content d^2, and the offers p_1 and p_2 into the probability that the seller accepts p_2. $d^2(q,p^1) = \emptyset$ means that the seller discloses nothing. Buyer 1's strategy σ^1 is a probability distribution over $\mathbb{R}_+$. Buyer 2's strategy σ^2 is a mapping from the disclosure content $d^2(q,p^1)$ to a probability distribution of prices over $\mathbb{R}_+$.

The solution concept I use is perfect Bayesian equilibrium (PBE), which requires the strategies of the players to be optimal both on- and off-path. Specifically, in this context, a PBE includes the seller's acceptance rules and disclosure rule, every buyer's pricing strategy and belief. It satisfies the following conditions: (1) The seller's strategy $\left\{a^1, d^2, a^2\right\}$ maximizes her payoff given the buyers' pricing strategies and belief updating process. (2) Buyer t's pricing strategy σ^t maximizes his payoff conditional on his own belief, the seller's acceptance and disclosure rules, and the other buyer's pricing strategy. (3) Buyer t's belief is updated (whenever possible) according to the Bayes' rule based on the seller's and all other buyers' strategy. From now on, I will refer to PBE as the "equilibrium".

I present a preliminary result first. As in other dynamic games in continuous-type settings, the seller's acceptance rule follows a cutoff rule. In every equilibrium, for a price offer p_t, there is a cutoff type q_t, such that the seller with a type below the threshold q_t accepts p_t, while the seller with a type above the threshold rejects it. This is the *skimming property* in the bargaining literature, and it is summarized in Lemma 1.

Lemma 1 *In every equilibrium, the seller's acceptance rule follows a cutoff rule. If a type-q seller accepts offer p_t with positive probability, then any lower seller types accepts offer p_t with probability* 1.

Lemma 1 is one of the standard results for bargaining games in continuous-type space settings. The supremum over all types accepting an offer is called the cutoff type. To make the statement concretely, I follow the convention that the cutoff type accepts the offer with probability 1. From Lemma 1, the equilibrium can be described by a triple $(d^t, p_t^*, \widehat{q_t})$ where $t = 1,2$. Here d^t is the seller's disclosure rule at the beginning of period t, and I let $d^1 = \emptyset$ to make notation consistent. p_t^* is the equilibrium price made by buyer t, and $\widehat{q_t}$ is the cutoff type accepting offer p_t^*. A type-q seller trades in period t where $\widehat{q_t} \geq q > \widehat{q_{t-1}}$.

2.4 Equilibrium Construction and Welfare Analysis

In this section, I present my main result. The analysis is focused on the case where discounting is small. When the seller is patient enough, she is less willing to accept a high offer today. She would rather reject the offer and disclose it as a good signal in the next period, which could change the seller's incentive of accepting an offer. This is different from the models of mandatory disclosure or mandatory non-disclosure, in which the seller does not have the disclosure option. I also briefly characterize the equilibrium when discounting is large in the two-period model.

2.4.1 Equilibrium Construction

I present the characterization of equilibrium via a series of lemmas and propositions. Lemma 2 presents a general statement about the type of equilibrium in the optional disclosure model.

Lemma 2 *In every equilibrium, buyer* 1 *and buyer* 2 *must play pure strategies.*

Lemma 2 highlights that in every equilibrium the buyers cannot play mixed strategies. Suppose this is not true, and buyer 1 is now randomizing among a set of offers. Among these offers, higher ones induce higher cutoff types in period 1. If all of buyer 1's offers are not disclosed on the equilibrium path, buyer 2 will have a non-degenerate belief of the cutoff type in period 1. For buyer 2, it is possible to trade with a high seller type or a lower one. Taking these possibilities into account, buyer 2 submits his offer. For a high seller type, given the disclosure option, she would have the incentive to justify her type by disclosing the highest offer she could receive in period 1. By doing so, she could signal an extra positive message to buyer 2 about her type. This suggests that the seller rejected the offer not because the offer was too low, but because her type is too high. Moreover, instead of having a non-degenerate belief of all possible cutoff types, buyer 2 would update his belief towards a higher cutoff type and submit a better price. Applying the same intuition, I conclude that there would be unraveling for buyer 1's offers. Buyer 2 would update his belief conditional on the disclosed offer, and best respond to it. In this situation, given buyer 2's strategy and the seller's disclosure rule, there is no reason for buyer 1 to randomize. Buyer 2 will also play a pure strategy if buyer 1 is not randomizing.

Applying the result in Lemma 2, I am able to derive other equilibrium characteristics. Lemma 3 indicates that there can be common actions in disclosure.

Lemma 3 *In every equilibrium, if a type-q seller prefers disclosing offer p_1 when p_1 is rejected, every seller type rejecting p_1 prefers disclosing p_1. In addition, if the type-q seller strictly prefers disclosing p_1, and p_1 is the equilibrium offer, then $p_1 \in d^2(q, \{p_1\})$ for any type q rejecting p_1.*

Lemma 2 allows me to focus on the case where buyers are only playing pure strategies. If a particular seller type prefers to disclose p_1 rather than conceal it, her continuation value of disclosing p_1 is higher than concealing p_1. This suggests that buyer 2 would submit a higher price when p_1 is revealed. Every type must have the same incentive to disclose this offer if she rejects it. Moreover, there is pooling in disclosure actions for the equilibrium offer. Otherwise, if a low type and a high type took diverging disclosure actions when rejecting the equilibrium offer, a future buyer could distinguish them according to their actions, and thus serve them differently. This creates an incentive for a low seller type to take the same action as a high seller type. Lemma 3 states that different seller types share the same preference of disclosing or concealing p_1 when p_1 is rejected. Correspondingly, if p_1 is disclosed, buyer 2 will form a belief of the cutoff type that discloses p_1.

There are two extreme disclosure rules, i.e., always disclose and always conceal. Lemma 4 states that always concealing past offers is not an equilibrium disclosure rule.

Lemma 4 *In every equilibrium,* $d^2\left(q, p^1\right) = \emptyset$ *for any* $p^1 \in P^1$ *and any* q *is not an equilibrium disclosure rule.*

This result is not surprising. For any offer p_1 higher than the equilibrium offer, the seller is inclined to reveal this offer if she rejects it. This can enhance buyer 2's belief about her type, thereby improving the price in period 2. In conclusion, always concealing any previous offer cannot be the disclosure rule in equilibrium.

In contrast, always disclosing past offers can be the disclosure rule in equilibrium. It can be sustained when buyer 2 has a negative belief of the seller's type under non-disclosure. Specifically, if non-disclosure happens, buyer 2 believes that the seller is of the lowest type. In this case, the seller always has the incentive to disclose the price offer made in period 1; this scenario is formally stated in Proposition 6. For notation purpose, I denote $\delta^* = 1 - \frac{1}{2\alpha}$, and $q_L = \frac{\delta A^2 + (2A-1)(1-\delta)\alpha}{2\delta A^2 + (2A-1)(2(1-\delta)\alpha - 1)}\underline{q}$.[6]

[6] I assume that $q_L \geq \underline{q}$ under the parameter values.

Proposition 6 *In the optional disclosure model, $d^2\left(q,p^1\right)=p^1$ for any $p^1\in P^1$ and any q can be the disclosure rule in equilibrium if buyer 2 has the belief that $G_2\left(\underline{q}\right)=1$ when $d^2\left(q,p^1\right)=\emptyset$. The cutoff types are $\widehat{q_2}=\frac{A}{2A-1}\widehat{q_1}$ and $\widehat{q_1}=q_L$, respectively.*

There are other equilibria in which buyer 2 has a more optimistic belief when non-disclosure happens besides the equilibrium with full disclosure. From Lemma 2, we know that in any equilibrium, both buyers must play pure strategies; thus, there is no randomization of prices in equilibrium. To make the equilibrium characterization concrete, I assume that every seller type either discloses the equilibrium offer or conceals it, i.e., there is no randomization in disclosing the equilibrium offer.[7] From Lemma 3, we know that all seller types share the same incentive of whether or not to disclose a rejected offer. Accordingly, buyer 2 will form a belief that the seller type is above a certain cutoff type based on what he knows about buyer 1's price. Applying the results derived previously in the lemmas, I can characterize a novel set of equilibria in the optional disclosure model.

In the new set of equilibria, the common feature is that all seller types have a common target price in every period. This price is the disclosure threshold. When an offer below or equal to this target price is made, the offer will not be disclosed after it is rejected. However, when an offer above this price is made, the seller has the incentive to disclose it after she rejects it. The seller takes different actions for different offers. She discretely discloses the information that is positive in some sense. Buyer 1's equilibrium price offer p_1^* matches the value of the seller's disclosure threshold. The seller of type below the cutoff $\widehat{q_1}$ is willing to accept p_1^* in period 1. As for buyer 2, he has a consistent belief of the seller's type in equilibrium; he chooses p_2^*, which

[7] I make this assumption to simplify the equilibrium characterization. Drawn from Lemma 2, I conclude that there are only equilibria in which buyer 1 and 2 play pure strategies. The equilibrium offer is perfectly predicted by buyer 2. If on the equilibrium path a seller type randomizes between disclosing and withholding p_1^* with probability α and $1-\alpha$, it must be the situation in which buyer 2 will offer the same price whether p_1^* is revealed or not. I can construct an equilibrium for any arbitrary $\alpha\in(0,1)$, since the exact probability that the seller mixes between disclosing and concealing the equilibrium offer does not affect the equilibrium construction. In the last part of the proof of Proposition 7, I show that if this assumption is removed, for any equilibrium in which some seller types randomize between disclosing and concealing p_1^*, there is an equilibrium generating the same equilibrium prices and cutoff types. In that equilibrium, all seller types conceal p_1^* on the equilibrium path, and other parts of the strategy profile remain unchanged. Therefore I make this assumption.

results in a cutoff type $\widehat{q_2}$ that maximizes his payoff. Moreover, there are multiple equilibria. Proposition 7 summarizes the details.

Proposition 7 *There is another class of equilibria.*

(1) Buyer 1 *submits* p_1^* *and trades with the seller type below* $\widehat{q_1}$. *The seller's disclosure rule is* $d^2(q,\{p_1\}) = \emptyset$ *if* $p_1 \le \widehat{p_1}$ *for any* q. *Here,* p_1^* *matches the value of* $\widehat{p_1}$ *in equilibrium. For any* $p_1 > \widehat{p_1}$, *any seller type has the incentive to disclose* p_1 *if she rejects it. Buyer* 2 *offers* p_2^* *and trades with the type below* $\widehat{q_2}$ *where* $\widehat{q_2} = \frac{A}{2A-1}\widehat{q_1}$.

(2) There is $\delta^* = 1 - \frac{1}{2\alpha}$ *such that:*

(i) For all $\delta \le \delta^*$, *the cutoff type in period 1 satisfies* $\widehat{q_1} \in \left[q_L, \frac{(2A-1)(1-\delta)\alpha}{\delta A^2+(2A-1)(2\alpha(1-\delta)-1)}\underline{q}\right]$.
(ii) For all $\delta > \delta^*$, *the cutoff type in period 1 satisfies* $\widehat{q_1} \in \left[q_L, \frac{(2A-1)}{2\delta A^2+(2A-1)(2\alpha(1-\delta)-1)}\underline{q}\right]$.

Part (i) of Proposition 7 captures the other class of equilibria. The intuition for Part (i) follows directly from Lemmas 1-4. This two-period game only has equilibria in pure strategies. On the equilibrium path, buyer 2 has a degenerate belief of the cutoff type. For any offer higher than the equilibrium price, the seller has the incentive to disclose the offer to buyer 2 after she rejects it. This disclosure incentive is shared by all types who reject the offer. Disclosure sends a signal to buyer 2 and causes him to update his belief towards a higher cutoff type. However, when offers lower than the equilibrium price are made, the seller does not want to disclose this information; she wants to pretend that the game is still on the equilibrium path.

In Figure 2.2, I illustrate the seller's decision rule in one equilibrium with $\alpha = 0.8$, $A = 0.92$, $\delta = 0.6$, $\underline{q} = 10$, $\overline{q} = 20$. There is a kink point on the line that represents the cutoff type accepting the price offer. This kink point is at the disclosure threshold. This is because the seller has different disclosure actions depending on whether the offer is above or below the disclosure threshold. Below the disclosure threshold, as the price offered in period 1 is concealed, buyer 2's price offer will remain unchanged. The seller is more willing to accept a higher price in period 1, and the cutoff type is more sensitive to changes in p_1. Therefore, the line representing the cutoff type has a steeper slope below the disclosure threshold. Above the disclosure threshold, the seller has the incentive to report high prices to buyer 2, and p_2 is contingent on the reported p_1. In this case, a higher p_1 serves relatively fewer types, and the slope of the line representing the cutoff type becomes flatter. Here, notice that there is no disclosure for the equilibrium offer. This is because buyer 2 maintains a consistent belief of the seller's type in a pure strategy equilibrium. There is no need to convey an extra message, and buyer 2 will not change his belief, even if the seller reveals p_1^*. Furthermore, there is no reason to reveal offers below the equilibrium price, as this would result in buyer 2 updating his belief downwards.

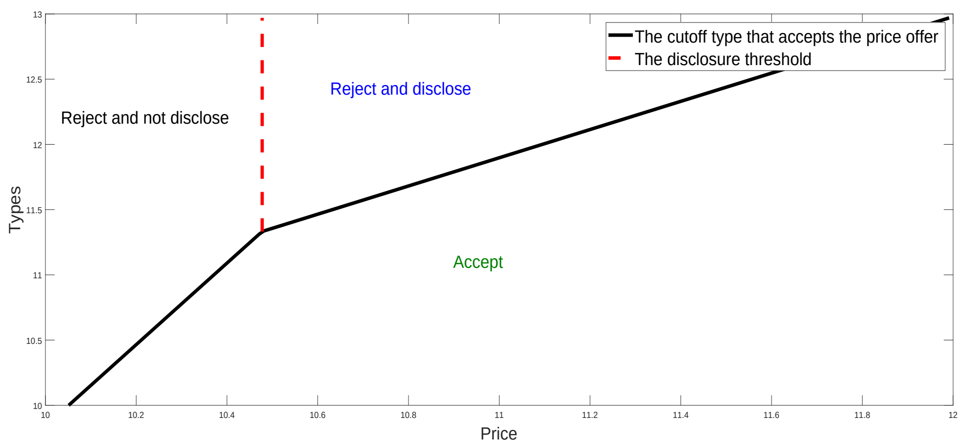

Figure 2.2. The seller's disclosure and acceptance rules

As for buyer 1, the equilibrium price maximizes his payoff given the seller's and buyer 2's strategy. To analyze buyer 1's pricing strategy, I characterize his payoff function, which is $\pi_1(q_1) = \frac{q_1 - \underline{q}}{\Delta q}\left(\frac{q_1 + \underline{q}}{2} - p_1\right)$. p_1 satisfies the equation below:

$$p_1 - \alpha q_1 = \delta(p_2 - \alpha q_1), \quad \text{where } p_2 = \begin{cases} p_2^*, & \text{when } p_1 \leq p_1^* \\ \frac{A^2}{2A-1} q_1, & \text{when } p_1 > p_1^* \end{cases}$$

In Figure 2.3, I present buyer 1's payoff function when the disclosure threshold is equal to 10.067, for the same parameters as before ($\alpha = 0.8$, $\delta = 0.6$, $A = 0.92$, $\underline{q} = 10$, $\overline{q} = 20$). Note that $\delta > \delta^*$ holds.

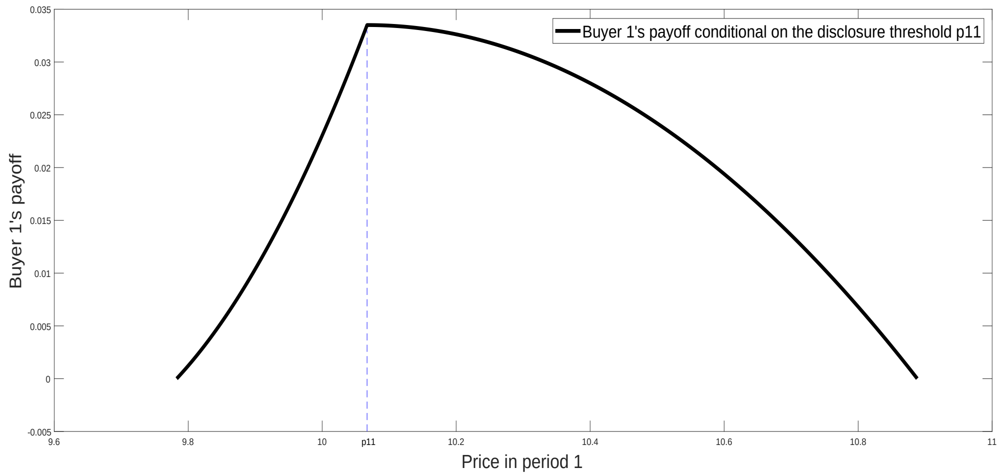

Figure 2.3. Buyer 1's payoff when the disclosure threshold is p_{11}

When $p_1 < p_1^*$, buyer 1 has a strong incentive to raise p_1. According to the seller's strategy, the price in period 1 will not be disclosed, if it is made below the disclosure threshold. Therefore, a small increase in p_1 will not be reported to buyer 2, and buyer 2's price will not be affected. Conditional on trade, buyer 1's marginal gain in average quality is $\frac{1}{2}$ when trading with extra high types, while his cost is $(1-\delta)\,\alpha$. When $\delta > \delta^*$, conditional on trade, the marginal gain from trading more high types outweighs the marginal cost of doing so. Moreover, a higher p_1 improves the chance of trade. Therefore, buyer 1 will set a price no lower than the disclosure threshold. This follows from the first part of buyer 1's payoff function where $p_1 \leq p_1^*$, which is also captured by the increasing part of buyer 1's payoff function in Figure 2.3. However, buyer 1 does not want to submit a price strictly above the disclosure threshold, as a future buyer will be informed and adjust his price upward accordingly. A higher p_1 now appears less attractive to the seller. In period 2, buyer 2 will propose $p_2 = \frac{A^2}{2A-1} q_1$, given his belief that the cutoff type is q_1. Conditional on trade, the marginal cost of trading extra high types becomes to $(1-\delta)\,\alpha + \delta \frac{A^2}{2A-1}$, while the marginal gain in average quality is still $\frac{1}{2}$. Buyer 1 has no incentive to offer a price higher than the disclosure threshold, as $(1-\delta)\alpha+\delta\frac{A^2}{2A-1}>\frac{1}{2}$. This follows from the second part of buyer 1's payoff function where $p_1>p_1^*$. In Figure 2.3, it is captured by the decreasing part of buyer 1's payoff function, which is to the right of the disclosure threshold.

I find that there are multiple equilibria, as the seller can pick different prices as the disclosure thresholds. In fact, there is a range of prices that can be sustained as the equilibrium disclosure thresholds, and they constitute the set of equilibrium prices in period 1.

When $\delta>\delta^*$, I have two constraints to construct the class of equilibria described in Proposition 7. One constraint is that buyer 1 has a non-negative payoff. This pins down the upper bound of the equilibrium cutoff type in period 1, which exists in the situation where buyer 1's surplus is zero. The other constraint is that buyer 1's payoff decreases when $p_1>p_1^*$. This pins down the lower bound of the cutoff type in period 1, which is equal to q_L. In Figure 2.4 below, I present buyer 1's payoff function under different price disclosure thresholds for the same set of parameters as in Figure 2.3. Note that $\delta>\delta^*$. As is shown in Figure 2.4, in any equilibrium, buyer 1's payoff reaches its maximum when p_1 matches the value of the seller's disclosure threshold. Any price in the interval $[p_{11},p_{15}]$ can be supported as the equilibrium disclosure threshold. Moreover, as δ decreases, the range of the equilibrium disclosure thresholds shrinks. This is because the seller becomes less willing to trade in period 1 when she is more patient. Buyer 1 needs to pay more in order to trade with the seller, and the seller will pick a higher disclosure threshold.

When $\delta\leq\delta^*$ and p_1 is below the disclosure threshold, the marginal benefit of raising p_1 is not always greater than the cost. In this situation, there is a price level and a corresponding cutoff type that maximize buyer 1's payoff when p_1 is not disclosed; this cutoff type is the upper bound of the equilibrium cutoff type in period 1. The lower bound of the cutoff type in period 1 is still q_L.

I can also characterize the situation off-path when $p_1>p_1^*$. Here, p_1 will be disclosed. Buyer 2 will form a certain belief of the cutoff type in period 1; in this scenario, suppose this

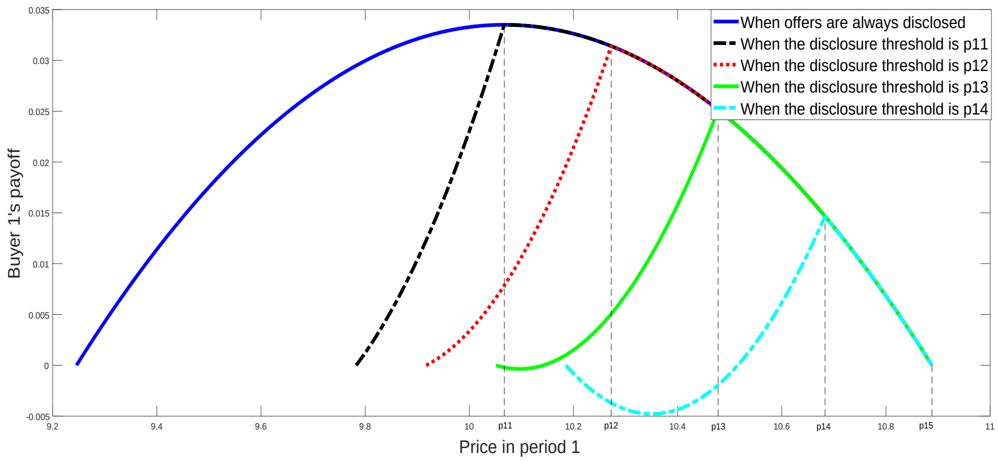

Figure 2.4. Buyer 1's payoff under different disclosure thresholds

belief is q_1. Given q_1, buyer 2 best responds by targeting the cutoff type q_2 where $q_2 = \frac{A}{2A-1}q_1$; p_2 is equal to Aq_2. The belief of the cutoff type q_1 must satisfy the condition that q_1 is indifferent between accepting p_1 and waiting for p_2. This shows that

$$p_1 - \alpha q_1 = \delta \left(\frac{A^2}{2A-1} q_1 - \alpha q_1 \right) \tag{2.1}$$

Equation (2.1) uniquely determines that $q_1 = \frac{p_1}{\delta \frac{A^2}{2A-1} + \alpha(1-\delta)}$. Given buyer 2's strategy, the actual cutoff type accepting p_1 is indeed q_1. Note that in this construction, buyer 2's strategy maximizes his payoff conditional on the belief q_1. As $p_1 > p_1^*$, buyer 2 is aware of this deviation.

In Figure 2.5, I illustrate the welfare split between the seller and the buyers for the equilibria when $\delta > \delta^*$. I can also rank the equilibria according to the overall welfare level, and this result is summarized in Proposition 8.

Proposition 8 *All the equilibria can be ranked by their welfare levels.*

(1) When $\delta \leq \delta^*$, the equilibrium with the highest welfare level is the one where $\widehat{q_1} = \frac{(2A-1)(1-\delta)\alpha}{\delta A^2 + (2A-1)(2\alpha(1-\delta)-1)} \underline{q}$. The equilibrium with the lowest welfare level is the one where $\widehat{q_1} = q_L$.

(2) When $\delta > \delta^*$, the equilibrium with the highest welfare level is the one where $\widehat{q_1} = \frac{(2A-1)}{2\delta A^2 + (2A-1)(2\alpha(1-\delta)-1)} \underline{q}$. In this equilibrium, buyer 1 has a payoff of zero. The equilibrium with

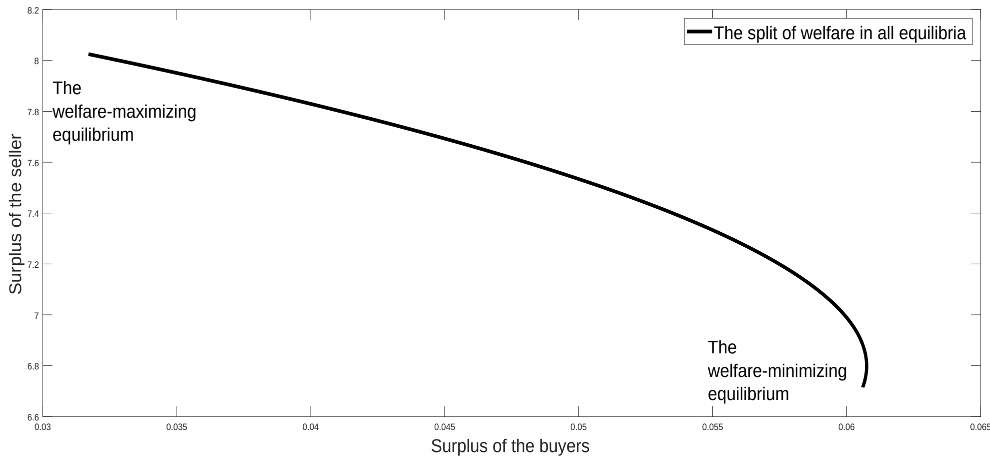

Figure 2.5. Welfare split between the seller and buyers

the lowest welfare level is the one where $\widehat{q_1} = q_L$.

The total welfare level in the game is monotonic in $\widehat{q_1}$; therefore, the equilibrium with the highest $\widehat{q_1}$ corresponds to the one that has the highest trading efficiency. When $\delta \leq \delta^*$, in any equilibrium, buyer 1 and buyer 2 both end up with non-zero surpluses. When $\delta > \delta^*$, the equilibrium cutoff in period 1 reaches the maximum when buyer 1 has an expected payoff of zero, which corresponds to the upper-left point in Figure 2.5. This is also when the overall welfare is maximized. Here, the seller's expected welfare reaches its maximum level, so does buyer 2's.

2.4.2 Welfare Analysis and Implications

In this section, I compare the equilibrium welfare levels in the optional disclosure model with the ones in two benchmark models where the observability of previous offers is exogenously determined. [23] and [18] present the results under two extreme information structures. One is a transparent market, where past offers are always observable to future buyers. The other is an opaque market, where previous offers are always unobservable. In these two benchmark models, the information structure is exogenous, and buyer 2 updates his belief in a relatively mechanical way. Before I present the welfare comparison, I first review the researchers' results under these two information structures in the current context. The equilibrium characteristics are summarized in Proposition 9 and 10, found below. For notation purpose, I refer to the benchmark model where the previous offers are observable as the BMO model and the benchmark model where the previous offers are unobservable as the BMU model.

Proposition 9 *In the BMO model, there is a unique equilibrium. The cutoff types are* $q_2^* = \frac{A}{2A-1} q_1^*$ *and* $q_1^* = q_L$*, respectively.*

When the past offers are observable, there is a unique equilibrium, and it is in pure strategies. This is because the current buyer's price offer can perfectly affect the future buyer's pricing strategy. Buyer 2's belief of the seller's type depends on p_1. Specifically, the belief about the remaining type at period 2 is above a single cutoff q_1. The construction of this unique equilibrium follows from backward induction.

However, if the past offers are never observable, I find that the equilibrium result is completely different. Here buyer 2's belief of the seller's type is independent of buyer 1's actual offer. I use $K_1(\cdot)$ to represent the cumulative distribution function of the cutoff type induced by the equilibrium price in period 1 and $K_2(\cdot)$ to represent the cumulative distribution function of the cutoff type induced by the equilibrium price in period 2.

Proposition 10 *In the BMU model,*

(1) When $\delta \leq \delta^*$*, there is a unique pure strategy equilibrium; the cutoff types in the two periods are* $q_2^{**} = \frac{A}{2A-1} q_1^{**}$ *and* $q_1^{**} = \frac{(2A-1)(1-\delta)\alpha}{\delta A^2 + (2A-1)(2(1-\delta)\alpha - 1)} \underline{q}$*.*

(2) When $\delta > \delta^*$*, there is no pure strategy equilibrium. Buyer 1 mixes between the prices that induce a cutoff* $\underline{q}$ *and cutoffs in interval* $\left[\underline{q}_2, \overline{q}_2\right]$*, where* $\underline{q}_2 = \frac{A}{2A-1} \underline{q}$ *and* $\overline{q}_2 = \frac{1}{2\delta A + 2\alpha(1-\delta) - 1} \underline{q}$*. Buyer 2 mixes between the prices that induce cutoffs* $\underline{q}_2$ *and* $\overline{q}_2$ *with probability* $K_2\left(\underline{q}_2\right) = \frac{1-2\alpha(1-\delta)}{2\delta A}$.

The type of equilibrium of the BMU model depends on the discount factor. Conditional on trade, the marginal gain of increasing p_1 is $\frac{1}{2}$ while the marginal cost is $\alpha(1-\delta)$. A pure strategy equilibrium exists when δ is small. However, when δ is large, there cannot be any pure strategy equilibrium, as buyer 1 could increase his surplus by deviating from the original p_1 and offering a higher price. This deviation would be unknown to buyer 2. However, in my model with the disclosure option, the seller adopts different disclosure actions above and below the disclosure threshold. Any deviation above the disclosure threshold will be known to buyer 2.

The market operates quite differently under these two regimes. In period 2, the continuation of the equilibrium of the BMU model does not depend on the offer history, while in the BMO model it does. Therefore, given $\delta > \delta^*$, when offers are observable to the future buyer, the seller is more willing to reject a high offer today, which gives her the opportunity to receive a better offer in the next period. That's why the two benchmark models have very different equilibria.

When $\delta \leq \delta^*$, the welfare-maximizing equilibrium of the optional disclosure model has the same cutoff as the unique equilibrium of the BMU model, while the welfare-minimizing equilibrium has the same cutoff as the unique equilibrium of the BMO model. Proposition 11 compares the welfare-maximizing equilibrium with the equilibria in the two benchmark models when $\delta > \delta^*$.

Proposition 11 *Given $\delta > \delta^*$,*

(1) There exists A^* such that, for $A > A^*$, the (expected) trading price in period 2 is strictly higher in the welfare-maximizing equilibrium of the optional disclosure model than in any equilibrium of the BMU model.

(2) There is A^{**} and $\overline{q}^*$ such that, for $A > A^{**}$ and $\overline{q} \leq \overline{q}^*$, the (expected) trading price in period 1 is also strictly higher in the welfare-maximizing equilibrium than in any equilibrium of the BMU model. In addition, if $A > A^*$ holds, the welfare-maximizing equilibrium ex-ante Pareto dominates any equilibrium of the BMU model.

(3) Any equilibrium of the optional disclosure model yields a (weakly) higher social surplus than the equilibrium of the BMO model.

In the welfare-maximizing equilibrium of the optional disclosure model, as well as any equilibrium of the BMU model, buyer 1 receives a payoff of zero. As for the seller, since the price is higher for both periods in the welfare-maximizing equilibrium of the optional disclosure model under some conditions, every seller type is ex-ante better off in the model with optional disclosure. Specifically, these sets of conditions require that A is large enough[8]; moreover, the variation in the quality of the good cannot be too large. As for buyer 2, he ends up with a strictly better payoff in the optional disclosure model. Consequently, these conditions altogether point toward part (2) in Proposition 11. Part (3) is fairly straightforward to see.

In the two-period model, I start to see the difference that emerges when the disclosure option is allowed. Here, I discover that the seller has more flexibility when she sends signals the prospective buyer. From buyer 2's perspective, the optional disclosure model can generate new beliefs, compared with both the BMU model and the BMO model.[9] Among those beliefs, some are more "optimistic" than others. If the belief is optimistic, buyer 2 believes the seller sets a pretty high price as her disclosure threshold. Buyer 1 would make an offer exactly equal to the disclosure threshold, and a low seller type would accept this offer. It indicates that the seller rejecting buyer 1's offer must have an item of considerable quality. This generates those optimistic beliefs and makes buyer 2 willing to offer a good price. Meanwhile, the welfare-maximizing equilibrium corresponds to the condition where buyer 2 is the most optimistic. In this equilibrium, the belief of the cutoff type in period 1 is so optimistic that buyer 2 will offer a sufficiently good price which could be even higher than the second period expected trading price in the BMU model. Conditional on a sufficiently high price in period 2, the trading price in period 1 could be higher in the optional disclosure model as well. As a result, more types have the chance to trade in the optional disclosure model. Allowing flexibility in the disclosure option can generate more optimistic beliefs, which potentially leads to Pareto improvement in

[8] In the proof of Proposition 11, I also discuss other possible conditions. Actually, when $A \geq \delta + \alpha(1-\delta)$, part (1) holds naturally for any $\alpha > \frac{1}{2}$ and $\delta > \delta^*$.

[9] Note that in the BMU model, when $\delta > \delta^*$, buyer 1 mixes between different prices in equilibrium. Buyer 2 has a non-degenerate belief of the cutoff type seller. For buyer 2, the cutoff type in period 1 may be a low type or a relatively high type. Taking all these into account, buyer 2 offers his price.

the trading outcome.

Here, I also obtain new insights from the seller's perspective. In this model, the buyer is the party who makes the offer, and he has all the bargaining power. However, allowing the disclosure option also gives the seller the ability to influence the outcome of the game. The disclosure option essentially gives the seller the power to control the amount of information released to the future buyer, which influences the future buyer's belief and his interpretation of what has happened in the past. By choosing her favorite disclosure threshold, the seller affects buyer 2's response to buyer 1's price and buyer 1's incentive of making an offer. This disclosure option benefits the seller in terms of the trading prices, which are strictly higher in some equilibria in my model than in any equilibrium in the BMO model and the BMU model.

Even though non-disclosure happens on the equilibrium path, the possibility of disclosure enables the seller to convey positive messages to the future buyer, which plays a crucial role in the equilibrium construction and the welfare results. Although in some real marketplace, we often do not observe disclosure behaviors of past offers. It may not be a good idea to disallow the disclosure option completely. Creating such a non-transparent market impedes not only the transmission of negative signals between the seller and the buyer but also positive ones. This could exert a negative impact on the total welfare level as it is shown in Proposition 11.

Compared with the results in the disclosure literature, the optional disclosure model has a set of equilibria where the equilibrium offer is not disclosed on the equilibrium path. The reason is that the disclosure content here is type-independent. All seller types have the same disclosure content, i.e., the price history, which is not a perfect signal of the seller's type. Therefore, disclosure is not a dominant strategy for the highest seller type; thus, I can construct an equilibrium where the equilibrium offer is not disclosed on the equilibrium path. However, the price history does influence a future buyer's belief of the seller's type. The seller has the incentive to disclose a high offer that will lead a future buyer to improve his belief of the seller's type.

2.5 Discussions

In this paper, I assume that the seller type follows a uniform distribution for the purpose of welfare analysis. The equilibrium construction is very similar when I have a more generalized distribution of types. In this case, the seller still maintains the cutoff-type acceptance rule in equilibrium. When δ is large enough, there is still one equilibrium with full disclosure. In addition, there can be multiple equilibria where the equilibrium offers are not disclosed. In those equilibria, disclosure happens only for the higher-than-equilibrium offers.

Another aspect that is worth of attention is the disclosure rule. In this paper, the seller has the flexibility to decide the equilibrium disclosure threshold, which brings the multiplicity of equilibria here. However, among these equilibria, some are more desirable than others. From a mechanism design perspective, policy makers can apply restrictions to the disclosure rule to select the equilibrium in which they are mostly interested. Specifically, instead of letting the seller choose the disclosure price, policy makers may select the disclosure threshold. For example, in the two-period model, when the seller is patient, policy makers can impose non-disclosure for a price under a threshold $\widehat{p_1}$ and voluntary disclosure (or mandatory disclosure) for all other prices to select out the welfare-maximizing equilibrium. Here, $\widehat{p_1}$ is equal to the equilibrium price p_1^* in the welfare-maximizing equilibrium. In this case, the cutoff type in period 1 is uniquely determined as $\frac{(2A-1)}{2\delta A^2+(2A-1)(2\alpha(1-\delta)-1)}\underline{q}$, and the welfare-maximizing equilibrium can be reached. Adopting a non-disclosure policy for lower prices and a voluntary (or mandatory) disclosure policy for higher prices could encourage the transmission of positive messages in the market while suppress negative ones, which potentially enhances trading efficiency.

A similar argument can be extended to the situation where there are possibly multiple trading units. For example, eBay in 2013 changed its disclosure policy regarding previous purchases. Instead of displaying the exact accepted prices for recent purchases, only those purchases that met the seller's asking price are displayed with the accepted amount. For those transactions that are processed in "Best offer" or "Special offer" (below the asking price), the selling prices are hidden, and only the existence of those offers is disclosed. Ebay also withholds the amount of a previously rejected "Best offer" bid and only reveals that the bid was rejected. This change in eBay's disclosure policy ensures the transmission of positive information about the trading item and promotes trading efficiency.

2.6 Extensions

This section discusses several extensions of the model. My current results are robust under some modeling variations. In Section 2.6.1, I extend the model to an infinite-period setting and discuss the equilibrium properties. In Section 2.6.2, I consider the situation of allowing buyers to inspect historical offers at some cost. In Section 2.6.3, I consider a variation when the offers may be leaked in adjacent periods with a small probability by nature. In Section 2.6.4, I discuss the situation where there are multiple buyers in every period.

2.6.1 The Infinite-Period Model

In this subsection, I present the equilibrium result under an infinite-period setting when $\delta > \delta^*$ and $\underline{q} \leq (2\alpha - 1)\overline{q}$. Here, the seller is patient, and there is a large variation in the asset's quality.[10] There is no final period, and the seller always has a chance to meet another buyer. In this setting, I define $p^{t-1} = \{p_1, \ldots, p_{t-1}\}$ as the history of offers at period t. P^{t-1} denotes the set of all possible price histories at period t, and $P^0 = \emptyset$. Buyer t has a belief about the seller's type, and the cumulative distribution function of the belief is $G_t(q)$. In the equilibrium construction, applying a similar argument as in Lemma 2 in the two-period model, a mixed strategy cannot be supported in equilibrium for any buyer unless one of the buyer's offers is accepted by all seller types. This property is formally restated as Lemma 8 in Appendix B.

Specifically, I focus my analysis on a tractable equilibrium called the *Simple Threshold* (ST) equilibrium in which the seller adopts a threshold rule and only discloses offers strictly

[10] When there is a small variation in the quality of the good, every buyer's optimal strategy is to serve all seller types. The analysis for this circumstance can be found in Appendix B.

above the disclosure threshold. Equilibrium offers are not disclosed on the equilibrium path. This is similar to the equilibria constructed in the two-period model. I formally show that for any other equilibrium, there exists an ST equilibrium that has the same cutoff type and trading price in every period. Hence, it is reasonable to focus only on the ST equilibrium. I formally define the *Simple Threshold* (ST) equilibrium as follows.

Definition 1 *The seller's disclosure rule d^t is called a threshold rule if there exists threshold functions: $\widehat{p_{m,t}}(\cdot):\left[\underline{q},\overline{q}\right]\times P^{t-1}\longrightarrow R_+$ for any $m\leq t-1$, such that $p_m\in d^t\left(q,p^{t-1}\right)$ if $p_m>\widehat{p_{m,t}}\left(q,p^{t-1}\right)$. Otherwise $p_m\notin d^t\left(q,p^{t-1}\right)$. If a strategy profile $(d^t,p_t^*,\widehat{q_t})$ in which d^t is a threshold rule constitutes an equilibrium, and every equilibrium offer is not disclosed on the equilibrium path, it is said to be an ST equilibrium.*

In an ST equilibrium, the seller has a threshold function $\widehat{p_{m,t}}(\cdot)$ in period t, which is a mapping from the seller's type and the price history to R_+. She discloses offer p_m made in period m to buyer t if and only if $p_m>\widehat{p_{m,t}}\left(q,p^{t-1}\right)$. No equilibrium offer is disclosed. Due to a similar argument explained in Lemma 4, any seller type has the incentive to disclose a price higher than the equilibrium price when rejecting it. Therefore, the value of the disclosure threshold $\widehat{p_{m,t}}(q,\cdot)$ on the equilibrium path matches the equilibrium price p_t^*. From Lemma 5, I also know that the threshold $\widehat{p_{m,t}}$ is invariant in every period $t>m$.

Lemma 5 *When $\delta>\delta^*$, in any ST equilibrium, if $p_m\in d^{m+1}(q,p^m)$, then $p_m\in d^t\left(q,p^{t-1}\right)$ for any $t>m+1$. Additionally if $p_m\notin d^{m+1}(q,p^m)$, then $p_m\notin d^t\left(q,p^{t-1}\right)$ for any $t>m+1$.*

Lemma 5 characterizes the invariance of the disclosure threshold across time. On the equilibrium path, if the seller discloses offer p_m to buyer $m+1$, then the seller must also disclose p_m to any buyer t ($t>m+1$). If the disclosure threshold $\widehat{p_{m,t}}\left(q,p^{t-1}\right)$ is higher than $\widehat{p_{m,m+1}}(q,p^t)$ for some $t>m+1$, it creates room for buyer m to make a small deviation that will not be reported to some upcoming buyers. Some future prices will remain unchanged. Therefore, the benefit of raising p_m could outweigh the cost. Lemma 5 allows me to simplify the notation of

$\widehat{p_{m,t}}\left(q, p^{t-1}\right)$, as it is invariant for all $t > m$ in any ST equilibrium. Thus, I denote the disclosure threshold for p_m by $\widehat{p_m}\left(q, p^{t-1}\right)$. In Theorem 2, I claim that for any equilibrium of the game, there exists an ST equilibrium that is payoff-equivalent to the participants. This allows me to focus on characterizing the ST equilibrium.

Theorem 2 *When $\delta > \delta^*$ and $\underline{q} < (2\alpha - 1)\overline{q}$, for any equilibrium of the infinite-horizon game, there is an ST equilibrium that induces the same cutoff type and price in every period.*

The proof of Theorem 2 is constructive and is deferred to Appendix B, but I will sketch the main idea here. In the payoff-equivalent ST equilibrium, the seller's disclosure strategy is constructed as $p_m \in d^t\left(q, p^{t-1}\right)$ only if $p_m > \widehat{p_m}\left(q, p^{t-1}\right)$, for $m \leq t-1$. The seller selects the disclosure thresholds such that $\widehat{p_1} = p_1^*$ and $\widehat{p_t} = p_t^*$ if $p_m \leq \widehat{p_m}$ for all $m \leq t-1$. Moreover, whenever there is a period $m \leq t-1$ such that $p_m > \widehat{p_m}$, the seller reports p_m to buyer t. The disclosure threshold for offers in period t is now raised to $\widehat{p_t} = p_t^* + \varepsilon$ where $\varepsilon = p_m - \widehat{p_m}$. In this construction, it is straightforward to check that the equilibrium offer is not disclosed; in addition, the equilibrium offer matches the value of the disclosure threshold, in every period. There is no deviation from either the seller or buyers, and it is an ST equilibrium. Here, the equilibrium cutoff types are the same as those in the original equilibrium.

Proposition 12 describes the ST equilibria.

Proposition 12 *When $\delta > \delta^*$ and $\underline{q} < (2\alpha - 1)\overline{q}$, there are two types of ST equilibria.*

(1) In a type-1 ST equilibrium, there exists $n^\star$ such that the equilibrium cutoff types remain constant after period $n^\star$, i.e., $\widehat{q_1} \leq \cdots \leq \widehat{q_{n^\star-1}} \leq \widehat{q_{n^\star}} = \widehat{q_{n^\star+1}} = \cdots$. Buyer $n^\star$ offers $\alpha\widehat{q_{n^\star}}$ and any future buyer t will make an offer no more than $\alpha\widehat{q_{n^\star}}$.

(2) In a type-2 ST equilibrium, there are two phases. There is a smallest period $n^{\star\star}$ such that $\overline{q}$ is served with positive probability. In phase 1, $\widehat{q_1} \leq \widehat{q_2} \leq \cdots \leq \widehat{q_{n^{\star\star}-1}}$ and $\widehat{q_{n^{\star\star}-1}} = (2\alpha - 1)\overline{q}$. In phase 2, buyer t ($t \geq n^{\star\star}$) randomizes between $\alpha\overline{q}$ and p_t^L where $p_t^L \leq p^*_{n^{\star\star}-1}$. The

probability of offering $\alpha\overline{q}$ is λ such that the following condition holds.

$$p^*_{n^{\star\star}-1} - \alpha\widehat{q_{n^{\star\star}-1}} = \delta\left(\lambda\,\alpha\overline{q} + (1-\lambda)\,p^*_{n^{\star\star}-1} - \alpha\widehat{q_{n^{\star\star}-1}}\right)$$

In a type-1 ST equilibrium, the seller sets her disclosure threshold as $\alpha\widehat{q_{n^\star}}$ for some $\widehat{q_{n^\star}} < \overline{q}$ in all following periods after period $n^\star$. Future buyers will observe any offer higher than $\alpha\widehat{q_{n^\star}}$. Under this circumstance, from some intermediate period onward, every offer that potentially results in trade will be revealed to all buyers in the future. This situation is quite similar to the BMO model. According to the intuition found in [23], all future buyers will not offer prices higher than $\alpha\widehat{q_{n^\star}}$ because offering higher prices will trigger more aggressive prices in the future. The seller can gain more reputation to reject a current high offer and disclose it to the next buyer. She will wait for good offers in the future, which leads to an impasse in trading after some period. If the seller sets her disclosure threshold as $\alpha\underline{q}$ in every period, any serious offer will eventually be revealed to all buyers, and I will have the same equilibrium result as in [23].

In a type-2 ST equilibrium, every seller type trades eventually. The seller sets a non-decreasing sequence of prices as the disclosure thresholds, and all types trade across time. In period $n^{\star\star}-1$, the cutoff type is $(2\alpha-1)\overline{q}$. After period $n^{\star\star}-1$, the seller sets the disclosure threshold at $p^*_{n^{\star\star}-1}$, which is the trading price in period $n^{\star\star}-1$. All future buyers randomize between $\alpha\overline{q}$ and p^L_t, where $p^L_t \leq p^*_{n^{\star\star}-1}$. $\alpha\overline{q}$ is accepted by all seller types, while p^L_t is rejected by all remaining types. These two offers both give the buyer a payoff of zero. The disclosure threshold has the property that $\widehat{p_t} = p^*_t$ when $t \leq n^{\star\star}-1$, and $\widehat{p_t} = p_{n^{\star\star}-1}$ when $t > n^{\star\star}-1$. Here, every buyer has no incentive to deviate, as a higher offer would be reported to buyers in the future, which could result in more aggressive offers. Additionally, a lower offer would be concealed, which would lead to a loss of trades and therefore impair his surplus.

2.6.2 Costly Inquiry

Another variation is to allow buyers to inspect the price history at some cost. In this case, it is evident that a buyer will never pay any price to inspect the historical offers in equilibrium. This is because offers are made in pure strategies unless the highest type is served. The possible randomization occurs only when the highest seller type is served. In the infinite-horizon game, I show that in equilibrium a buyer only randomizes between exactly two offers if he plays a mixed strategy, one of which is $\alpha\bar{q}$. Therefore, the offer history is entirely predictable from the buyer's perspective, and he refuses to pay anything to inquire about past offers. Every equilibrium of the original model also exists under this modified setting.

2.6.3 Leakage of Past Offers

Another possible extension is that the offers may be leaked in adjacent periods with a small probability ε by nature. This is actually a practical modification of the original model, as the past transaction prices may be leaked by some anonymous information sources in real-world trading. Here, the original model in the paper becomes a special case with $\varepsilon = 0$. The equilibrium result when $\delta > 1 - \frac{1}{2\alpha}$ is robust with a sufficiently small perturbation ε. As long as ε is sufficiently small, the current buyer would have relatively little chance to learn the actual offer made in the previous period. In the two-period model, if buyer 1 lowered his price today, there is a very large chance that this reduced amount would not be caught by buyer 2, and therefore the next period offer would remain unchanged. There is a range of high seller types who accept the original offer now would reject this reduced offer and switch to buyer 2. In this case, the probability of trade is reduced. For a small enough ε, the range of these switching seller types is so great that the drop in average quality of the good outweighs the reduction in price. In this case, reducing the price below the disclosure threshold would not benefit buyer 1. In summary, buyer 1 still has the incentive to raise his price to match the value of the seller's disclosure threshold. Any equilibrium of the original model also exists under this variation.

2.6.4 Multiple Buyers

The model can also be extended to have competition on the buyer's side. Instead of having a single buyer in each period, I can allow multiple buyers into the market and have them make simultaneous offers to the seller. In this scenario, when $\delta > \delta^*$, the equilibrium construction is very similar. The only difference is that every buyer ends up with a payoff of zero, due to a Bertrand competition setup. As a result, the equilibrium is unique in the two-period model.

2.7 Conclusion

In my modeling exercise with optional disclosure, I discuss the equilibrium construction in both the two-period model and the infinite-period model. Compared with the BMO model and the BMU model, the optional disclosure model has a new set of equilibria generated by allowing for flexibility in disclosure. This is due to the fact that buyers have a new class of equilibrium beliefs. These new beliefs originate from the seller's threshold rule in disclosing previous offers. Such beliefs do not exist in either the BMO model or the BMU model. As for the seller, she also utilizes this flexibility in disclosure, taking different actions for different prices to promote trade.

Furthermore, I do welfare analysis and compare the generated social welfare in these three models when there are two trading periods. Here I find that it is in policy makers' interest to utilize the disclosure option, especially when the item's quality does not vary broadly. With optional disclosure, information is filtered through a threshold rule and transmitted more efficiently, which enhances trading efficiency in the market under some conditions. Policy makers may not necessarily impose mandatory disclosure or concealment on every past offer. More importantly, although there is a multiplicity of equilibria, policy makers can apply restrictions to the disclosure rules in order to select the equilibrium in which they are interested. Adopting a non-disclosure policy for lower prices and a voluntary (or mandatory) disclosure policy for higher prices could potentially enhance trading efficiency. This idea is also reflected in eBay's display of previous offers.

Information about past transactions is a very crucial part of the market information structure. This paper focuses on the seller's incentive to voluntarily disclose past rejected prices and provides policy implications for designing or regulating the market information structure. More generally, there are other aspects of transactional information. How to design the optimal disclosure policy for general transactional information is still an interesting and open question to study.

Chapter 3

Long-Term Contracting with a Present-biased Agent

People's time preferences affect their tradeoffs between payoffs in a long-run contracting environment. In this paper, I consider a long-term contracting problem between a monopolistic seller and a present-biased buyer with asymmetric information in a Markov environment. The analysis is focused on the situation in which the buyer and the seller are fully aware of the degree of inconsistent discounting. I find that time inconsistency affects the optimal allocations through a novel unconditional effect, which is essentially the cross-period marginal effect of the current type on future types in the Markov environment. The optimal contract still possesses the *no distortion at the top* feature, and the allocations for the low type are always distorted for any realization of types in the history. However, the principal's expected surplus from contracting with the present-biased agent could be strictly worse than contracting with the time consistent agent.

3.1 Introduction

People's time preferences influence their inter-temporal decision making. Standard theories in the long-term contracting environment are based on the discounted-utility (DU) model. This model assumes that consumers evaluate payoffs at a consistent discount rate, and it is represented by an exponential discount function. However, growing evidence from field studies and experiments has documented that people's time preferences do not follow the standard exponential discounting assumption. Moreover, they exhibit a fairly inconsistent pattern; people tend to be more patient in the long run than in the short run. Such preferences have been called present-biased preferences or time-inconsistent preferences, which are often captured by quasi-hyperbolic discount functions in the literature ([48]; [42]; [30]; [38]).

With present-biased preferences, the discount rate between two consecutive periods falls across time. People are inclined to underestimate the trade-off between payoffs at future periods when making decisions beforehand, and this leads them to give in to current temptations and procrastinate on difficult tasks. A consequence of present-biased preferences is that people are often behind schedule. What is considered as optimal today may not turn out to be optimal when the future comes. This issue becomes a real problem in a unilateral commitment situation where the consumer can leave the contract anytime while the seller fully commits.[1] The consumer tends to withdraw from the contract at an earlier stage than what he has planned as he underestimates the tradeoffs at later periods. The following example illustrates this situation.

Example 1 *A cellular phone company provides service to a present-biased consumer continuously for three periods. The consumer can walk away from this relationship at any* $t \geq 1$. *The service itself gives the consumer 6 utils each period, and the company charges* $p = (p_1, p_2, p_3) = (6, 7, 5)$. *The consumer has the present-bias parameter* $\beta = 0.9$ *and his per-period payoff satisfies* $u_t = 6 - p_t$ *if the contract is accepted. I assume* $\delta = 1$ *for simplicity. The consumer plans to finish the three-period contract when contracting in period* 1. *However, in fact, he exits after period* 1.

A natural question to ask is how should a profit-maximizing firm modify the contract if it is aware that the consumers are present-biased. How does time inconsistency affect the firm's optimal selling strategies? To answer these questions, I study a dynamic monopolistic screening

[1] The unilateral commitment setting seems to be the most appropriate for many marketplaces including the insurance market and the internet plan market.

problem with a present-biased agent. The model incorporates the framework of [5], and it is a multi-period model with the participation of a buyer (agent, he) and a seller (principal, she). The buyer repeatedly purchases some nondurable good from the seller. In every period, the buyer has private information about his type, which follows a Markovian evolution process. I extend the existing literature in dynamic mechanism design by allowing the agent's time preference to be present-biased.

The agent's time inconsistency is modeled by a quasi-hyperbolic discount function. The discount rate is $\beta\delta$ between the current period and the next period, and it is δ for any other subsequent pair of consecutive periods. In the model, I assume that the agent is sophisticated, and the degree parameter β is known both to the buyer and the seller. This directly relates to the previous literature studying contracting problems between a principal and an agent where the degree of inconsistency is known to the principal. [39] consider a moral hazard problem with a time-inconsistent agent. In their model, the agent has private information about his cost of accomplishing a task, and waiting is optimal when the private cost of the agent is too high. However, as the principal cannot distinguish between inefficient procrastination and an efficient delay, he cannot implement the first-best contract. In [13], a monopolistic firm designs a contract with a two-part tariff for the present-biased agents. When the goods have immediate costs and delayed benefits (e.g., health clubs), the monopoly charges a price below the marginal cost and a high entry fee. When the agent is fully sophisticated, the firm can still achieve the first-best profit level. However, the result in [13] heavily relies on the assumption that there is no asymmetric information between the principal and the agent. In [16] and [15], they use the same utility function as [20]. In a static setting, the firm designs the choice menu that the agent can choose from, which screens the agent's willingness to pay. The firm can extract all the surplus when the agent wants to buy a larger quantity or a higher quality good. Another approach is provided by [14]: they consider the situation where the principal provides contracts to agents with different abilities to forecast changes in their future tastes. In this context, the principal knows the agent's degree of inconsistency but does not know if the agent is aware of his inconsistency. The authors find that the optimal menu of contracts serves as a commitment device for the relatively sophisticated agents but exploits the naive agents. My paper distinguishes itself from the previous literature through its dynamic features. In my model, I assume that the degree of inconsistency is known to both the seller and the buyer. The seller designs the long-term contract in order to screen the buyer's valuation of the good. Since the buyer's current valuation has an impact on his future valuation, the seller needs to take account of this impact when designing the optimal long-run contract for the present-biased buyer.

Specifically, I characterize the profit-maximizing contract for a monopolistic seller who interacts with a present-biased agent. This contract is very different from the contract for a time consistent agent. In [5], where the buyer is time consistent, the contract becomes efficient once the buyer reports himself to be a high type. Only a persistent low type continues to receive an inefficient allocation. The distortion in the contract is due to the marginal effect of the current type on future types, which is the per period marginal effect in the Markov environment captured by the "impulse response" function in [40]. However, the optimal contract for a present-biased agent lacks this feature, and the degree of the distortion depends on the degree of time inconsistency. Regarding time preferences, the monopolist is more patient than the buyer. This brings an additional residual effect which captures the distortions in allocations due to the non-transferable utilities between the buyer and the seller. The seller can take advantage of the difference in discounting by imposing payments at later periods. However, the seller's hands are tied, as he also needs to provide the buyer an incentive to stay in the contract until the end. Accordingly, the characterization of the optimal contract offers insights that help us better understand present-biased preferences. Time inconsistency affects the optimal allocations through a novel unconditional effect, which is essentially the cross-period marginal effect of the current type on future types in the Markov environment. The optimal scheme combines the "impulse response" (also referred to as the conditional effect in my paper), the novel unconditional effect and the residual effect.

This paper is also related to the research papers on dynamic mechanism design. In the standard mechanism design problem, the Revelation Principle allows us to simplify the problem and only focus on designing the direct mechanism. This paper also concentrates on the direct mechanism. In the model, I incorporate the conventional setup from [36] and [5] with one principal and one agent. Regarding the methodology, I follow the standard dynamic mechanism design approach ([4]; [10]; [6]; [40]) to solve the multi-period model.

The rest of the paper is organized as follows. Section 3.2 describes the model. In Section 3.3, I characterize the optimal contract for a three-period model where there are binary types. I also discuss the properties of the contract and compare it to the benchmark model where the buyer is time consistent. Section 3.4 studies a three-period model with continuous types. Section 3.5 presents the optimal contract for a T-period model where there are binary types. Section 3.6 concludes.

3.2 The Model

I consider a three-period model. In each period, the buyer (agent) purchases a non-durable good from the seller (principal), and the contract on quantities and transfers (x_s, t_s) is signed in the first period. The buyer has the option to exit the contract at the beginning of the second or the third period if his expected continuation payoff of staying in the contract falls below his value of the outside option, i.e., 0.[2] This suggests that I will impose the ex-post participation constraints in all three periods. The buyer's private valuations of the good in period 1, 2 and 3 are denoted by θ_1, θ_2 and θ_3, respectively. He receives a payoff of $\theta_s x_s - t_s$ per period when consuming x_s. θ_1 is known only to the buyer at the time of contracting, and θ_2 and θ_3 are realized at the end of period 1 and period 2, respectively. The common knowledge is the distribution of the buyer's type, i.e. $F_1(\theta_1)$, $F_2(\theta_2|\theta_1)$ and $F_3(\theta_3|\theta_2)$. Here, I assume the evolution of the buyer's type follows a

[2] There are two reasons to provide the exit option to the buyer. One reason is that if the exit option is not provided, the seller can make an infinitely large profit by making transfers to the buyer in the first two periods and charge a tremendous amount in the last period. The other reason is that this assumption is appropriate in many practical contracting environments. Usually, the seller fully commits to the terms of the contract while the buyer can leave and take other better deals. Sometimes, exiting the contract costs the buyer a fixed penalty c. Here I normalize c, and consider the simplest case where $c = 0$.

Markov process. In all three periods, the support of the type space is $\left[\underline{\theta}, \overline{\theta}\right]$, and I assume that all the PDF and CDF functions are continuous differentiable. The inverse hazard rates $\frac{1-F_1(\cdot)}{f_1(\cdot)}$, $\frac{1-F_2(\cdot|\theta_1)}{f_2(\cdot|\theta_1)}$ and $\frac{1-F_3(\cdot|\theta_2)}{f_3(\cdot|\theta_2)}$ are non-decreasing in the corresponding variables. In period s, the seller incurs a loss of $\frac{1}{2}x_s^2$ when she produces x_s units. Under this specification, notice that the efficient allocation in period s is exactly θ_s.

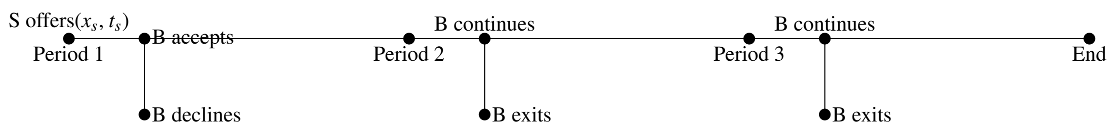

Figure 3.1. The timeline of the three-period game

In my model, the buyer is present-biased with a $\beta - \delta$ discounting function while the seller is time consistent with a discount factor δ. I further assume the degree of time inconsistency is known to the buyer and the seller. The inter-temporal utility functions are given below.

$$\begin{aligned} \text{Buyer's utility} &: \theta_1 x_1 - t_1 + \beta\delta\left(\theta_2 x_2 - t_2\right) + \beta\delta^2\left(\theta_3 x_3 - t_3\right) \\ \text{Seller's utility} &: t_1 - \frac{1}{2}x_1^2 + \delta\left(t_2 - \frac{1}{2}x_2^2\right) + \delta^2\left(t_3 - \frac{1}{2}x_3^2\right) \end{aligned}$$

Following [36], I focus on the direct mechanism of reporting types. Let $x_1(\theta_1)$, $x_2(\theta_2|\theta_1)$ and $x_3(\theta_3|\theta_1,\theta_2)$ denote the trading quantities in the three periods, respectively, and $t_1(\theta_1)$, $t_2(\theta_2|\theta_1)$ and $t_3(\theta_3|\theta_1,\theta_2)$ the corresponding payments. The seller's problem can be expressed as:

$$\max_{x_1,x_2,x_3,t_1,t_2,t_3} E_{\theta_1}\left[t_1 - \frac{1}{2}x_1^2\right] + \delta E_{\theta_1,\theta_2}\left[t_2 - \frac{1}{2}x_2^2\right] + \delta^2 E_{\theta_1,\theta_2,\theta_3}\left[t_3 - \frac{1}{2}x_3^2\right]$$

subject to the incentive compatibility (IC) constraints and the individual participation (IR) constraints.

$$\begin{aligned}
\text{IC1} &: \quad V_1(\theta_1) \geq V_1(\theta_1', \theta_1) \text{ for all } \theta_1 \text{ and } \theta_1' \\
\text{IC2} &: \quad V_2(\theta_2|\theta_1) \geq V_2(\theta_2', \theta_2|\theta_1) \text{ for all } \theta_1 \text{ and all } \theta_2,\ \theta_2' \\
\text{IC3} &: \quad V_3(\theta_3|\theta_1, \theta_2) \geq V_3(\theta_3', \theta_3|\theta_1, \theta_2) \text{ for all } \theta_1,\ \theta_2 \text{ and all } \theta_3,\ \theta_3' \\
\text{IRs} &: \quad V_1(\theta_1) \geq 0,\ V_2(\theta_2|\theta_1) \geq 0 \text{ and } V_3(\theta_3|\theta_1, \theta_2) \geq 0 \text{ for all } \theta_1,\ \theta_2,\ \theta_3
\end{aligned}$$

where $V_1(\theta_1', \theta_1)$ represents the buyer's expected payoff at the contracting stage when he reports his type to be θ_1' conditional on his true type θ_1. The other $V_2(\cdot)$ and $V_3(\cdot)$ have similar meanings. Specifically,

$$\begin{aligned}
V_3(\theta_3', \theta_3|\theta_1, \theta_2) &= \theta_3 x_3(\theta_3'|\theta_1, \theta_2) - t_3(\theta_3'|\theta_1, \theta_2) \\
V_2(\theta_2', \theta_2|\theta_1) &= \theta_2 x_2(\theta_2'|\theta_1) - t_2(\theta_2'|\theta_1) + \beta\delta E_{\theta_3}\left[\theta_3 x_3(\theta_3|\theta_1, \theta_2') - t_3(\theta_3|\theta_1, \theta_2')\,|\theta_1, \theta_2\right] \\
V_1(\theta_1', \theta_1) &= \theta_1 x_1(\theta_1') - t_1(\theta_1') + \beta\delta E_{\theta_2}\left[\theta_2 x_2(\theta_2|\theta_1') - t_2(\theta_2|\theta_1')\,|\theta_1\right] \\
&\quad + \beta\delta^2 E_{\theta_2, \theta_3}\left[\theta_3 x_3(\theta_3|\theta_1', \theta_2) - t_3(\theta_3|\theta_1', \theta_2)\,|\theta_1\right]
\end{aligned}$$

Under the IR and IC constraints, the buyer truthfully reports his type on the equilibrium path. The seller's mechanism induces a dynamic Bayesian game, and I use perfect Bayesian equilibrium (PBE) as the solution concept for this paper.

3.3 A Binary-Type Model

In this section, I will focus on the situation in which the type space is binary, and the analysis for a more general model is provided in the subsequent section. Each period, there are two possible types, $\overline{\theta}$ and $\underline{\theta}$. I let $\Delta\theta = \overline{\theta} - \underline{\theta}$. The buyer's type evolves according to a discrete Markov chain. The probability of a high type conditional on a high type is α while the probability of a low type conditional on a low type is γ. I further assume that types are positively correlated, i.e., $\Pr\left(\overline{\theta}|\overline{\theta}\right)-\Pr\left(\overline{\theta}|\underline{\theta}\right)=\alpha+\gamma-1\geq 0$. In period 1, the seller has a prior $\left(\mu_{\overline{\theta}},\mu_{\underline{\theta}}\right)=(\mu,1-\mu)$ on the buyer's type in period 1. As before, the efficient allocation in period s is exactly θ_s. For now, I assume that $\beta\left(\alpha+\gamma-1\right)-\left(1-\beta\right)\left(1-\gamma\right)=\alpha\beta+\gamma-1\geq 0$. This corresponds to the assumption that the overall effect of θ_1 on θ_3 is positive in Section 3.4.

The setup of the binary-type model is very similar to the model in [5] except that the agent is present-biased, and the periods are finite. In the study of the dynamic model in [5], it is assumed that every type is served with a positive quantity, which is guaranteed by the assumption that $\Delta\theta$ cannot be too large. I continue to maintain this assumption here, and the specific condition can be easily derived from the optimal allocations. Before presenting my results, I will introduce the optimal contract for the time consistent agent first, which follows from [5].

Let h_s represent the history of reported types at period s. h_s is defined as $h_1=\varnothing$ and $h_s=\left\{h_{s-1},\theta'_{s-1}\right\}$ where θ'_{s-1} is the reported type in period $s-1$. H_s refers to the set of possible histories at time s. Let $h_s^L=\{\underline{\theta},\underline{\theta},\ldots,\underline{\theta}\}$ denote the history of type where the agent reports $\underline{\theta}$ in the previous $s-1$ periods. I also refer to h_s^L as the lower branch. Under the optimal contract, the seller's surplus is maximized subject to the IR and the IC constraints.

Proposition 13 *When $\beta=1$, the allocation rule in the optimal contract is characterized by the following supply function, given the reported type is θ.*

$$x_s^*(\theta|h_s)=\begin{cases}\overline{\theta}, & \text{if } \theta=\overline{\theta}\\ \underline{\theta}, & \text{if } \theta=\underline{\theta} \text{ and } h_s\in H_s\backslash h_s^L\\ \underline{\theta}-\Delta\theta\frac{\mu}{1-\mu}\left[\frac{\alpha+\gamma-1}{\gamma}\right]^{s-1}, & \text{if } \theta=\underline{\theta} \text{ and } h_s=h_s^L\end{cases}$$

From the allocation rule, we can see that the allocation becomes efficient once the agent reports himself to be a high type. However, this property fails when the agent is present-biased. Before I present the allocation rule for the present-biased agent, let me introduce a preliminary result. The IR and IC constraints can be simplified.

Lemma 6 *In the binary-type model,*

(1) The IR constraints for the high type $\overline{\theta}$ in the last two periods are redundant.

(2) The IR constraints for the low type $\underline{\theta}$ always hold with equality in the last two periods.

(3) The IC constraints for the high type $\overline{\theta}$ always hold with equality in the last two periods.

The results in Lemma 6 follows from the *single crossing property* in the signaling game. I solve the model under the IC constraints for the high type and the IR constraints for the low type primarily. This corresponds to the first-order contract in [6]. I still need to check the IC constraints for the low type in all periods and the IR constraint for the high type in period 1. I will come to that part later.

Definition 2 *A contract is called first-order optimal if it maximizes the seller's profit under the IR constraints for the low type and the IC constraints for the high type in all periods.*

Proposition 14 *The allocation rule in the first-order optimal contract is as follows. For any discount factor δ, the allocation for the high type is always at the efficient level $\overline{\theta}$. However, the allocation for the low type $\underline{\theta}$ is always distorted downwards. The payment scheme is pinned* *done by the IR constraints for the low type and the IC constraints for the high type.*

$$\begin{aligned}
x_1\left(\overline{\theta}\right) &= x_2\left(\overline{\theta}|\theta_1\right) = x_3\left(\overline{\theta}|\theta_1,\theta_2\right) = \overline{\theta} \\
x_1\left(\underline{\theta}\right) &= \underline{\theta} - \frac{\mu}{1-\mu}\Delta\theta \\
x_2\left(\underline{\theta}|\overline{\theta}\right) &= \underline{\theta} - \frac{\alpha}{1-\alpha}\left(1-\beta\right)\Delta\theta \\
x_2\left(\underline{\theta}|\underline{\theta}\right) &= \underline{\theta} - \beta\frac{\mu\left(\alpha+\gamma-1\right)}{\left(1-\mu\right)\gamma}\Delta\theta - \left(1-\beta\right)\frac{1-\gamma}{\gamma}\Delta\theta \\
x_3\left(\underline{\theta}|\overline{\theta},\overline{\theta}\right) &= \underline{\theta} - \left(1-\beta\right)^2\frac{\alpha}{1-\alpha}\Delta\theta \\
x_3\left(\underline{\theta}|\overline{\theta},\underline{\theta}\right) &= \underline{\theta} - \beta\left(1-\beta\right)\frac{\alpha\left(\alpha+\gamma-1\right)}{\left(1-\alpha\right)\gamma}\Delta\theta - \left(1-\beta\right)^2\frac{1-\gamma}{\gamma}\Delta\theta \\
x_3\left(\underline{\theta}|\underline{\theta},\overline{\theta}\right) &= \underline{\theta} - \beta\left(1-\beta\right)\frac{\mu\alpha\left(\alpha+\gamma-1\right)}{\left(1-\mu\right)\left(1-\gamma\right)\left(1-\alpha\right)}\Delta\theta - \left(1-\beta\right)^2\frac{\alpha}{1-\alpha}\Delta\theta \\
x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right) &= \underline{\theta} - \beta^2\frac{\mu\left(\alpha+\gamma-1\right)^2}{\left(1-\mu\right)\gamma^2}\Delta\theta + \beta\left(1-\beta\right)\frac{\mu\left(1-\gamma\right)\left(\alpha+\gamma-1\right)}{\left(1-\mu\right)\gamma^2}\Delta\theta \\
&\quad -\beta\left(1-\beta\right)\frac{\left(\alpha+\gamma-1\right)\left(1-\gamma\right)}{\gamma^2}\Delta\theta - \left(1-\beta\right)^2\frac{1-\gamma}{\gamma}\Delta\theta
\end{aligned}$$

It is useful to compare the allocations for the present-biased agent with the ones for the time consistent agent. Remember that when the agent is time consistent, the contract becomes efficient in all following periods as soon as the agent reports a high type. The allocation is inefficient only when the agent reports the low type repeatedly in the history. Distortions are introduced only to extract more surplus from the high types; in addition, the information rent paid to the high type depends on the allocations in the following periods. Since the IC constraint for the high type is binding, the high type would end up with the same surplus if he falsely reported a low type. Thus, only the allocations for persistent low types in the lower branch are distorted downwards. The quantities are chosen efficiently conditional on a high type report.

However, this result only holds partially when the agent is present-biased. The agent still receives an efficient allocation when he reports a high type. This is not surprising as the principal distorts the allocations in order to induce truthful revelation of the high types. There is no need to distort the allocations for the high types. The allocations for the low types will always be inefficient even there exists a high type report in the history. This is different from the results in [5]. When $\beta \to 1$, the allocations for the present-biased agent converge to the allocations for the time consistent agent.

In a Markovian environment, types are correlated across all periods. The distortion to the allocation in the current period depends on the incentive for truth-telling in the current and previous periods. When $\beta = 1$, utilities are transferable between the buyer and the seller; therefore, the seller can extract the total information rents from the high type in the future periods ex-ante. Moreover, since the type in the current period has an impact on the types in the future periods, there is a causation chain that links the information (type) in each period together, which is the one step chain going from period 1 to period 3 via period 2. This chain is referred to as the “impulse response” in [40] and the “informativeness measure” in [10]. It describes how responsive the buyer’s current type is to his previous type through the type evolution process. In my paper, I also refer to this effect as the conditional effect.

When $\beta < 1$, there is a wedge between the discount factors of the buyer and seller. Here, utilities are not perfectly transferable, and the ex-post participation constraints bind. The seller cannot extract the total information rents in later periods ex-ante. For example, in period 1, he can only extract a proportion, β, of the rent in period 2. For any future period, there is always a residual amount of the information rent at that period left to the buyer; I call this phenomenon the residual effect. Notice that the residual effect happens in all periods except in period 1.

In the game, the buyer discounts period 2 by $\beta\delta$ in period 1 and discounts period 3 by $\beta\delta$ in period 2. This is as if the buyer discounts period 3 by $\beta^2\delta^2$ in period 1 through the conditional effect. However, the buyer discounts period 3 only by $\beta\delta^2$ from the perspective of period 1, and time inconsistency adds another layer to my model. The allocation in period 3 is further distorted to account for the inconsistency, which is through the unconditional effect. The unconditional effect of θ_1 on θ_3 captures how θ_3 is affected by θ_1 without knowing θ_2, and it is the cross-period causation chain going from θ_1 to θ_3. Here, β of θ_3 is responsive to θ_1 through the conditional effect, and $1-\beta$ of it goes through the unconditional effect. A weighted mix of the conditional effect and the unconditional effect captures the informativeness of θ_3 given θ_1 in period 1.

The information chain for the time consistent agent

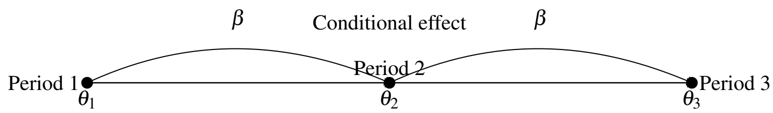

The information chain for the present-biased agent

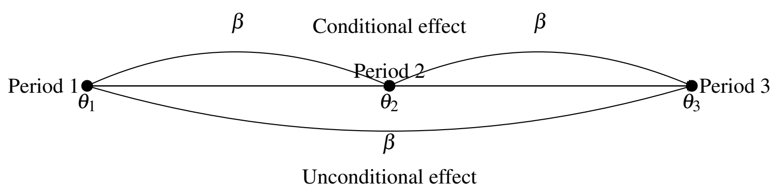

In the binary-type model, only the allocations for the low types are distorted. Through the conditional effect, the unconditional effect and the residual effect, the allocations in the future periods will be affected if the buyer reports a low type in the current period. It is obvious that there is no distortion to the supply in period 3 when the buyer reports $\overline{\theta}$. When the buyer reports $\overline{\theta}$ initially, the allocation in period 1 is not distorted. Here, there is no subsequent distortion to the allocation in period 3 through the conditional effect or the unconditional effect. When the buyer reports $\overline{\theta}$ in period 2, the supply in period 2 is not distorted; in addition, the distortion to the allocation in period 3 through the conditional effect also vanishes. This suggests that the conditional effect on the allocation in period 3 persists only when $\underline{\theta}$ is reported in all three periods. Specifically in $x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right)$, the conditional effect is captured by $\frac{\mu(\alpha+\gamma-1)^2}{(1-\mu)\gamma^2}\Delta\theta$. However, the unconditional effect persists regardless of the intermediate type θ_2. It is described by the term $\frac{\mu\alpha(\alpha+\gamma-1)}{(1-\mu)(1-\gamma)(1-\alpha)}\Delta\theta$ in $x_3\left(\underline{\theta}|\underline{\theta},\overline{\theta}\right)$ and the term $\frac{\mu(1-\gamma)(\alpha+\gamma-1)}{(1-\mu)\gamma^2}\Delta\theta$ in $x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right)$. In summary, the information causation chains going through $\overline{\theta}$ are all inactive. The residual effect persists in period 2 or 3 if $\underline{\theta}$ is reported. Notice that in [5], there is only the conditional effect in the model.

The first-order optimal contract satisfies the other constraints when β is large enough. Here, all other constraints can be transformed into functions of β. In the extreme case where

$\beta = 1$, given any parameters α, γ and $\mu \in (0,1)$, all the other IC and IR constraints hold with strict inequalities under the first-order contract. As the other constraints are continuous in β, when β is sufficiently large, these additional IC and IR constraints are satisfied as well; therefore, the first-order contract maximizes the seller's payoff subject to all the IC and IR constraints. This is formally stated in Proposition 15.

Proposition 15 *For any parameters α, γ and $\mu \in (0,1)$, there exists a $\underline{\beta}$ such that when $\beta \geq \underline{\beta}$, the IC constraints for type $\underline{\theta}$ hold, and the IR constraints for type $\overline{\theta}$ also hold. As a result, the first-order contract is indeed optimal.*

Given β sufficiently large, I can also analyze the seller's expected profit in my model. It turns out that this profit is increasing in β when β sufficiently large. The next proposition summarizes this result.

Proposition 16 *There exists a β' such that the seller's expected profit is an increasing function of β when $\beta \geq \beta'$. Consequently, the seller's expected profit from contracting with a present-biased buyer is strictly less than her expected profit from contracting with a time consistent agent.*

The seller's expected profit is the difference between the expected social surplus and the high-type buyer's information rent. However, without further assumptions on the distribution parameters —i.e., α, γ and μ —it is hard to compare the magnitude of the social surplus and the information rent in my model with those in the model where the agent is time consistent. Moreover, if we take a closer look at the allocations for $\underline{\theta}$ in period 2 and period 3, the allocations are still below the efficient level even when the buyer has previously reported a high type. Given β large enough, the inefficiency in these allocations decreases when β increases. This impact dominates the other indeterminacy in the seller's profit function when β is sufficiently large, which eventually leads to the result in Proposition 16.

3.4 Extensions

3.4.1 A Continuous-Type Model

In this section, I study the model under a continuous type space. The result gives a general description of the optimal contract and provides quantitative measures of the unconditional effect. The standard approach of dynamic mechanism design includes solving a relaxed problem under the local constraints first and finding conditions such that the local constraints are sufficient for implementation. I follow this convention and characterize the three local incentive compatibility constraints in Lemma 9, 10 and 11 in Appendix B. I further transform the original problem using the three local constraints. Lemma 7 shows the tranformed principal's problem.

Lemma 7 *For any direct mechanism that satisfies Lemma 9, Lemma 10 and Lemma 11, the seller's objective function can be written as*

$$\begin{aligned} E\pi \quad = \quad & \int_{\underline{\theta}}^{\overline{\theta}} J_1(\theta_1) f_1(\theta_1) d\theta_1 + \delta \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} J_2(\theta_1,\theta_2) f_2(\theta_2|\theta_1) f_1(\theta_1) d\theta_2 d\theta_1 \\ & + \delta^2 \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} J_3(\theta_1,\theta_2,\theta_3) f_3(\theta_3|\theta_2) f_2(\theta_2|\theta_1) f_1(\theta_1) d\theta_3 d\theta_2 d\theta_1 - R(\underline{\theta}) \end{aligned}$$

where

$$\begin{aligned} J_1(\theta_1) &= \theta_1 x_1 - \frac{1}{2}x_1^2 - \frac{1-F_1(\theta_1)}{f_1(\theta_1)} x_1 \\ J_2(\theta_1,\theta_2) &= \theta_2 x_2 - \frac{1}{2}x_2^2 + \beta\left(-\frac{\partial F_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)}\right)\frac{1-F_1(\theta_1)}{f_1(\theta_1)} x_2 - (1-\beta)\frac{1-F_2(\theta_2|\theta_1)}{f_2(\theta_2|\theta_1)} x_2 \\ J_3(\theta_1,\theta_2,\theta_3) &= \theta_3 x_3 - \frac{1}{2}x_3^2 - \beta^2 \frac{\partial F_3(\theta_3|\theta_2)/\partial\theta_2}{f_3(\theta_3|\theta_2)} \frac{\partial F_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)} \frac{1-F_1(\theta_1)}{f_1(\theta_1)} x_3 \\ & \quad - (\beta-\beta^2)\frac{\partial f_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)} \frac{1-F_3(\theta_3|\theta_2)}{f_3(\theta_3|\theta_2)} \frac{1-F_1(\theta_1)}{f_1(\theta_1)} x_3 \\ & \quad - (\beta-\beta^2)\left(-\frac{\partial F_3(\theta_3|\theta_2)/\partial\theta_2}{f_3(\theta_3|\theta_2)}\right)\frac{1-F_2(\theta_2|\theta_1)}{f_2(\theta_2|\theta_1)} x_3 - (1-\beta)^2 \frac{1-F_3(\theta_3|\theta_2)}{f_3(\theta_3|\theta_2)} x_3 \end{aligned}$$

Before I solve the problem, let me take a look at the components in the seller's surplus. The virtual surplus $J_1(\theta_1)$ in period 1 is equal to the social surplus minus the distortion. The distortion is the familiar inverse hazard rate, and it measures the amount of the information rent conceded to the higher types above θ_1. Through the stochastic dependence of types, θ_1 has a cascade effect on the information rent in all future periods.

For $J_2(\theta_1, \theta_2)$, the term $\frac{\partial F_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)}$ is the " impulse response" in [40]), which captures the marginal effect of θ_1 on θ_2. Moreover, as the buyer discounts his payoff in period 2 further by β, $\frac{\partial F_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)}$ is multiplied by β. The whole term $\beta\left(-\frac{\partial F_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)}\right)\frac{1-F_1(\theta_1)}{f_1(\theta_1)}x_1$ shows the rent conceded in period 2 in order to induce truth-telling in period 1. With greater informativeness, this amount becomes larger. The presence of the other term is due to the residual effect. The information rent for higher types in the second period is $\frac{1-F_2(\theta_2|\theta_1)}{f_2(\theta_2|\theta_1)}x_2$. To leave the agent at least his reservation value 0, the seller can extract up to $\beta\frac{1-F_2(\theta_2|\theta_1)}{f_2(\theta_2|\theta_1)}x_2$ ex-ante, in period 1. The remaining term in the second period virtual surplus is $(1-\beta)\frac{1-F_2(\theta_2|\theta_1)}{f_2(\theta_2|\theta_1)}x_2$. In these two periods, the distortions in the virtual surplus compared to those in the time consistent model ([40]) are due to the disagreement in time preferences of the buyer and the seller. The nature of time inconsistency begins to show itself in the virtual surplus function in period 3.

In period 3, there are three sources of distortions that affect the virtual surplus. The first one is from the conditional effect. In my model, the term $\frac{\partial F_3(\theta_3|\theta_2)/\partial\theta_2}{f_3(\theta_2|\theta_1)}\frac{\partial F_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)}$ represents the conditional effect of θ_1 on θ_3. For a present-biased agent, he discounts the payoff in period 2 further by β from the perspective of period 1, and the payoff in period 3 further by β from the perspective of period 2. The whole term $\beta^2\frac{\partial F_3(\theta_3|\theta_2)/\partial\theta_2}{f_3(\theta_2|\theta_1)}\frac{\partial F_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)}\frac{1-F_1(\theta_1)}{f_1(\theta_1)}x_3$ shows the rent conceded in period 3 to induce truth-telling in period 1, through the conditional effect. The second source is the residual effect. There is the rent $(1-\beta)\frac{1-F_2(\theta_2|\theta_1)}{f_2(\theta_2|\theta_1)}x_2$ left in $J_2(\theta_1, \theta_2)$. Through the impact of θ_2 on θ_3, $\beta(1-\beta)\left(-\frac{\partial F_3(\theta_3|\theta_2)/\partial\theta_2}{f_3(\theta_3|\theta_2)}\right)\frac{1-F_2(\theta_2|\theta_1)}{f_2(\theta_2|\theta_1)}x_3$ is conceded to the buyer, in period 3. Moreover, there is also a pure residual effect in period 3, which is captured by the term $(1-\beta)^2\frac{1-F_3(\theta_3|\theta_2)}{f_3(\theta_3|\theta_2)}x_3$. The last source of distortion is the unconditional effect. From the perspective of period 1, the buyer discounts his payoff in period 3 only by β. As for the total effect of θ_1 on θ_3, $\beta^2\frac{\partial F_3(\theta_3|\theta_2)/\partial\theta_2}{f_3(\theta_2|\theta_1)}\frac{\partial F_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)}$ captures the marginal response of θ_3 to θ_1 through the conditional effect. However, due to the inconsistent time preference, a proportion, $\beta-\beta^2$, of the total effect still exists in period 1. This amount is captured by the unconditional effect of θ_1 on θ_3 without knowing θ_2. It is described by $\frac{\partial F_{3|1}(\theta_3|\theta_1)}{\partial\theta_1}$, and $\frac{\partial f_2(\theta_2|\theta_1)\partial\theta_1}{\partial\theta_1}(1-F_3(\theta_3|\theta_2))$ is the corresponding density function.[3] $\left(\beta-\beta^2\right)\frac{\partial f_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)}\frac{1-F_3(\theta_3|\theta_2)}{f_3(\theta_3|\theta_2)}\frac{1-F_1(\theta_1)}{f_1(\theta_1)}x_3$ is the distortion from the unconditional effect of θ_1 on θ_3.

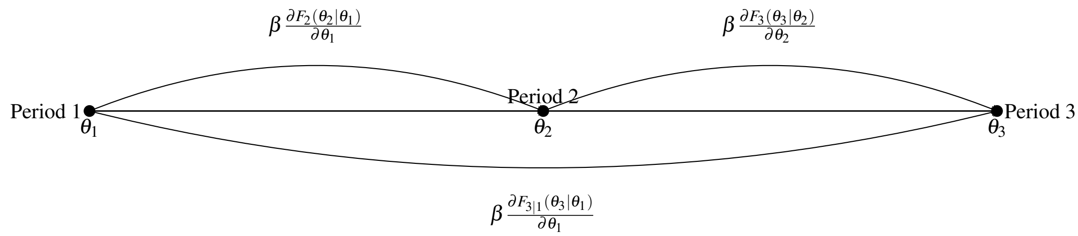

The seller's surplus also depends on the surplus of the lowest types, which is captured by the function $R\left(\underline{\theta}\right)$. The specific expression of $R\left(\underline{\theta}\right)$ can be found in the proof of Lemma 7, and it is an increasing function of $V_1\left(\underline{\theta}\right)$, $V_2\left(\underline{\theta}|\theta_1\right)$ and $V_3\left(\underline{\theta}|\theta_1,\theta_2\right)$. Therefore, I set $V_1\left(\underline{\theta}\right)=0$, $V_2\left(\underline{\theta}|\theta_1\right)=0$ for any θ_1, and $V_3\left(\underline{\theta}|\theta_1,\theta_2\right)=0$ for any θ_1, θ_2. Thus, the IR constraints for the low type always hold.

I further adopt the convention in [10] and assume that the type space $\left[\underline{\theta},\overline{\theta}\right]$ of θ_s is ordered by first-order stochastic dominance (FSD).[4] When $\theta_1>\theta_1'$, it implies that $F_2\left(\theta_2|\theta_1\right)\le F_2\left(\theta_2|\theta_1'\right)$ for all θ_2, and with strict inequality for some θ_2. When $\theta_2>\theta_2'$, I have that $F_3\left(\theta_3|\theta_2\right)\le F_3\left(\theta_3|\theta_2'\right)$ for all θ_3, and with strict inequality for some θ_3. In addition, I further assume that the overall effect of θ_1 on θ_3 is positive, i.e., $\beta\frac{\partial F_3(\theta_3|\theta_2)/\partial\theta_2}{f_3(\theta_3|\theta_2)}\frac{\partial F_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)}+$

[3]Notice that the unconditional effect of θ_1 on θ_3 without knowing θ_2 has the property:

$$\begin{aligned}
1-F_{3|1}\left(\theta_3|\theta_1\right) &= \int_{\theta_3}^{+\infty} f_3\left(s|\theta_1\right)ds=\int_{\theta_3}^{+\infty}\frac{f_{1,3}\left(\theta_1,s\right)}{f_1\left(\theta_1\right)}ds \\
&= \int_{\theta_3}^{+\infty}\frac{\int_{-\infty}^{+\infty} f_3\left(s|\theta_2\right)f_2\left(\theta_2|\theta_1\right)f_1\left(\theta_1\right)d\theta_2}{f_1\left(\theta_1\right)}ds \\
&= \int_{\theta_3}^{+\infty}\int_{-\infty}^{+\infty} f_3\left(s|\theta_2\right)f_2\left(\theta_2|\theta_1\right)d\theta_2 ds \\
&= \int_{-\infty}^{+\infty} f_2\left(\theta_2|\theta_1\right)\left(1-F\left(\theta_3|\theta_2\right)\right)d\theta_2
\end{aligned}$$

Under the differentiability assumptions of the density functions, I have:

$$\frac{\partial F_{3|1}\left(\theta_3|\theta_1\right)}{\partial\theta_1}=-\int_{-\infty}^{+\infty}\frac{\partial f_2\left(\theta_2|\theta_1\right)}{\partial\theta_1}\left(1-F\left(\theta_3|\theta_2\right)\right)d\theta_2$$

[4]Here, the assumption that types are positively correlated in the binary-type model is essentially the discrete version of this FSD assumption.

$(1-\beta)\frac{\partial f_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)}\frac{1-F_3(\theta_3|\theta_2)}{f_3(\theta_3|\theta_2)} \geq 0$. I maximize the seller's expected profit point-wise. One condition to make the local constraints sufficient for implementation is the strong monotonicity of $x_2^*(\theta_2|\theta_1)$ and $x_3^*(\theta_3|\theta_1,\theta_2)$ in all θ_1, θ_2 and θ_3. The optimal mechanism is summarized in the following proposition.

Proposition 17 *Under the assumptions made in this subsection,*

$$\begin{aligned}
x_1^*(\theta_1) &= \max\left\{0, \theta_1 - \frac{1-F_1(\theta_1)}{f_1(\theta_1)}\right\} \\
x_2^*(\theta_2|\theta_1) &= \max\left\{0, \theta_2 - \beta\left(-\frac{\partial F_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)}\right)\frac{1-F_1(\theta_1)}{f_1(\theta_1)} - (1-\beta)\frac{1-F_2(\theta_2)}{f_2(\theta_2)}\right\} \\
x_3^*(\theta_3|\theta_1,\theta_2) &= \max\{0, \theta_1 - \beta^2\frac{\partial F_3(\theta_3|\theta_2)/\partial\theta_2}{f_3(\theta_3|\theta_2)}\frac{\partial F_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)}\frac{1-F_1(\theta_1)}{f_1(\theta_1)} \\
&\quad -\left(\beta-\beta^2\right)\left(-\frac{\partial F_3(\theta_3|\theta_2)/\partial\theta_2}{f_3(\theta_3|\theta_2)}\right)\frac{1-F_2(\theta_2|\theta_1)}{f_2(\theta_2|\theta_1)} - (1-\beta)^2\frac{1-F_3(\theta_3)}{f_3(\theta_3)} \\
&\quad -\left(\beta-\beta^2\right)\frac{\partial f_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)}\frac{1-F_3(\theta_3|\theta_2)}{f_3(\theta_3|\theta_2)}\frac{1-F_1(\theta_1)}{f_1(\theta_1)}\}
\end{aligned}$$

$t_1^*(\theta_1)$, $t_2^*(\theta_2|\theta_1)$ and $t_3^*(\theta_3|\theta_1,\theta_2)$ are constructed from the envelope conditions. If $x_2^*(\theta_2|\theta_1)$ is non-decreasing in both θ_1 and θ_2, and $x_3^*(\theta_3|\theta_1,\theta_2)$ is non-decreasing in all θ_1, θ_2 and θ_3, then the above allocations and payments consist the optimal mechanism.

I provide the proof of Proposition 17 in Appendix C under the strong monotonicity of $x_2^*(\theta_2|\theta_1)$ and $x_3^*(\theta_3|\theta_1,\theta_2)$ in θ_1, θ_2 and θ_3. There can be weaker assumptions, and I will try to figure them out in my future research.

It is not hard to connect the optimal allocations in the binary-type model with the allocations derived in the general model. If I define the discrete inverse hazard rate function as $\frac{1-F(\theta_s|\theta_{s-1})}{P(\theta_s|\theta_{s-1})}\Delta\theta$, the discrete version of $\frac{\partial F_s(\theta_s|\theta_{s-1})/\partial\theta_{s-1}}{f_s(\theta_s|\theta_{s-1})}$ as $\frac{\left[F_s(\theta_s|\theta_{s-1})-F_s\left(\theta_s|\theta'_{s-1}\right)\right]}{P_s(\theta_s|\theta_{s-1})\left[\theta_{s-1}-\theta'_{s-1}\right]}\Delta\theta$, and the discrete version of $\frac{\partial f_s(\theta_s|\theta_{s-1})/\partial\theta_{s-1}}{f_s(\theta_s|\theta_{s-1})}$ as $\frac{P_s(\theta_s|\theta_{s-1})-P_s\left(\theta_s|\theta'_{s-1}\right)}{P_s(\theta_s|\theta_{s-1})\left[\theta_{s-1}-\theta'_{s-1}\right]}\Delta\theta$, it is straightforward to see that the optimal allocations in the binary-type model are consistent with those in the continuous-type model. [5]

3.4.2 A T-Period Model

In this subsection, I present the optimal contract in a T-period binary-type model. Here I still maintain the assumptions that every type is served with a positive quantity, and β is sufficiently large to guarantee that the local constraints are sufficient for implementation. Before I present my result, let me first talk about the notations. $x_s(\theta_s|h_s)$ and $t_s(\theta_s|h_s)$ are the corresponding allocation and transfer in period s, conditional on the history of reported types h_s. For any h_s, let N be the total number of $\underline{\theta}$ reported in the history before period s. I use n to represent the nth $\underline{\theta}$ in the history, and m_n labels the exact period when the nth $\underline{\theta}$ was reported. For example, if the reported types are $\left\{\overline{\theta},\overline{\theta},\underline{\theta},\overline{\theta},\overline{\theta},\underline{\theta},\underline{\theta}\right\}$ by period 7, there are two $\underline{\theta}$ in the history. For the first $\underline{\theta}$, $m_1 = 3$; for the second $\underline{\theta}$, $m_2 = 6$.

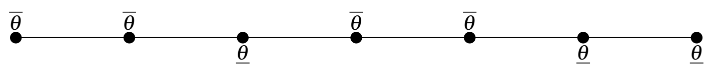

Furthermore, let k_n denote the inverse hazard rate in period m_n when the nth $\underline{\theta}$ is reported. Let $p_{n,n+i}$ denote the general informativeness measure of $\theta_{m_{n+i}}$ given θ_{m_n}. When period m_n and period m_{n+i} are adjacent periods, $p_{n,n+i}$ measures the conditional effect of θ_{m_n} on $\theta_{m_{n+i}}$. If

[5] Notice that the re-defined discrete counterparts take the following forms.

$$\frac{1-F(\theta_s|\theta_{s-1})}{P(\theta_s|\theta_{s-1})} = \begin{cases} 0, & \text{if } \theta_s = \overline{\theta} \\ \frac{\mu}{1-\mu}, & \text{if } \theta_s = \underline{\theta}, s = 1 \\ \frac{\alpha}{1-\alpha}, & \text{if } \theta_s = \underline{\theta}, \theta_{s-1} = \overline{\theta} \\ \frac{1-\gamma}{\gamma}, & \text{if } \theta_s = \underline{\theta}, \theta_{s-1} = \underline{\theta} \end{cases}$$

$$\frac{\left[F_s(\theta_s|\theta_{s-1}) - F_s\left(\theta_s|\theta'_{s-1}\right)\right]\Delta\theta}{P_s(\theta_s|\theta_{s-1})\left[\theta_{s-1}-\theta'_{s-1}\right]} = \begin{cases} 0, & \text{if } \theta_s = \overline{\theta} \\ -\frac{\alpha+\gamma-1}{1-\alpha}, & \text{if } \theta_s = \underline{\theta}, \theta_{s-1} = \overline{\theta} \\ -\frac{\alpha+\gamma-1}{\gamma}, & \text{if } \theta_s = \underline{\theta}, \theta_{s-1} = \underline{\theta} \end{cases}$$

$$\frac{P_s(\theta_s|\theta_{s-1}) - P_s\left(\theta_s|\theta'_{s-1}\right)}{P_s(\theta_s|\theta_{s-1})\left[\theta_{s-1}-\theta'_{s-1}\right]} = \begin{cases} \frac{\alpha+\gamma-1}{\alpha\Delta\theta}, & \text{if } \theta_s = \overline{\theta}, \theta_{s-1} = \overline{\theta} \\ \frac{\alpha+\gamma-1}{(1-\gamma)\Delta\theta}, & \text{if } \theta_s = \overline{\theta}, \theta_{s-1} = \underline{\theta} \\ -\frac{\alpha+\gamma-1}{(1-\alpha)\Delta\theta}, & \text{if } \theta_s = \underline{\theta}, \theta_{s-1} = \overline{\theta} \\ -\frac{\alpha+\gamma-1}{\gamma\Delta\theta}, & \text{if } \theta_s = \underline{\theta}, \theta_{s-1} = \underline{\theta} \end{cases}$$

there are intermediate periods between period m_n and period m_{n+i}, $p_{n,n+i}$ corresponds to the unconditional effect. $p_{n,n+i}$ can also be viewed as the generalized "impulse response".

Proposition 18 *There exists $\underline{\beta}_T$ such that for $\beta \geq \underline{\beta}_T$, the allocation rule listed below maximizes the seller's expected profit under all the IR and IC constraints. The allocation for the high type is always at the efficient level $\overline{\theta}$. However, the allocation for the low type is always distorted.*

$$\begin{aligned} x_s\left(\overline{\theta}|h_s\right) &= \overline{\theta} \\ x_s\left(\underline{\theta}|h_s\right) &= \underline{\theta} - \sum_{n=1}^{N} k_n A_n - k_{N+1}\left(1-\beta\right)^{s-1} \end{aligned}$$

where N is the total number of $\underline{\theta}$ in the history up to period s, and

$$\begin{aligned} A_n &= \sum_{j=0}^{N-n} \beta^{j+1}\left(1-\beta\right)^{s-j-2} \sum_{\substack{\sigma_1,\ldots\sigma_j\in\{n+1,\ldots,N\} \\ \sigma_1<\sigma_2<\cdots<\sigma_j}} p_{n,\sigma_1} p_{\sigma_1,\sigma_2}\cdots p_{\sigma_{j-1},\sigma_j} p_{\sigma_j,N+1} \\ k_1 &= \begin{cases} \frac{\alpha}{1-\alpha}, & \text{if } \theta_1 = \overline{\theta} \\ \frac{\mu}{1-\mu}, & \text{if } \theta_1 = \underline{\theta} \end{cases} \\ k_n &= \begin{cases} \frac{\alpha}{1-\alpha}, & \text{if } \theta_{m_n-1} = \overline{\theta} \\ \frac{1-\gamma}{\gamma},, & \text{if } \theta_{m_n-1} = \underline{\theta} \end{cases} \\ p_{n,n+i} &= \begin{cases} \frac{(\alpha+\gamma-1)\alpha}{(1-\gamma)(1-\alpha)}, & \text{if } \theta_{m_n+1} = \overline{\theta}, i=1 \\ \frac{\alpha+\gamma-1}{\gamma}, & \text{if } \theta_{m_n+1} = \underline{\theta}, i=1 \\ \frac{\alpha+\gamma-1}{1-\gamma} k_{n+i}, & \text{if } \theta_{m_n+1} = \overline{\theta}, i>1 \\ -\frac{\alpha+\gamma-1}{\gamma} k_{n+i}, & \text{if } \theta_{m_n+1} = \underline{\theta}, i>1 \end{cases} \end{aligned}$$

In the binary-type model, the Markov chain has stationary transition probabilities; thus, the informativeness measure of $\theta_{m_{n+i}}$ given θ_{m_n} can be summarized as $p_{n,n+i}$, which is characterized in Proposition 18. Similar to the three-period model, the distortions in the allocations are a weighted sum of the conditional effect, the unconditional effect, and the residual effect. $k_n A_n$ represents the overall distortion in $x_s\left(\underline{\theta}|h_s\right)$ due to inducing truth-telling in period m_n. A_n

is essentially an exhaustive summation of all the information causation chains from θ_{m_n} to θ_s through the conditional and the unconditional effects. The gross amount of the distortion in $x_s(\underline{\theta}|h_s)$ is the summation of all $k_n A_n$ plus the term representing the pure residual effect in period s. The graph below illustrates the existing information causation chains for the former 7-period example.

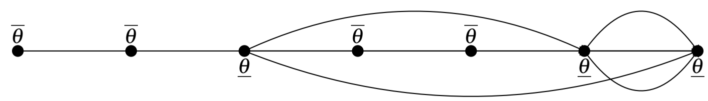

In the T-period model, the *no distortion at the top* principle still holds. When $\theta_s = \underline{\theta}$, the allocation in period s is distorted. Similar to the argument in the three-period model, the information causation chains going through $\overline{\theta}$ are all inactive, and the distortion only depends on the period when the buyer reports a low type.

3.5 Conclusion

This paper extends the existing behavioral literature and studies a model of the long-term contractual relationship between a monopolistic seller and a present-biased buyer. Specifically, the buyer has private information of his type, and the type evolves according to a Markov process. The optimal contract still possesses the *no distortion at the top* property. However, the allocations provided to the low type are never Pareto efficient. When the degree of inconsistency is not severe, the expected profit of the seller from contracting with a present-biased agent is strictly less than from contracting with a time consistent agent.

Another contribution of the paper is the interpretation and characterization of how time inconsistency influences the optimal contract in a long-term contracting environment. Time inconsistency affects the optimal contract through a novel unconditional effect, which is essentially the cross-period marginal effect of the current type on future types in the Markov environment. As the utilities are not perfectly transferable between the two parties, there is also the residual effect capturing the remaining information rent left to the buyer due to the difference in the discount factors.

In the future research, it would be interesting to study the long-term contracting problem with present-biased agents under other market structures, such as a perfect competitive market or an oligopoly market. We may further see how the adverse effect of time inconsistency relates to the degree of competition in the market. Additionally, it would be worthwhile to find more evidence of time inconsistency in real life long-term contracts.

Appendix A

Proofs for Chapter 1

In Appendix A, I provide detailed proofs for most of the propositions, lemmas and theorems in Chapter 1.

Proof of Proposition 1. In the baseline model, the buyer's valuation of the good is $\overline{\theta}$, if the good receives a positive signal from the certifier. As long as $0 < F_s \leq \overline{\theta}$, only the high seller type has an incentive to certify. In the equilibrium, the profit-maximizing certifier picks $F_s \in (0, \overline{\theta}]$.

Firstly, I will show that the high seller type certifies with probability 1. Suppose the high type certifies with probability $\gamma \in (0,1)$ in equilibrium, for a fixed fee $F_s \leq \overline{\theta}$. She can always receive the G signal if she certifies; thus, she charges $\overline{\theta}$ and receives $\overline{\theta} - F_s$ when she certifies. To make the high type willing to randomize, her payoff must be $\overline{\theta} - F_s$ without certification. Let $p' = \overline{\theta} - F_s$. When not certifying, the high type sets a price equal to p'. The low type would mimic the high type and charge p'. Since the buyer cannot differentiate the types without certification, both seller types would end up with the same payoff, which is equal to p'. Accordingly, the seller would have no incentive to exert effort in production. In the subgame initiated by F_s, $e = 0$. From the consistency of beliefs, the buyer would be willing to pay 0 for the good. This indicates $\overline{\theta} - F_s = 0$. Since the seller would not certify, the certifier would have a surplus of 0. Nevertheless, the certifier could slightly lower F_s so that $\overline{\theta} - F_s > p'$. The high type would certify with probability 1 in this situation, and the seller would exert positive effort. The certifier's payoff would be positive. I construct a deviation here. In conclusion, the high type must certify with probability 1, and the equilibrium is fully separating.

Furthermore, the buyer believes that the good is of low type if there is no certification. For a given F_s, the seller's payoff satisfies $U_s = e\left(\overline{\theta} - F_s\right) + (1-e)\,0 - c(e)$, for any effort level $e \in [0,1]$. The seller's first order condition is

$$c'(e) = \overline{\theta} - F_s$$

Given $c''(e) \geq 0$, the second order condition is also satisfied. From the first order condition, I derive $\frac{de}{dF_s} = -\frac{1}{c''(e)}$. As for the certifier, $U_c = eF_s$ for $0 \leq F_s \leq \overline{\theta} - \frac{c(e)}{e}$. The upper bound of F_s is derived from the seller's participation constraint, i.e., $U_s \geq 0$. The certifier's first order condition is

$$e - \frac{F_s}{c''(e)} = 0$$

From the first order condition, I conclude that e_s^* satisfies $c'(e_s^*) = \overline{\theta} - e_s^* c''(e_s^*)$. The optimal certification fee satisfies $F_s^* = e_s^* c''(e_s^*)$. I also need to check the participation constraint is satisfied. Here $U_s = e_s^*\left(\overline{\theta} - F_s^*\right) - c(e_s^*) = e_s^* c'(e_s^*) - c(e_s^*)$. It is an increasing function of e_s^* given $c''(e_s^*) \geq 0$. Therefore, $U_s \geq 0 \times c'(0) - c(0) = 0$. The participation constraint holds.

■

Proof of Proposition 2. In the full model, the certifier picks the signal scheme besides the certification fee. I solve the game backwards. The buyer purchases the good only when $E(\theta|s) \geq p$. Following the discussion in Section 1.4.2, I find that there are two possible equilibrium types. I will show later that the semi-pooling equilibrium does not exist.

One possible equilibrium is a separating equilibrium. In this equilibrium, only the high type opts for certification. The seller types are fully separated. This suggests that $\pi_s(G|\underline{\theta})\, p_G + (1 - \pi_s(G|\underline{\theta}))\, p_B < F_s$ and $p_G = \overline{\theta}$. For a similar analysis established in the proof of Proposition 1, the high type certifies with probability 1. Moreover, $\pi_s\left(G|\overline{\theta}\right) = 1$. Otherwise, $E(\theta|B) = \overline{\theta}$. The low type seller would deviate to certify, which is inconsistent with the assumption.

In this case, the buyer believes the good is low quality without certification. The rest of the construction is the same as in the proof of Proposition 1, and the detail is omitted here. In summary, the seller's equilibrium effort level e_s^* satisfies the condition that $c'(e_s^*) = \overline{\theta} - e_s^* c''(e_s^*)$. The buyer accepts the asking price if the price is $\underline{\theta}$, or if the price is $\overline{\theta}$ and the signal is G. In addition, the certifier charges $e_s^* c''(e_s^*)$ and implements the signal scheme that $\pi_s^*\left(G|\overline{\theta}\right) = 1$ and $\overline{\theta}\pi_s^*(G|\underline{\theta}) < F_s^*$.

Another possible equilibrium is a pooling equilibrium. Both the high type and the low type certify. Suppose the high type charges a price p. The low type seller must charge the same price with positive probability. Suppose the buyer has a belief $(\mu_{sG}, 1-\mu_{sG})$ when the signal is G, and a belief $(\mu_s, 1-\mu_s)$ when there is no signal. Here, $\mu_{sG} > \mu_s$ must hold; otherwise, the signal is useless. If the asking price is equal to p_G, the buyer will purchase the good only when the signal is G. Without certification, the buyer will accept the price only when it is below $\mu_s\overline{\theta} + (1-\mu_s)\underline{\theta}$. The high-type seller will charge $p = p_G$ if she plans to certify. The low-type seller is willing to certify as long as $\pi_s(G|\underline{\theta})p_G - F_s \geq \mu_s\overline{\theta} + (1-\mu_s)\underline{\theta}$. Furthermore, $\pi_s\left(G|\overline{\theta}\right) > \pi_s(G|\underline{\theta})$. Therefore, the certifier picks $F_s = \pi_s(G|\underline{\theta})p_G - \mu_s\overline{\theta} - (1-\mu_s)\underline{\theta}$. Here, the low-type seller is indifferent between certifying or not while the high type strictly prefers certification. Therefore $\mu_s = 0$ and $F_s = \pi_s(G|\underline{\theta})p_G$. For now, suppose the low type certifies with probability 1. Under this construction, the seller's payoff satisfies $U_s = p_G\left(e\pi_s\left(G|\overline{\theta}\right) + (1-e)\pi_s(G|\underline{\theta})\right) - F_s - c(e)$. The first order condition shows that

$$c'(e) = p_G\left(\pi_s\left(G|\overline{\theta}\right) - \pi_s(G|\underline{\theta})\right)$$

From the consistency of beliefs, $\mu_{sG} = \frac{e\pi_s\left(G|\overline{\theta}\right)}{e\pi_s\left(G|\overline{\theta}\right)+(1-e)\pi_s(G|\underline{\theta})}$. The certifier's payoff satisfies $U_C = F_s = \pi_s(G|\underline{\theta})p_G$, where F_s satisfies the participation constraint that $U_s \geq 0$. The optimality conditions are given as follows.

$$\frac{dU_c}{d\pi_s\left(G|\overline{\theta}\right)} = \frac{\left(\overline{\theta} - p_G\right)e\pi_s(G|\underline{\theta})}{e\pi_s\left(G|\overline{\theta}\right) + (1-e)\pi_s(G|\underline{\theta})} + \pi_s(G|\underline{\theta})\frac{dp_G}{de}\frac{de}{d\pi_s\left(G|\overline{\theta}\right)} \tag{A.1}$$

$$\frac{dU_c}{d\pi_s(G|\underline{\theta})} = \frac{e\pi_s\left(G|\overline{\theta}\right)p_G}{e\pi_s\left(G|\overline{\theta}\right)+(1-e)\,\pi_s(G|\underline{\theta})} + \pi_s(G|\underline{\theta})\frac{dp_G}{de}\frac{de}{d\pi_s(G|\underline{\theta})} \tag{A.2}$$

where $\frac{dp_G}{de} = \frac{\left(\overline{\theta}-p_G\right)\pi_s\left(G|\overline{\theta}\right)+p_G\pi_s(G|\underline{\theta})}{e\pi_s\left(G|\overline{\theta}\right)+(1-e)\pi_s(G|\underline{\theta})} \geq 0$. I also derive $\frac{de}{d\pi_s\left(G|\overline{\theta}\right)}$ and $\frac{de}{d\pi_s(G|\underline{\theta})}$.

$$\begin{aligned}\frac{de}{d\pi_s\left(G|\overline{\theta}\right)} &= \left(c''(e)-\frac{\overline{\theta}\pi_s\left(G|\overline{\theta}\right)\pi_s(G|\underline{\theta})\left(\pi_s\left(G|\overline{\theta}\right)-\pi_s(G|\underline{\theta})\right)}{\left(e\pi_s\left(G|\overline{\theta}\right)+(1-e)\,\pi_s(G|\underline{\theta})\right)^2}\right)^{-1} \\ &\quad \times\frac{e\overline{\theta}\left(\pi_s\left(G|\overline{\theta}\right)^2-(1-e)\left(\pi_s\left(G|\overline{\theta}\right)-\pi_s(G|\underline{\theta})\right)^2\right)}{\left(e\pi_s\left(G|\overline{\theta}\right)+(1-e)\,\pi_s(G|\underline{\theta})\right)^2}\end{aligned}$$

$$\frac{de}{d\pi_s(G|\underline{\theta})} = -\left(c''(e)-\frac{\overline{\theta}\pi_s\left(G|\overline{\theta}\right)\pi_s(G|\underline{\theta})\left(\pi_s\left(G|\overline{\theta}\right)-\pi_s(G|\underline{\theta})\right)}{\left(e\pi_s\left(G|\overline{\theta}\right)+(1-e)\,\pi_s(G|\underline{\theta})\right)^2}\right)^{-1}\frac{e\overline{\theta}\pi_s\left(G|\overline{\theta}\right)^2}{\left(e\pi_s\left(G|\overline{\theta}\right)+(1-e)\,\pi_s(G|\underline{\theta})\right)^2}$$

If $c''(e) \geq \frac{\overline{\theta}\pi_s\left(G|\overline{\theta}\right)\pi_s(G|\underline{\theta})\left(\pi_s\left(G|\overline{\theta}\right)-\pi_s(G|\underline{\theta})\right)}{\left(e\pi_s\left(G|\overline{\theta}\right)+(1-e)\pi_s(G|\underline{\theta})\right)^2}$, I get $\pi_s^*\left(G|\overline{\theta}\right)=1$. Otherwise, I have

$$\begin{aligned}\frac{dU_c}{d\pi_s\left(G|\overline{\theta}\right)}+\frac{dU_c}{d\pi_s(G|\underline{\theta})} &= \frac{\left(\overline{\theta}-p_G\right)e\pi_s(G|\underline{\theta})+e\pi_s\left(G|\overline{\theta}\right)p_G}{e\pi_s\left(G|\overline{\theta}\right)+(1-e)\,\pi_s(G|\underline{\theta})} \\ &\quad -\pi_s(G|\underline{\theta})\frac{dp_G}{de}\frac{e\overline{\theta}\frac{(1-e)\left(\pi_s\left(G|\overline{\theta}\right)-\pi_s(G|\underline{\theta})\right)^2}{\left(e\pi_s\left(G|\overline{\theta}\right)+(1-e)\pi_s(G|\underline{\theta})\right)^2}}{c''(e)-\frac{\overline{\theta}\pi_s\left(G|\overline{\theta}\right)\pi_s(G|\underline{\theta})\left(\pi_s\left(G|\overline{\theta}\right)-\pi_s(G|\underline{\theta})\right)}{\left(e\pi_s\left(G|\overline{\theta}\right)+(1-e)\pi_s(G|\underline{\theta})\right)^2}}\end{aligned} \tag{A.3}$$

which is still positive. Therefore, $\pi_s^*\left(G|\overline{\theta}\right)=1$. In equilibrium, I must have $\pi_s^*(G|\underline{\theta}) = \frac{e_s^*\overline{\theta}-e_s^*c'(e_s^*)}{e_s^*\overline{\theta}+(1-e_s^*)c'(e_s^*)}$ and $F_s^* = \pi_s^*(G|\underline{\theta})\,p_G = e_s^*\left(\overline{\theta}-c'(e_s^*)\right)$. The effort level satisfies the condition that $c'(e_s^*) = \overline{\theta}-e_s^*c''(e_s^*)$. I need to check the participation constraint.

$$\begin{aligned}U_s &= p_G^*\left(e_s^*\pi_s^*\left(G|\overline{\theta}\right)+(1-e_s^*)\,\pi_s^*(G|\underline{\theta})\right)-\pi_s^*(G|\underline{\theta})\,p_G^*-c(e_s^*) \\ &= p_G^*e_s^*\left(\pi_s^*\left(G|\overline{\theta}\right)-\pi_s^*(G|\underline{\theta})\right)-c(e_s^*) \\ &= e_s^*c'(e_s^*)-c(e_s^*)\end{aligned}$$

This is positive from the proof of Proposition 1. The participation constraint is also satisfied.

The last step is to show that the low type must certify with probability 1. Notice that the low type opts for certification only when she charges p_G. Suppose she certifies with probability $\gamma < 1$. This indicates that she trades at p_G with probability $\gamma\pi_s\left(G|\underline{\theta}\right)$. The certifier's payoff is $U_c = \left(e + (1-e)\gamma\right)\pi_s\left(G|\underline{\theta}\right)p_G$ where $p_G = \frac{\overline{\theta}e\pi_s\left(G|\overline{\theta}\right)}{e\pi_s\left(G|\overline{\theta}\right)+\gamma(1-e)\pi_s\left(G|\underline{\theta}\right)}$. The seller's payoff is

$$U_s = p_G\left(e\pi_s\left(G|\overline{\theta}\right) + (1-e)\gamma\pi_s\left(G|\underline{\theta}\right)\right) - F_s - c(e)$$

The seller's first order condition is $c'(e) = p_G\left(\pi_s\left(G|\overline{\theta}\right) - \gamma\pi_s\left(G|\underline{\theta}\right)\right)$. Similar to the case where the low type certifies with probability 1, I can solve the optimal signal scheme $\pi_s^*\left(G|\overline{\theta}\right)(\gamma)$ and $\pi_s^*\left(G|\underline{\theta}\right)(\gamma)$, which are functions of γ. The certifier's payoff, U_c, is also a continuous function of γ for $\gamma \in [0,1]$. Using the Envelope theorem, I have

$$\frac{dU_c}{d\gamma} = (1-e)\pi_s^*\left(G|\underline{\theta}\right)(\gamma)\,p_G\frac{e\left(\pi_s^*\left(G|\overline{\theta}\right)(\gamma) - \pi_s^*\left(G|\underline{\theta}\right)(\gamma)\right)}{e\pi_s^*\left(G|\overline{\theta}\right)(\gamma) + (1-e)\gamma\pi_s^*\left(G|\underline{\theta}\right)(\gamma)}$$

For any $e > 1$, $\frac{dU_c}{d\gamma}$ is strictly positive. Therefore, the certifier could implement the signal scheme $\pi_s^*\left(G|\overline{\theta}\right)$ and $\pi_s^*\left(G|\underline{\theta}\right)$, and a fee $F_s^* - \varepsilon$, which are derived under $\gamma = 1$. Certification is always obtained under this mechanism. For a small enough ε, this mechanism generates a higher surplus for the certifier than the current mechanism does. Therefore, this type of equilibrium must be a completely pooling one.

The two equilibrium types generate the same equilibrium effort level, so do the certifier's surplus. Therefore, they both exist.

■

Proof of Corollary 1. This result is directed obtained from $c'\left(e_s^*\right) = \overline{\theta} - e_s^* c''\left(e_s^*\right)$. ■

Proof of Proposition 3. The proof is done in the analysis of Section 1.5.1, and it is omitted here. ■

Proof of Proposition 4. In any equilibrium of the full model, the buyer certifies only when $U_b(p|c_b = 1) \geq U_b(p|c_b = 0)$. If the buyer certifies, he will accept the good only when the signal is good. There are still two equilibrium types. In one type of equilibrium, the buyer's belief is so pessimistic that he believes the good definitely has bad quality ($\mu_b = 0$). In this type of equilibrium, certification is not acquired, which is similar to the baseline model.

In the other equilibrium type, the certifier is hired. For now, let me consider the situation in which the certifier is hired with probability 1. I will show the uniqueness of this equilibrium later. Let the high type's asking price be p. The low-type seller wants to mimic the high type's asking price. If $\pi(\underline{\theta}|G) > 0$, the low type must do it with probability 1. If $\pi(\underline{\theta}|G) = 0$, the low type is indifferent between p and 0. For now, let me construct the equilibrium where the low type always charges p. I will prove that it is the only equilibrium later.

Following the argument about the pricing strategies of the seller and the certifier in Section 1.5.2, I have $U_b\left(\pi_b\left(\overline{\theta}|G\right)\overline{p}|c_b = 1\right) = U_b\left(\pi_b\left(\overline{\theta}|G\right)\overline{p}|c_b = 0\right)$. The fee has to satisfy the following expression.

$$F_b = \mu_b \pi_b\left(G|\overline{\theta}\right)\overline{\theta} - \frac{\mu_b \pi_b\left(G|\overline{\theta}\right) + (1-\mu_b)\pi_b\left(G|\underline{\theta}\right)}{\pi_b\left(G|\overline{\theta}\right) + \left(1-\pi_b\left(G|\overline{\theta}\right)\right)\left(\mu_b \pi_b\left(G|\overline{\theta}\right) + (1-\mu_b)\pi_b\left(G|\underline{\theta}\right)\right)}\mu_b\overline{\theta} \quad \text{(A.4)}$$

This indicates that $p = \overline{p}$ on the equilibrium path. The optimality condition of the seller can be characterized as follows.

$$c'(e) = \overline{p}\left(\pi_b\left(G|\overline{\theta}\right) - \pi_b\left(G|\underline{\theta}\right)\right) \quad \text{(A.5)}$$

From the consistency of beliefs, I have $\mu_b = e$ and $p_G = \frac{e\overline{\theta}\pi\left(G|\overline{\theta}\right)+(1-e)\underline{\theta}\pi(G|\underline{\theta})}{e\pi\left(\overline{\theta}|G\right)+(1-e)\pi(G|\underline{\theta})}$. I derive the seller's effort choice using equation (A.5). The certifier's payoff is $U_c = F_b$. There are two optimality conditions presented as follows.

$$\frac{dU_c}{d\pi_b\left(G|\overline{\theta}\right)} = e\overline{\theta}\left(1 - \frac{\left(e\pi_b\left(G|\overline{\theta}\right) + (1-e)\pi_b\left(G|\underline{\theta}\right)\right)^2 - (1-e)\pi_b\left(G|\underline{\theta}\right)}{\left(\pi_b\left(G|\overline{\theta}\right) + \left(1-\pi_b\left(G|\overline{\theta}\right)\right)\left(e\pi_b\left(G|\overline{\theta}\right) + (1-e)\pi_b\left(G|\underline{\theta}\right)\right)\right)^2}\right) + \frac{dU_c}{de}\frac{de}{d\pi_b\left(G|\overline{\theta}\right)} \quad \text{(A.6)}$$

$$\frac{dU_c}{d\pi_b(G|\underline{\theta})} = -\frac{(1-e)\pi_b(G|\overline{\theta})e\overline{\theta}}{\left(\pi_b(G|\overline{\theta}) + (1-\pi_b(G|\overline{\theta}))\left(e\pi_b(G|\overline{\theta}) + (1-e)\pi_b(G|\underline{\theta})\right)\right)^2} + \frac{dU_c}{de}\frac{de}{d\pi_b(G|\underline{\theta})} \tag{A.7}$$

In addition, $\frac{dU_c}{de} = \overline{\theta}(1-2e)\left(\pi_b\left(G|\overline{\theta}\right) - \pi_b(G|\underline{\theta})\right)$. $\frac{de}{d\pi_b\left(G|\overline{\theta}\right)}$ and $\frac{de}{d\pi_b(G|\underline{\theta})}$ are derived as follows.

$$\frac{de}{d\pi_b\left(G|\overline{\theta}\right)} = M^{-1}\frac{e\left(\pi_b\left(G|\overline{\theta}\right) - \pi_b(G|\underline{\theta})\right)^2 + \pi_b(G|\underline{\theta})(2-\pi_b(G|\underline{\theta}))}{\left(\pi_b\left(G|\overline{\theta}\right) + \left(1-\pi_b\left(G|\overline{\theta}\right)\right)\left(e\pi_b\left(G|\overline{\theta}\right) + (1-e)\pi_b(G|\underline{\theta})\right)\right)^2}e\overline{\theta} \tag{A.8}$$

$$\frac{de}{d\pi_b(G|\underline{\theta})} = -M^{-1}\frac{\pi_b\left(G|\overline{\theta}\right)\left(2-\pi_b\left(G|\overline{\theta}\right)\right)}{\left(\pi_b\left(G|\overline{\theta}\right) + \left(1-\pi_b\left(G|\overline{\theta}\right)\right)\left(e\pi_b\left(G|\overline{\theta}\right) + (1-e)\pi_b(G|\underline{\theta})\right)\right)^2}e\overline{\theta} \tag{A.9}$$

Here, $M = c''(e) - \frac{\left(\pi_b\left(G|\overline{\theta}\right)+\left(1-\pi_b\left(G|\overline{\theta}\right)\right)\pi_b(G|\underline{\theta})\right)\left(\pi_b\left(G|\overline{\theta}\right)-\pi_b(G|\underline{\theta})\right)}{\left(\pi_b\left(G|\overline{\theta}\right)+\left(1-\pi_b\left(G|\overline{\theta}\right)\right)\left(e\pi_b\left(G|\overline{\theta}\right)+(1-e)\pi_b(G|\underline{\theta})\right)\right)^2}\overline{\theta}$. Therefore, $\frac{dU_c}{de}M^{-1}$ determines the sign of the indirect effect. If it is positive, $\pi\left(G|\overline{\theta}\right) = 1$ and $\pi(G|\underline{\theta}) = 0$. If it is negative, I get equation (A.10) by summing up equation (A.6) and equation (A.7).

$$\begin{aligned}\frac{dU_c}{d\pi_b\left(G|\overline{\theta}\right)} + \frac{dU_c}{d\pi_b(G|\underline{\theta})} &= e\overline{\theta}\left(1 - \frac{\left(e\pi_b\left(G|\overline{\theta}\right) + (1-e)\pi_b(G|\underline{\theta})\right)^2 + (1-e)\left(\pi_b\left(G|\overline{\theta}\right) - \pi_b(G|\underline{\theta})\right)}{\left(\pi_b\left(G|\overline{\theta}\right) + \left(1-\pi_b\left(G|\overline{\theta}\right)\right)\left(e\pi_b\left(G|\overline{\theta}\right) + (1-e)\pi_b(G|\underline{\theta})\right)\right)^2}\right) \\ &\quad -\frac{dU_c}{de}M^{-1}K \end{aligned} \tag{A.10}$$

where $K = \frac{\left(\pi_b\left(G|\overline{\theta}\right)-\pi_b(G|\underline{\theta})\right)\left(2-\left(e\pi_b\left(G|\overline{\theta}\right)+(1-e)\pi_b(G|\underline{\theta})\right)-\pi_b\left(G|\overline{\theta}\right)\right)}{\left(\pi_b\left(G|\overline{\theta}\right)+\left(1-\pi_b\left(G|\overline{\theta}\right)\right)\left(e\pi_b\left(G|\overline{\theta}\right)+(1-e)\pi_b(G|\underline{\theta})\right)\right)}e\overline{\theta}$. When $\frac{dU_c}{de}M^{-1} \leq 0$, since $\frac{\left(e\pi_b\left(G|\overline{\theta}\right)+(1-e)\pi_b(G|\underline{\theta})\right)^2+(1-e)\left(\pi_b\left(G|\overline{\theta}\right)-\pi_b(G|\underline{\theta})\right)}{\left(\pi_b\left(G|\overline{\theta}\right)+\left(1-\pi_b\left(G|\overline{\theta}\right)\right)\left(e\pi_b\left(G|\overline{\theta}\right)+(1-e)\pi_b(G|\underline{\theta})\right)\right)^2} \leq 1$, equation (A.10) is positive. This indicates that $\frac{dU_c}{d\pi_b\left(G|\overline{\theta}\right)} > 0$ and $\frac{dU_c}{d\pi_b(G|\underline{\theta})} \leq 0$. Therefore, I still have $\pi_b\left(G|\overline{\theta}\right) = 1$. I can simplify equation (A.7) further.

$$\frac{dU_c}{d\pi_b(G|\underline{\theta})} = -(1-e)e\overline{\theta} - \frac{\overline{\theta}(1-2e)(1-\pi_b(G|\underline{\theta}))}{c''(e) - (1-\pi_b(G|\underline{\theta}))\overline{\theta}}e\overline{\theta} \tag{A.11}$$

In addition, $U_b(\overline{p}|c_b = 1) = U_b(\overline{p}|c_b = 0) = 0$ shows that $\overline{p} = e\overline{\theta}$. Whether there is an interior solution or a corner solution depends on $c(e)$. Note that $\tilde{e}$ denotes the effort level at the corner solution where $c'(\tilde{e}) = \tilde{e}\overline{\theta}$. If $\frac{dU_c}{d\pi_b(G|\underline{\theta})} > 0$ when $\pi_b(G|\underline{\theta}) = 0$, I only have the interior solution. This essentially requires $\varepsilon_{c'(e)} \in \left(\frac{\tilde{e}}{1-\tilde{e}}, 1\right)$ (or $\varepsilon_{c'(e)} \in \left(1, \frac{\tilde{e}}{1-\tilde{e}}\right)$). The derived signal scheme satisfies $\pi_b^*(G|\underline{\theta}) = 1 - \frac{c''(e_b^*)(1-e_b^*)}{e_b^*\overline{\theta}}$, and the equilibrium effort level satisfies $c'(e_b^*) = (1-e_b^*)c''(e_b^*)$. I also have $U_c^* = (1-e_b^*)c'(e_b^*)$. In other situations, there may exist a corner solution where $\pi_b(G|\underline{\theta}) = 0$. Since $c'(e) = e\overline{\theta}(1-\pi_b(G|\underline{\theta}))$, equation (A.11) can be further simplified.

$$\frac{dU_c}{d\pi_b(G|\underline{\theta})} = -e\overline{\theta}\frac{\varepsilon_{c'(e)} - \frac{e}{1-e}}{\varepsilon_{c'(e)-1}}$$

If $\frac{dU_c}{d\pi_b(G|\underline{\theta})} \leq 0$ for all e, I only have the corner solution that $\pi_b^*(G|\underline{\theta}) = 0$. This essentially requires $\varepsilon_{c'(e)} < \min\{\frac{e}{1-e}, 1\}$ or $\varepsilon_{c'(e)} > \max\{\frac{e}{1-e}, 1\}$ for all e. For all other cost functions, I need to compare the certifier's payoff at the interior solution and at the corner solution.

From now on, I will prove the uniqueness of this equilibrium. Firstly, suppose the low type mimics the high type's asking price p with probability $\alpha < 1$. In this case, the buyer must certify with probability 1. Otherwise, the low type would charge p with probability 1. Let $\left(\widetilde{\mu_b}, 1-\widetilde{\mu_b}\right)$ denote the buyer's belief when the asking price is p. Let $(\mu_b, 1-\mu_b)$ be the buyer's prior belief after the production process. Using the Bayes's rule, I get that $\widetilde{\mu_b} = \frac{\mu_b}{\mu_b+(1-\mu_b)\alpha}$. Using the same argument derived under $\alpha = 1$, I can show that $\pi_b\left(G|\overline{\theta}\right) = 1$ and $U_c = (e+\alpha(1-e))F_b = e\left(1-\frac{e}{e+\alpha(1-e)}\right)\overline{\theta}(1-\pi_b(G|\underline{\theta})(\alpha))$, where $\pi_b(G|\underline{\theta})(\alpha)$ depends on α. U_c is a continuous and differentiable function of α. Using the Envelop theorem, I have

$$\frac{dU_c}{d\alpha} = \overline{\theta}(1-e)\left(\frac{e}{e+\alpha(1-e)}\right)^2(1-\pi_b(G|\underline{\theta})(\alpha)) \geq 0$$

The certifier's payoff is maximized when $\alpha = 1$. Therefore, the certifier could slightly modify the optimal mechanism $\pi_b^*\left(G|\overline{\theta}\right)(1)$ and $\pi_b^*(G|\underline{\theta})(1)$ derived under $\alpha = 1$. Because U_c is continuous in γ, there exists a mechanism consisting a signal scheme $\pi_b\left(G|\overline{\theta}\right)(1)$ and

$\pi_b(G|\underline{\theta})(1)+\varepsilon$, and a fee F_b such that, for a small enough ε, the generated U_c is arbitrarily close to the U_c^* derived under $\alpha = 1$. This mechanism could lead to a higher surplus for the certifier than the current mechanism where $\alpha < 1$. In conclusion, both seller types must always charge the same price.

Moreover, another possible case is when the buyer certifies with probability $\beta < 1$. In this situation, the low type always charges the same price as the high type does. Using the same argument derived under $\beta = 1$, I can show that $\pi_b\left(G|\overline{\theta}\right) = 1$ and $U_c = \beta\overline{\theta}e(1-e)(1-\pi_b(G|\underline{\theta})(\beta))$ where $\pi_b(G|\underline{\theta})(\beta)$ depends on β. U_c is a continuous and differentiable function of β for $\beta \in [0,1]$. Using the Envelop theorem, I have

$$\frac{dU_c}{d\beta} = \overline{\theta}e(1-e)(1-\pi_b(G|\underline{\theta})(\beta))^2 \geq 0$$

The certifier's payoff is maximized at $\beta = 1$. Therefore, the certifier could slightly modify the certification scheme to be $\pi_b^*\left(G|\overline{\theta}\right)(1)$ and $\pi_b^*(G|\underline{\theta})(1)$, and a fee $F_b^* - \varepsilon$. Under this mechanism, the buyer always certifies. For a small enough ε, this mechanism could generate a higher surplus for the certifier than the current mechanism. Therefore the buyer must certify the good with probability 1. This concludes the proof of the uniqueness.

■

Proof of Corollary 2. In the proof of Proposition 4, I derive the condition of the induced effort level, which is $c'\left(e_b^*\right) = e_b^*\overline{\theta}\left(1-\pi_b^*(G|\underline{\theta})\right) \leq \overline{\theta}$. Therefore, the effort level is below the first-best. ■

Proof of Theorem 1. The welfare comparison depends on the magnitude of e_s^* and e_b^*. The larger one of these two leads to the model with a higher social welfare in equilibrium. From Proposition 1 and 2, I can calculate the equilibrium effort levels in these two models. Specifically, $\varepsilon_{c'(e_s^*)} = \frac{\overline{\theta}}{c'(e_s^*)} - 1$ holds in equilibrium in the seller-certification model, and $\varepsilon_{c'\left(e_b^*\right)} = \frac{e_b^*}{1-e_b^*}$ holds at the interior solution in the buyer-certification model. $\frac{\overline{\theta}}{c'(e)} - 1$ is decreasing in e, while $\frac{e}{1-e}$

is increasing in e. Therefore, $\frac{\overline{\theta}}{c'(e)}-1$ and $\frac{e}{1-e}$ have only one intersection. At the intersection, the effort level satisfies $c(e)=(1-e)\overline{\theta}$. I denote this effort level by e_1, and $B=\frac{e_1}{1-e_1}$. For any $e>e_1$, $\frac{e}{1-e}$ has a higher value than $\frac{\overline{\theta}}{c'(e)}-1$. For any $e<e_1$, $\frac{\overline{\theta}}{c'(e)}-1$ has a higher value than $\frac{e}{1-e}$. Therefore, if $\varepsilon_{c'(e)}>B$, e_b^* is larger than e_s^*; if $\varepsilon_{c'(e)}<B$, e_b^* is smaller than e_s^*. Moreover, in the buyer-certification game, the equilibrium effort level at the corner solution satisfies $c'(\tilde{e})=\tilde{e}\overline{\theta}$. Here, there exists $B'=\frac{1}{\tilde{e}}-1$ such that if $\varepsilon_{c'(e)}>B'$ for all e, $c'(e_s^*)<\frac{\overline{\theta}}{1+B'}=\tilde{e}\overline{\theta}=c'(\tilde{e})$; thus, $e_s^*<\tilde{e}$. If $\varepsilon_{c'(e)}<B'$, $e_s^*>\tilde{e}$. Let $A=\max\{B,B'\}$ and $A'=\min\{B,B'\}$. In summary, if $\varepsilon_{c'(e)}\leq A'$ for all e, the seller-certification model yields a higher social surplus. If $\varepsilon_{c'(e)}\geq A$ for all e, the buyer-certification model yields a higher social surplus.

Regarding the certifier's payoff, using the results derived in the proof of Proposition 4, I find $F_b^*=\left(1-e_b^*\right)c'\left(e_b^*\right)$. Since $c'\left(e_b^*\right)\leq e_b^*\overline{\theta}$ in the buyer-certification model, $F_b^*<e_b^*\left(\overline{\theta}-c'\left(e_b^*\right)\right)$. Nevertheless, in the seller-certification model, the certifier ends up with the payoff $e_s^*\left(\overline{\theta}-c'\left(e_s^*\right)\right)$ where $e_s^*c''\left(e_s^*\right)=\overline{\theta}-c'\left(e_s^*\right)$. As e_s^* maximizes the value of the function $e\left(\overline{\theta}-c'(e)\right)$, the certifier is always better off in the seller-certification game.

■

Proof of Proposition 5. Here, the seller's asking price is based on the actual signal she receives. There are still two possible equilibrium types. The construction of the separating equilibrium is the same as in Proposition 2. I will not repeat it here. The construction of the pooling equilibrium has different optimality conditions. The buyer's prior belief of the good's quality is $(\mu_s,1-\mu_s)$. Conditional on the signal G, the price the seller can charge is $p_G=\frac{\mu_s\pi_s\left(G|\overline{\theta}\right)\overline{\theta}}{\mu_s\pi_s\left(G|\overline{\theta}\right)+(1-\mu_s)\pi_s(G|\underline{\theta})}$; conditional on the signal B, the price the seller can charge is $p_B=\frac{\mu_s\pi_s\left(B|\overline{\theta}\right)\overline{\theta}}{\mu_s\pi_s\left(B|\overline{\theta}\right)+(1-\mu_s)\pi_s(B|\underline{\theta})}$. Conditional on certifying, the seller's payoff satisfies

$$U_s=p_G\left(e\pi_s\left(G|\overline{\theta}\right)+(1-e)\pi_s\left(G|\underline{\theta}\right)\right)+p_B\left(e\pi_s\left(B|\overline{\theta}\right)+(1-e)\pi_s\left(B|\underline{\theta}\right)\right)-c(e)-F_s$$

The optimality condition satisfies

$$c'(e) = (p_G - p_B)\left(\pi_s\left(G|\overline{\theta}\right) - \pi_s\left(G|\underline{\theta}\right)\right)$$

In the subgame that is initiated by the certification scheme, $\mu_s = e$ from the consistency of beliefs. The maximum fee the certifier can charge leaves the low-type seller a payoff of zero. Therefore $U_c = F_s = \sum_{s\in\{G,B\}} p_s \pi_s(s|\underline{\theta})$. The low-type seller must certify with probability 1. Otherwise, the certifier can lower the fee slightly so that certification is guaranteed. The Lagrangian of the certifier's maximization problem is

$$L = \sum_{s\in\{G,B\}} p_s \pi_s(s|\underline{\theta}) + \lambda\left((p_G - p_B)\left(\pi_s\left(G|\overline{\theta}\right) - \pi_s(G|\underline{\theta})\right) - c'(e)\right)$$

Moreover, $\sum_{s\in\{G,B\}} \pi_s(s|\theta) = 1$ for any $\theta \in \left\{\overline{\theta}, \underline{\theta}\right\}$. I derive the following first order conditions:

$$\begin{aligned}\frac{dL}{d\pi_s\left(G|\overline{\theta}\right)} &= \frac{\left(\overline{\theta} - p_G\right) e\pi_s(G|\underline{\theta})}{e\pi_s\left(G|\overline{\theta}\right) + (1-e)\pi_s(G|\underline{\theta})} + \lambda\left[p_G + \frac{\left(\overline{\theta} - p_G\right) e\left(\pi_s\left(G|\overline{\theta}\right) - \pi_s(G|\underline{\theta})\right)}{e\pi_s\left(G|\overline{\theta}\right) + (1-e)\pi_s(G|\underline{\theta})}\right] \\ &\quad - \frac{\left(\overline{\theta} - p_B\right) e\pi_s(B|\underline{\theta})}{e\pi_s\left(B|\overline{\theta}\right) + (1-e)\pi_s(B|\underline{\theta})} - \lambda\left[p_B + \frac{\left(\overline{\theta} - p_B\right) e\left(\pi_s\left(B|\overline{\theta}\right) - \pi_s(B|\underline{\theta})\right)}{e\pi_s\left(B|\overline{\theta}\right) + (1-e)\pi_s(B|\underline{\theta})}\right]\end{aligned}$$

$$\begin{aligned}\frac{dL}{d\pi_s(G|\underline{\theta})} &= \frac{p_G e\pi_s\left(G|\overline{\theta}\right)}{e\pi_s\left(G|\overline{\theta}\right) + (1-e)\pi_s(G|\underline{\theta})} - \frac{\lambda p_G \pi_s\left(G|\overline{\theta}\right)}{e\pi_s\left(G|\overline{\theta}\right) + (1-e)\pi_s(G|\underline{\theta})} \\ &\quad - \frac{p_B e\pi_s\left(B|\overline{\theta}\right)}{e\pi_s\left(B|\overline{\theta}\right) + (1-e)\pi_s(B|\underline{\theta})} + \frac{\lambda p_B \pi_s\left(B|\overline{\theta}\right)}{e\pi_s\left(B|\overline{\theta}\right) + (1-e)\pi_s(B|\underline{\theta})}\end{aligned}$$

If $\lambda \leq 0$ at the solution, I have $\frac{dL}{d\pi_s\left(G|\overline{\theta}\right)} < 0$. This suggests $\pi_s\left(G|\overline{\theta}\right) = 0$, which contradicts the assumption of the signals. Therefore I must have $\lambda \geq 0$. Moreover, I derive the following condition by summing up $\frac{dL}{d\pi_s\left(G|\overline{\theta}\right)}$ and $\frac{dL}{d\pi_s(G|\underline{\theta})}$.

$$\frac{dL}{d\pi_s\left(G|\overline{\theta}\right)}+\frac{dL}{d\pi_s\left(G|\underline{\theta}\right)} = \left(\frac{p_G-e\overline{\theta}}{e\pi_s\left(G|\overline{\theta}\right)+(1-e)\,\pi_s\left(G|\underline{\theta}\right)}-\frac{e\overline{\theta}-p_B}{e\pi_s\left(B|\overline{\theta}\right)+(1-e)\,\pi_s\left(B|\underline{\theta}\right)}\right) \times (e-\lambda)\left(\pi_s\left(G|\overline{\theta}\right)-\pi_s\left(G|\underline{\theta}\right)\right)$$

Here I must have $e=\lambda$ at the optimal condition. Otherwise if $e<\lambda$, $\frac{dL}{d\pi_s(G|\underline{\theta})}>0$. This suggests $\pi_s\left(G|\underline{\theta}\right)=1$, which makes the signal useless. If $e>\lambda$, in contrast, $\frac{dL}{d\pi_s(G|\underline{\theta})}<0$. I have $\pi_s\left(G|\underline{\theta}\right)=0$. However, $\frac{dL}{d\pi_s\left(G|\overline{\theta}\right)}<e\left(\overline{\theta}-p_B\right)-\frac{\left(\overline{\theta}-p_B\right)e\left(1-e\pi_s\left(G|\overline{\theta}\right)\right)}{1-e\pi_s\left(G|\overline{\theta}\right)}=0$. This suggests $\pi_s\left(G|\overline{\theta}\right)=0$. Here, the signal structure is completely noisy, and the induced effort is zero. This cannot be the solution to the maximization problem. Therefore, I must have $e=\lambda$. In addition, there is a simplified optimality condition.

$$\frac{dL}{de}=\overline{\theta}-(p_G-p_B)\left(\pi_s\left(G|\overline{\theta}\right)-\pi_s\left(G|\underline{\theta}\right)\right)-ec''(e)=0$$

which is derived under $e=\lambda$. Since $c'(e)=(p_G-p_B)\left(\pi_s\left(G|\overline{\theta}\right)-\pi_s\left(G|\underline{\theta}\right)\right)$, the equilibrium effort level e_s^* must satisfy $c'(e_s^*)=\overline{\theta}-e_s^*c''(e_s^*)$, which is the same as the one characterized in Proposition 2. Moreover, the fee satisfies $F_s^*=\sum_{s\in\{G,B\}}p_s^*\pi_s^*\left(s|\underline{\theta}\right)$.

$$\begin{aligned}
F_s^* &= p_G^*\pi_s^*\left(G|\underline{\theta}\right)+p_B^*\pi_s^*\left(B|\underline{\theta}\right)\\
&= e_s^*\overline{\theta}\frac{e_s^*\pi_s^*\left(G|\overline{\theta}\right)\left(1-\pi_s^*\left(G|\overline{\theta}\right)\right)+(1-e_s^*)\,\pi_s^*\left(G|\underline{\theta}\right)\left(1-\pi_s^*\left(G|\underline{\theta}\right)\right)}{\left(e_s^*\pi_s^*\left(G|\overline{\theta}\right)+(1-e_s^*)\,\pi_s\left(G|\underline{\theta}\right)\right)\left(e_s^*\pi_s\left(B|\overline{\theta}\right)+(1-e_s^*)\,\pi_s^*\left(B|\underline{\theta}\right)\right)}\\
&= e_s^*\overline{\theta}\left(1-\frac{e_s^*(1-e_s^*)\left(\pi_s^*\left(G|\overline{\theta}\right)-\pi_s^*\left(G|\underline{\theta}\right)\right)^2}{\left(e_s^*\pi_s^*\left(G|\overline{\theta}\right)+(1-e_s^*)\,\pi_s^*\left(G|\underline{\theta}\right)\right)\left(e_s^*\pi_s^*\left(B|\overline{\theta}\right)+(1-e_s^*)\,\pi_s^*\left(B|\underline{\theta}\right)\right)}\right)\\
&= e_s^*\left(\overline{\theta}-c'(e_s^*)\right)
\end{aligned}$$

since $c'(e_s^*)=e_s^*\overline{\theta}\frac{e_s^*(1-e_s^*)\left(\pi_s^*\left(G|\overline{\theta}\right)-\pi_s^*\left(G|\underline{\theta}\right)\right)^2}{\left(e_s^*\pi_s^*\left(G|\overline{\theta}\right)+(1-e_s^*)\pi_s^*\left(G|\underline{\theta}\right)\right)\left(e_s^*\pi_s^*\left(B|\overline{\theta}\right)+(1-e_s^*)\pi_s^*\left(B|\underline{\theta}\right)\right)}$. Therefore, the certifier's payoff is also the same as in Proposition 2. I still need to check the participation constraint of the seller.

$$\begin{aligned}
U_s &= p_G^*\left(e_s^*\pi_s^*\left(G|\overline{\theta}\right)+\left(1-e_s^*\right)\pi_s^*\left(G|\underline{\theta}\right)\right)+p_B^*\left(e_s^*\pi_s^*\left(B|\overline{\theta}\right)+\left(1-e_s^*\right)\pi_s^*\left(B|\underline{\theta}\right)\right)\\
&\quad -\pi_s^*\left(G|\underline{\theta}\right)p_G^*-\pi_s^*\left(B|\underline{\theta}\right)p_B^*-c\left(e_s^*\right)\\
&= \left(p_G^*-p_B^*\right)e_s^*\left(\pi_s^*\left(G|\overline{\theta}\right)-\pi_s^*\left(G|\underline{\theta}\right)\right)-c\left(e_s^*\right)\\
&= e_s^*c'\left(e_s^*\right)-c\left(e_s^*\right)
\end{aligned}$$

Using the argument in the proof of Proposition 1, I conclude that $U_s \geq 0$. The participation constraint is also satisfied. ■

Appendix B

Proofs for Chapter 2

B.1 Proofs for Chapter 2.4

In Appendix B, I provide detailed proofs for the propositions, lemmas and theorems in Chapter 2.

Proof of Lemma 1. Consider an arbitrary equilibrium. In period 2, if a type-q seller accepts the offer p_2 with positive probability, I must have

$$p_2 - \alpha q \geq \delta (Aq - \alpha q) \Rightarrow p_2 \geq \delta Aq + (1-\delta)\alpha q$$

Therefore, for any seller type $q' < q$, $p_2 > \delta Aq' + (1-\delta)\alpha q'$. Any lower type seller must accept p_2 with probability 1.

In period 1, if a type-q seller accepts p_1 with positive probability, she must receive a higher surplus from accepting the offer than from rejecting and concealing it. Let $E(\max\{p_2, Aq\})$ be her expected payoff in period 2 when p_1 is concealed. I must have

$$p_1 - \alpha q \geq \delta (E(\max\{p_2, Aq\}) - \alpha q) \Rightarrow p_1 \geq \delta E(\max\{p_2, Aq\}) + (1-\delta)\alpha q$$

For any seller type $q' < q$, $E(\max\{p_2, Aq\}) \geq E(\max\{p_2, Aq'\})$. The seller's payoff in period 2 is monotonically increasing in q. Therefore, for any seller type $q' < q$, $p_1 > E(\max\{p_2, Aq'\}) + (1-\delta)\alpha q'$. Any lower seller type receives a strictly higher payoff from accepting offer p_1 than from rejecting and concealing it.

Moreover, a seller of type q also receives a higher surplus from accepting p_1 than from rejecting and disclosing it. Similarly, I can show that any lower seller type receives a strictly higher payoff from accepting p_1 than from rejecting and disclosing it. Therefore, any lower type must accept p_1 with probability 1.

The same reasoning also applies to the infinite-horizon game. ■

Proof of Lemma 2. Lemma 2 is proved by contradiction. In this two-period model, I assume that $\overline{q}$ is large enough so the cutoffs in both periods are strictly less than $\overline{q}$. Consider an equilibrium in which buyer 1 plays a mixed strategy. Buyer 2 must also randomize. When buyer 1 randomizes in a set of offers, different offers correspond to different cutoff types. As all offers must yield the same expected payoff, higher offers must be accepted by higher cutoff types. Let S_1 be the set of equilibrium offers buyer 1 randomizes in, and $S_1^C = \left\{p_1 | p_1 \in S_1, d^2\left(q, p^1\right) = \emptyset, \forall q\right\}$. Let $p_1^H = \sup S_1^C$ and $p_1^L = \inf S_1^C$. Let $\widehat{q_1}^H$ be the cutoff type induced by offer p_1^H, and $\widehat{q_1}^L$ be the cutoff type induced by offer p_1^L. For any offer $p_1 \in S_1^C$ and the corresponding cutoff type $\widehat{q_1}$, p_1^H and p_1 give the seller the same payoff. I have $\frac{\widehat{q_1}+\underline{q}}{2} - p_1 = \frac{\widehat{q_1}^H+\underline{q}}{2} - p_1^H$. Thus, I must have $\widehat{q_1}^H$ be the supremum of the cutoff types induced by buyer 1's offer in S_1^C, and $\widehat{q_1}^L$ be the infimum of these cutoffs. Ψ_1 represents the cdf of the equilibrium cutoffs in period 1 induced by offers in S_1^C. In period 2, let S_2 be the set of offers buyer 2 randomizes in if nothing is disclosed; $p_2^H = \sup S_1$, and $p_1^L = \inf S_1$. For any $p_2 \in S_2$, the corresponding cutoff type $\widehat{q_2}$ satisfies the condition that $A\widehat{q_2} = p_2$. Let $\widehat{q_2}^H = \frac{1}{A}p_2^H$ and $\widehat{q_2}^L = \frac{1}{A}p_2^L$. Ψ_2 represents the cdf of the equilibrium cutoffs in period 2 induced by offers in S_2.

Suppose S_1^C has more than 1 element. Buyer 2 randomizes in S_2. For the type-q seller, her expected payoff in period 2 if she discloses nothing is denoted by $V_2^C(q)$, where

$$V_2^C(q) = \int_q^{\widehat{q_2}^H} A\widetilde{q} d\Psi_2(\widetilde{q}) + Aq\Psi_2(q) \tag{B.1}$$

The cutoff $\widehat{q_1}^H$ must satisfy the following condition

$$p_1^H = \delta\left[\int_{\widehat{q_1}^H}^{\widehat{q_2}^H} A\widetilde{q}d\Psi_2(\widetilde{q}) + A\widehat{q_1}^H\Psi_2\left(\widehat{q_1}^H\right)\right] + \alpha(1-\delta)\widehat{q_1}^H \tag{B.2}$$

As $\frac{dV_2^C(q)}{dq} = A\Psi_2(q) \geq 0$, $V_2^C(q)$ is a non-decreasing function in q and strictly increasing in q when $\widehat{q_2}^L < q$. I establish my argument with a few steps.

Step 1: In this step, I will show S_1^C has at most three elements.

Picking any $p_1^M \in S_1^C$, I use $\widehat{q}_1^M$ to denote the corresponding cutoff type that is induced by p_1^M. Since p_1^H and p_1^M give buyer 1 the same payoff, I have $\frac{\widehat{q_1}^M + q}{2} - p_1^M = \frac{\widehat{q_1}^H + q}{2} - p_1^H$. As p_1^M is concealed, I have

$$p_1^M = \delta\left[\int_{\widehat{q_1}^M}^{\widehat{q_2}^H} A\widetilde{q}d\Psi_2(\widetilde{q}) + A\widehat{q_1}^M\Psi_2\left(\widehat{q_1}^M\right)\right] + \alpha(1-\delta)\widehat{q_1}^M \tag{B.3}$$

This suggests

$$p_1^H - \frac{1}{2}\widehat{q_1}^H = \delta\left[\int_{\widehat{q_1}^M}^{\widehat{q_2}^H} A\widetilde{q}d\Psi_2(\widetilde{q}) + A\widehat{q_1}^M\Psi_2\left(\widehat{q_1}^M\right)\right] + \left[\alpha(1-\delta) - \frac{1}{2}\right]\widehat{q_1}^M \tag{B.4}$$

The right hand side of the equation (B.4) is a convex function, and the first order derivative with respect to $\widehat{q_1}^M$ is increasing in $\widehat{q_1}^M$. Therefore, there are at most three elements in S_1^C.

Step 2: In this step I will show that if there is more than 1 element in S_1^C, $p_1^H \notin S_1^C$, and $p_1^H = \frac{\delta A^2}{2A-1}\widehat{q_1}^H + \alpha(1-\delta)\widehat{q_1}^H$. As a result, S_1^C has at most 1 element.

When p_1^H is disclosed, buyer 2 has a degenerate belief of the cutoff type. Suppose his belief of the cutoff type is $\widehat{q_1}'$, and offers price p_2', when p_1^H is disclosed. Given buyer 2's strategy, the seller discloses offer p_1^H if

$$p_2' \geq \int_{q}^{\widehat{q_2}^H} A\widetilde{q}d\Psi_2(\widetilde{q}) + Aq\Psi_2(q) \tag{B.5}$$

Case 1: Suppose $V_2^C\left(q_2^L\right) < p_2' < V_2^C\left(q_2^H\right)$. There exists a threshold type $\widehat{q}'$ such that type $\widehat{q}'$ is indifferent between disclosing p_1^H or concealing p_1^H, and $\widehat{q_2}^L < \widehat{q}' < \widehat{q_2}^H$ holds. Here all seller types below $\widehat{q}'$ prefers disclosing p_1^H and all seller types above $\widehat{q}'$ prefers concealing p_1^H. From the consistency of beliefs, if p_1^H is disclosed, buyer 2's belief of the seller's type would be uniformly distributed in interval $\left[\widehat{q_1}', \widehat{q}'\right]$.

Case 1.1: $\widehat{q}' \geq \frac{A}{2A-1}\widehat{q_1}'$ holds. In this case, buyer 2's payoff is a function of the cutoff type $\widehat{q_2}$ in period 2.

$$\pi_2\left(\widehat{q_2}\right) = \left(\frac{\widehat{q_2} + \widehat{q_1}'}{2} - A\widehat{q_2}\right)\frac{\widehat{q_2} - \widehat{q_1}'}{\widehat{q}' - \widehat{q_1}'}$$

$\pi_2\left(\widehat{q_2}\right)$ is maximized when $\widehat{q_2} = \frac{A}{2A-1}\widehat{q_1}'$. Moreover, given $\widehat{q_1}'$, buyer 2's price offer is $p_2' = \frac{A^2}{2A-1}\widehat{q_1}'$. From the consistency of beliefs,

$$p_1^H = \delta p_2' + \alpha\left(1-\delta\right)\widehat{q_1}' \tag{B.6}$$

I get $\widehat{q_1}' = \frac{p_1^H}{\delta\frac{A^2}{2A-1} + \alpha(1-\delta)}$. Using equation (B.2), I can compare the price p_2' and $V_2^C\left(\widehat{q_1}^H\right)$.

$$\begin{aligned}
p_2' - V_2^C\left(\widehat{q_1}^H\right) &= \frac{A^2}{\delta A^2 + \alpha\left(1-\delta\right)\left(2A-1\right)}\left[\delta V_2^C\left(\widehat{q_1}^H\right) + \alpha\left(1-\delta\right)\widehat{q_1}^H\right] - V_2^C\left(\widehat{q_1}^H\right) \\
&= \frac{\alpha\left(1-\delta\right)\left(2A-1\right)}{\delta A^2 + \alpha\left(1-\delta\right)\left(2A-1\right)}\left[\frac{A^2}{\left(2A-1\right)}\widehat{q_1}^H - V_2^C\left(\widehat{q_1}^H\right)\right] \\
&= \frac{A\alpha\left(1-\delta\right)\left(2A-1\right)}{\delta A^2 + \alpha\left(1-\delta\right)\left(2A-1\right)}\left[\frac{A}{\left(2A-1\right)}\widehat{q_1}^H - \int_{\widehat{q_1}^H}^{\widehat{q_2}^H}\widetilde{q}d\Psi_2\left(\widetilde{q}\right) - \widehat{q_1}^H\Psi_2\left(\widehat{q_1}^H\right)\right]
\end{aligned}$$

If $\widehat{q_1}^H \geq \widehat{q_2}^H$, $p_2' > V_2^C\left(\widehat{q_1}^H\right)$ holds naturally. Type $\widehat{q_1}^H$ will deviate to disclose p_1^H. If $\widehat{q_1}^H < \widehat{q_2}^H$, I will show that $\frac{A}{2A-1}\widehat{q_1}^H \geq \widehat{q_2}^H$. When the cutoff type in period 2 is q_2, buyer 2's payoff function is

$$\pi_2\left(q_2\right) = \int_{(\underline{q})}^{q_2}\int_{\underline{q}}^{q}\frac{q - Aq_2}{1 - \frac{\widetilde{q}-\underline{q}}{\overline{q}-\underline{q}}}d\Psi_1\left(\widetilde{q}\right)dq$$

Since $\widehat{q_2}^H$ is the cutoff in equilibrium, it must satisfy the following first order condition.

$$q_2^H\int_{\underline{q}}^{\widehat{q_1}^H}\frac{1}{1 - \frac{\widetilde{q}-\underline{q}}{\overline{q}-\underline{q}}}d\Psi_1\left(q_1\right) = \frac{A}{2A-1}\int_{\underline{q}}^{\widehat{q_1}^H}\frac{q_1}{1 - \frac{\widetilde{q}-\underline{q}}{\overline{q}-\underline{q}}}d\Psi_1\left(q_1\right) \tag{B.7}$$

From (B.7), I conclude $\frac{A}{2A-1}\widehat{q_1}^H \geq \widehat{q_2}^H$. Therefore $\frac{A}{(2A-1)}\widehat{q_1}^H > \int_{\widehat{q_1}^H}^{\widehat{q_2}^H} \widetilde{q} d\Psi_2(\widetilde{q}) + \widehat{q_1}^H \Psi_2\left(\widehat{q_1}^H\right)$. This suggest that $p_2' - V_2^C\left(\widehat{q_1}^H\right) > 0$. Therefore, type $\widehat{q_1}^H$ will deviate and disclose p_1^H when it is offered. Moreover, I still need to check that $\widehat{q}' \geq \frac{A}{2A-1}\widehat{q_1}'$. Since $p_2' > V_2^C(\widehat{q_1})$ for any $\widehat{q_1} \in \left(\widehat{q_1}^H, \frac{A}{2A-1}\widehat{q_1}^H\right)$, any seller type in $\left(\widehat{q_1}^H, \frac{A}{2A-1}\widehat{q_1}^H\right)$ would disclose offer p_1^H. The assumption that $\widehat{q}' \geq \frac{A}{2A-1}\widehat{q_1}'$ holds since $\widehat{q}' \geq \frac{A}{2A-1}\widehat{q_1}^H \geq \frac{A}{2A-1}\widehat{q_1}'$.

Case 1.2: $\widehat{q}' < \frac{A}{2A-1}\widehat{q_1}'$ holds. In this case, Buyer 2's offer is $p_2' = A\widehat{q}'$. Moreover, since $\widehat{q_2}^L < \widehat{q}' < \widehat{q_2}^H$, I have $V_2^C(\widehat{q}') > A\widehat{q}'$. However, type $\widehat{q}'$ is indifferent between disclosing or concealing p_1^H, and $V_2^C(\widehat{q}') = p_2' = A\widehat{q}'$, which contradicts the former argument.

From Case 1.1 and 1.2, I conclude $p_1^H \notin S_1^C$. Buyer 2's price is $\frac{A^2}{2A-1}\widehat{q_1}^H$ when p_1^H is disclosed and $p_1^H = \frac{\delta A^2}{2A-1}\widehat{q_1}^H + \alpha(1-\delta)\widehat{q_1}^H$.

Case 2: Suppose $p_2' \leq V_2^C\left(q_2^L\right)$. Buyer 2 belief of $\widehat{q_1}'$ cannot exceed q_2^L. His offer cannot exceed Aq_2^L. However, applying equation (B.5) and (B.2) together, I find that $p_2' > V_2^C\left(q_2^L\right)$. I find a contradiction.

Case 3: If $p_2' \geq V_2^C\left(q_2^H\right) = Aq_2^H$, any seller type in $\left(q_1^H, q_2^H\right)$ strictly prefers disclosing offer p_1^H, and any type $q \geq q_2^H$ weakly prefers disclosing offer p_1^H. Therefore $p_1^H \notin S_1^C$.

From the analysis in these three cases, I conclude that $p_1^H \notin S_1^C$. There is at most 1 elements in S_1^C. I will prove $p_1^H = \frac{\delta A^2}{2A-1}\widehat{q_1}^H + \alpha(1-\delta)\widehat{q_1}^H$.

If $d^2\left(q, \{p_1^H\}\right) = \{p_1^H\}$ for all q in equilibrium, buyer 2 believes that the cutoff seller type rejecting p_1^H is $\widehat{q_1}^H$. Any type higher than $\widehat{q_1}^H$ must also reject and disclose p_1^H. Therefore, buyer 2 will offer $\frac{A^2}{2A-1}\widehat{q_1}^H$. p_1^H satisfies the following indifference condition

$$p_1^H - \alpha\widehat{q_1}^H = \delta\left[\frac{A^2}{2A-1}\widehat{q_1}^H - \alpha\widehat{q_1}^H\right]$$

This suggests that $p_1^H = \frac{\delta A^2}{2A-1}\widehat{q_1}^H + \alpha(1-\delta)\widehat{q_1}^H$.

If $d^2\left(q, \{p_1^H\}\right) = \emptyset$ with probability $\beta(q) < 1$ for some q, buyer 2 must offer the same price whether p_1^H is disclosed or not. In this equilibrium, suppose that type q discloses the equilibrium offer p_1^H with probability $P(q, D)$. In period 2, when p_1^H is disclosed, buyer 2 forms a posterior belief of the seller's type with the density function $\phi_1(q) = \frac{\frac{1}{\bar{q}-\widehat{q_1}^H}P(q,D)}{\int_{\widehat{q_1}^H}^{\bar{q}} \frac{1}{\bar{q}-\widehat{q_1}^H}P(q,D)dq}$, and Φ_1 denotes the distribution function. When p_1^H is concealed, buyer 2 forms a posterior belief of the seller's type with the density function $\phi_2(q) = \frac{\frac{1}{\bar{q}-\widehat{q_1}^H}(1-P(q,D))}{\int_{\widehat{q_1}^H}^{\bar{q}} \frac{1}{\bar{q}-\widehat{q_1}^H}(1-P(q,D))dq}$, and Φ_2 denotes the distribution function. When p_1^H is disclosed, buyer 2's maximization problem depends on the cutoff type q_2 in period 2, which can be written as

$$\max \int_{\widehat{q_1}^H}^{q_2} (q - Aq_2)\,\phi_1(q)\,dq \tag{B.8}$$

The solution $\widehat{q_2}$ satisfies the condition that $(1-A)\,\widehat{q_2}\phi_1(\widehat{q_2}) = A\Phi_1(\widehat{q_2})$. In other words, it is $(1-A)\,\widehat{q_2}P(D,\widehat{q_2}) = A\int_{\widehat{q_1}^H}^{\widehat{q_2}} P(D,q)\,dq$. Buyer 2's price is $A\widehat{q_2}$. When p_1^H is concealed, buyer 2's maximization problem is $\max \int_{\widehat{q_1}^H}^{q_2} (q - Aq_2)\,\phi_2(q)\,dq$. Since $\widehat{q_2}$ is also the solution, I have $(1-A)\,\widehat{q_2}\phi_2(\widehat{q_2}) = A\Phi_2(\widehat{q_2})$. In addition, $(1-A)\,\widehat{q_2}\,(1-P(D,\widehat{q_2})) = A\int_{\widehat{q_1}^H}^{\widehat{q_2}} (1-P(D,q))\,dq$. Altogether, I have $\widehat{q_2} = \frac{A}{2A-1}\widehat{q_1}^H$. Therefore, $p_1^H = \frac{\delta A^2}{2A-1}\widehat{q_1}^H + \alpha(1-\delta)\,\widehat{q_1}^H$.

Let $p_1^S = \sup S_1$, and $\widehat{q_1}^S$ be the corresponding cutoff type. Using the arguments in Step 2, I get that $p_1^S \notin S_1^C$. Applying the same techniques in proving Step 2, I can show that $p_1^S = \frac{\delta A^2}{2A-1}\widehat{q_1}^S + \alpha(1-\delta)\,\widehat{q_1}^S$.

Step 3: Lastly, I will show that buyer 1 is not randomizing in equilibrium.

Case 1: I first consider the situation in which $S_1^C = \emptyset$. For any $p_1^M \in S_1$ and $p_1^M < p_1^S$, $p_1^M \in d^2\left(q, \{p_1^M\}\right)$ for some q. When p_1^M is disclosed, buyer 2 has a degenerate belief of the cutoff type; let p_2^M be his offer. Let $\widehat{q_1}^M$ be the cutoff type accepting offer p_1^M. Here,

$$p_1^M = \delta p_2^M + \alpha(1-\delta)\,\widehat{q_1}^M \tag{B.9}$$

Using the same techniques as in Step 2, I derive that buyer 2's offer p_2^M satisfies $p_2^M = \frac{A^2}{2A-1}\widehat{q_1}^M$. Since $\frac{\widehat{q_1}^M+q}{2} - p_1^M = \frac{\widehat{q_1}^S+q}{2} - p_1^S$, I get the following equation.

$$\frac{\widehat{q_1}^M}{2} - \alpha(1-\delta)\,\widehat{q_1}^M - \delta p_2^M = \left[\frac{1}{2} - \frac{\delta A^2}{2A-1} - \alpha(1-\delta)\right]\widehat{q_1}^S \tag{B.10}$$

The left hand side of (B.10) is actually $\left[\frac{1}{2}-\frac{\delta A^2}{2A-1}-\alpha\left(1-\delta\right)\right]\widehat{q_1}^M$. The right hand side of the equation is clearly smaller than the left hand side. Buyer 1 cannot randomize between p_1^M and p_1^H.

Case 2: If there exists $p_1^M \in S_1^C$, I have $p_1^M < p_1^S$. Since p_1^M is the only equilibrium offer not disclosed on path, buyer 2 has a degenerate belief of the cutoff type; let p_2^M be his offer. Let $\widehat{q_1}^M$ be the cutoff type accepting offer p_1^M. Here,

$$p_1^M = \delta p_2^M + \alpha\left(1-\delta\right)\widehat{q_1}^M \tag{B.11}$$

Using the technique in proving Case 1, I can show that buyer 1 cannot randomize between p_1^M and p_1^H.

Thus, I conclude that buyer 1 must play a pure strategy in equilibrium. In this case, buyer 2 has degenerate belief of the cutoff type, and he must also play a pure strategy. ■

Proof of Lemma 3. Lemma 2 allows me to only focus on the equilibrium where the two buyers play pure strategies. For a fixed p_1, let p_2^D denote buyer 2's price when p_1 is disclosed, and the corresponding cutoff type be $\widehat{q_2}^D$ ($p_2^D = A\widehat{q_2}^D$). Buyer 2's price when p_1 is concealed is denoted by p_2^C. Type-q seller's payoff in period 2 is $\max\left\{p_2, Aq\right\}$. If a type-$q$ seller prefers disclosing p_1, I must have:

$$\max\left\{p_2^D, Aq\right\} \geq \max\left\{p_2^C, Aq\right\} \tag{B.12}$$

(B.12) indicates that $p_2^D \geq p_2^C$. Therefore $\max\left\{p_2^D, Aq'\right\} \geq \max\left\{p_2^C, Aq'\right\}$ holds for all types rejecting p_1.

For the second part of Lemma 3, I have $\max\left\{p_2^D, Aq\right\} > \max\left\{p_2^C, Aq\right\}$. This suggests that $\max\left\{p_2^D, Aq'\right\} > \max\left\{p_2^C, Aq'\right\}$ for any $q' < q$; in addition, $p_2^D > p_2^C$ must hold. If p_1 is the equilibrium offer, suppose there exists a type q^M such that her disclosure rule in equilibrium satisfies $d^2\left(q^M, \{p_1\}\right) = \emptyset$ with positive probability. I must have $q^M \geq \widehat{q_2}^D$. When p_1 is not disclosed, buyer 2 believes that the seller's type must be greater than $\widehat{q_2}^D$. His offer p_2^C will be at least $A\widehat{q_2}^D = p_2^D$. However, this contradicts $p_2^D > p_2^C$. ■

Proof of Lemma 4. When the seller's strategy is $d^2\left(q, p^1\right) = \emptyset$ for any $p^1 \in P^1$ and q, any offer becomes unobservable to buyer 2. Following from Lemma 2, I only need to discuss the pure strategy equilibrium. In Proposition 10, I show that the BMU model does not have a pure strategy equilibrium when $\delta > 1 - \frac{1}{2\alpha}$. This intuitively shows that $d^2\left(q, p^1\right) = \emptyset$ for any $p^1 \in P^1$ and q cannot be sustained as the equilibrium disclosure rule when $\delta > 1 - \frac{1}{2\alpha}$. A general proof is presented below.

Consider an equilibrium in which all possible offers including the equilibrium offer p_1^* are concealed by all seller types. Let the corresponding cutoff accepting p_1^* be $\widehat{q_1}$. Buyer 2's offer must be $\frac{A^2}{2A-1}\widehat{q_1}$. However, for any $p_1 > p_1^*$, I will prove that buyer 2 would offer a price higher than $\frac{A^2}{2A-1}\widehat{q_1}$ if p_1 is known. Therefore, the seller would have the incentive to disclose it. Let me pick an arbitrary $p_1' > p_1^*$. Suppose that buyer 2's belief of the cutoff type is $\widehat{q_1}'$, and his offer is p_2' when p_1' is revealed. If $p_2' \leq \frac{A^2}{2A-1}\widehat{q_1}$, the price that buyer 2 offers when p_1' is concealed is higher than the price when p_1' is revealed. Therefore, any type prefers concealing p_1'. Suppose buyer 2's belief $\widehat{q_1}'$ of the cutoff type accepting p_1' is below $\widehat{q_1}$. Since

$$p_1' = \delta p_2' + \alpha(1-\delta)\widehat{q_1}' \tag{B.13}$$

I find $p_1' \leq \delta\frac{A^2}{2A-1}\widehat{q_1} + \alpha(1-\delta)\widehat{q_1}$, which contradicts $p_1' > p_1^*$. Therefore, $\widehat{q_1}'$ must be above $\widehat{q_1}$. Buyer 2 would offer a price greater than $\frac{A^2}{2A-1}\widehat{q_1}$, and p_1' would be disclosed by some seller types. Moreover, The proof also shows that buyer 2 would offer a higher price when a p_1 that is higher than the equilibrium price is disclosed. In conclusion, $d^2\left(q, p^1\right) = \emptyset$ for any $p^1 \in P^1$ and any q is not an equilibrium disclosure rule. ■

Proof of Proposition 6. When $d^2(q, p_1) = p^1$ for any $p^1 \in P^1$ and any q, any offer becomes observable to buyer 2. I first solve the equilibrium under this disclosure rule and check whether the players want to deviate.

In period 2, the cutoff type $\widehat{q_2}$ is indifferent between accepting the offer p_2 or rejecting and receiving the payoff $A\widehat{q_2}$. For buyer 2, $p_2 = A\widehat{q_2}$, and his payoff maximizing problem can be expressed as:

$$\max_{\widehat{q_2}} \frac{\widehat{q_2} - \widehat{q_1}}{\overline{q} - \widehat{q_1}} \left(\frac{\widehat{q_2} + \widehat{q_1}}{2} - p_2 \right)$$

The first order condition shows that $\widehat{q_2} = \frac{A}{2A-1}\widehat{q_1}$. At optimal, $\widehat{q_2}\left(\widehat{q_1}\right) = \frac{A}{2A-1}\widehat{q_1}$ is satisfied. From the consistency of beliefs, the cutoff type $\widehat{q_1}$ satisfies the condition below.

$$p_1 - \alpha\widehat{q_1} = \delta\left(p_2 - \alpha\widehat{q_1}\right) \Rightarrow p_1 = \delta p_2 + (1-\delta)\,\alpha\widehat{q_1} = \left(\delta\frac{A^2}{2A-1} + (1-\delta)\,\alpha\right)\widehat{q_1}$$

Buyer 1's payoff maximizing problem is:

$$\max_{\widehat{q_1}} \frac{\widehat{q_1} - \underline{q}}{\overline{q} - \underline{q}} \left(\frac{\widehat{q_1} + \underline{q}}{2} - p_1 \right)$$

Solving the first order condition, I get $\widehat{q_1} = \frac{\delta A^2 + (2A-1)(1-\delta)\alpha}{2\delta A^2 + (2A-1)(2(1-\delta)\alpha - 1)}\underline{q} = q_L$ and $\widehat{q_2} = \frac{A}{2A-1}\widehat{q_1}$.

I need to check whether the seller wants to deviate and conceal p_1. Given buyer 2's belief that non-disclosure suggests the seller being the lowest type, p_2 would be $\widehat{q_2} = \frac{A^2}{2A-1}\underline{q}$, conditional on non-disclosure. This price is strictly worse than the offer in equilibrium. Every seller type has no incentive to deviate. ■

Proof of Proposition 7. From Lemma 2, I only need to focus on the situation in which the two buyers play pure strategies.

Let the equilibrium price be p_t^* and the corresponding cutoff be $\widehat{q_t}$ in period t. Buyer 2's belief of the remaining type is above $\widehat{q_1}$, due to a pure strategy construction and the consistency of beliefs. When a price p_1' is disclosed, I denote buyer 2's belief of the cutoff in period 1 by $\widehat{q_1}'$ and his offer by p_2'. Moreover, I denote the cutoff type accepting p_2' by $\widehat{q_2}'$.

In Proposition 6, I discuss the equilibrium where any offer is disclosed by all seller types. This equilibrium is sustained if buyer 2 has the belief that the seller is of type $\underline{q}$ when no information is disclosed. There can be other equilibria in which buyer 2 has different beliefs of the seller's type when there is non-disclosure.

Case 1: I first consider the case in which $d^2(q,\{p_1^*\}) = \emptyset$ for some q. Using the argument in Lemma 3, I find that concealing p_1^* generates a higher payoff for every type that rejects p_1^*. Thus, I must have $d^2(q,\{p_1^*\}) = \emptyset$ if $p_1 = p_1^*$ for all q.[1] Buyer 2 maximizes his payoff by choosing $p_2^* = A\widehat{q_2}$ and $\widehat{q_2} = \frac{A}{2A-1}\widehat{q_1}$, where $\widehat{q_2}$ is the cutoff type accepting p_2^*. There is also the indifference condition, i.e., $p_1^* - \alpha\widehat{q_1} = \delta\left(A\frac{A}{2A-1}\widehat{q_1} - \alpha\widehat{q_1}\right)$. This shows that $p_1^* = \delta\frac{A^2}{2A-1}\widehat{q_1} + (1-\delta)\alpha\widehat{q_1}$.

If $p_1' > p_1^*$, the seller has the incentive to disclose p_1' when she rejects it. This is because buyer 2 will update his belief towards a higher cutoff type and provide $p_2' > p_2^*$, from the proof of Lemma 4. Thus, a price p_1' higher than the equilibrium offer will be disclosed eventually. In fact, all seller types have the incentive to do so, from Lemma 3. Buyer 2 maximizes his surplus when the cutoff type in period 2 is $\frac{A}{2A-1}\widehat{q_1}'$, and his price is $\frac{A^2}{2A-1}\widehat{q_1}'$. Given buyer 2's strategy, any seller type below $\frac{A}{2A-1}\widehat{q_1}'$ strictly prefers disclosing p_1'. All seller types weakly prefer disclosing offer p_1'. From the consistency of beliefs, $\widehat{q_1}'$ is pinned down by

$$p_1' - \alpha\widehat{q_1}' = \delta\left[\frac{A^2}{2A-1}\widehat{q_1}' - \alpha\widehat{q_1}'\right]$$

This shows that $\widehat{q_1}' = \frac{2A-1}{\delta A^2 + \alpha(1-\delta)(2A-1)}p_1'$. Moreover, if an offer $p_1' < p_1^*$ is disclosed, buyer 2 must set $p_2' < p_2^*$. Suppose not. If buyer 2 makes his price $p_2' \geq p_2^*$, his belief of the cutoff type in period 1 must be $\widehat{q_1}' \geq \widehat{q_1}$. This is inconsistent with $p_1' < p_1^*$. Therefore, any

[1] If in equilibrium some seller types disclose p_1^* while some seller types do not, buyer 2 must offer the same price whether p_1^* is disclosed or not. From the assumption made earlier, a seller type either discloses or conceals the equilibrium offer p_1^* with probability 1. Therefore, the support of buyer 2's belief when p_1^* is disclosed has no intersection with the support of buyer 2's belief when p_1^* is concealed. Suppose that buyer 2's belief of the seller's type is distributed according to $G_2(q|D)$ if p_1^* is disclosed, and $G_2(q|C)$ if p_1^* is concealed; in addition, $g_2(q|D)$ and $g_2(q|C)$ are the density functions, respectively. Let q_2' and q_2'' be the corresponding cutoffs in period 2 when p_1^* is disclosed or concealed. Buyer 2's optimal conditions are $(1-A)q_2'g_2(q_2'|D) = AG_2(q_2'|D)$ and $(1-A)q_2''g_2(q_2''|C) = AG_2(q_2''|C)$. Since $q_2' \neq q_2''$, buyer 2's price offer will be different whether p_1^* is disclosed or not. This is a contradiction, and I conclude that $d^2(q,\{p_1^*\}) = \emptyset$ for all q.

seller type $q < \widehat{q_2}$ would not disclose p_1' to buyer 2 because buyer 2 would respond with a lower price if he is informed. For any seller type $q \geq \widehat{q_2}$, her payoff would be Aq, and she would be indifferent between disclosing or concealing p_1'. However, her equilibrium strategy does not include disclosing p_1'. Otherwise, suppose there exists a seller of type q ($q \geq \widehat{q_2}$) such that her equilibrium strategy includes disclosing p_1'. Buyer 2 would believe that the seller's type is above $\widehat{q_2}$ if p_1' were disclosed, and he would set $p_2 \geq A\widehat{q_2}$, which is inconsistent with $p_2 < p_2^*$. Therefore, I conclude that any $p_1' < p_1^*$ is concealed by all seller types according to the seller's equilibrium strategy.

Given the seller's strategy and buyer 2's strategy, I must guarantee that buyer 1's payoff is maximized when making the equilibrium offer p_1^*. Any offer below p_1^* would be concealed while any offer above p_1^* would be disclosed. Buyer 1's payoff is a function of the cutoff type q_1 and is given below.

$$\pi_1(q_1) = \begin{cases} \frac{q_1 - \underline{q}}{\Delta q}\left(\frac{q_1+\underline{q}}{2} - (1-\delta)\alpha q_1 - \delta A \widehat{q_2}\right), & \text{when } p_1 \leq p_1^* \\ \frac{q_1 - \underline{q}}{\Delta q}\left(\frac{q_1+\underline{q}}{2} - (1-\delta)\alpha q_1 - \delta A \frac{A}{2A-1} q_1\right), & \text{when } p_1 > p_1^* \end{cases}$$

$$\frac{d\pi_1(q_1)}{dq_1} = \begin{cases} \frac{1}{\Delta q}\left(\frac{q_1+\underline{q}}{2} - (1-\delta)\alpha q_1 - \delta A \widehat{q_2}\right) + \frac{q_1 - \underline{q}}{\overline{q}-\underline{q}}\left[\frac{1}{2} - (1-\delta)\alpha\right], & \text{when } p_1 \leq p_1^* \\ \frac{1}{\Delta q}\left(\left(1 - 2(1-\delta)\alpha - \frac{2\delta A^2}{2A-1}\right) q_1 + \left(\frac{\delta A^2}{2A-1} + (1-\delta)\alpha\right)\underline{q}\right), & \text{when } p_1 > p_1^* \end{cases}$$

In this construction, I need $\frac{d\pi_1(q_1)}{dq_1} \geq 0$ for $p_1 \leq p_1^*$, and $\frac{d\pi_1(q_1)}{dq_1} \leq 0$ for $p_1 > p_1^*$ to guarantee that p_1^* is the best response. The second condition $\frac{d\pi_1(q_1)}{dq_1} \leq 0$ for $p_1 > p_1^*$ suggests $\widehat{q_1} \geq \frac{\delta A^2 + (2A-1)(1-\delta)\alpha}{2\delta A^2 + (2A-1)(2(1-\delta)\alpha - 1)}\underline{q} = q_L$. When $\frac{1}{2} - (1-\delta)\alpha \leq 0$, $\widehat{q_1} = \frac{(2A-1)(1-\delta)\alpha}{\delta A^2 + (2A-1)(2\alpha(1-\delta)-1)}\underline{q}$ maximizes buyer 1's payoff if $p_1 \leq p_1^*$. Therefore, the upper bound of $\widehat{q_1}$ is $\frac{(2A-1)(1-\delta)\alpha}{\delta A^2 + (2A-1)(2\alpha(1-\delta)-1)}\underline{q}$, and $\widehat{q_1} \in \left[q_L, \frac{(2A-1)(1-\delta)\alpha}{\delta A^2 + (2A-1)(2\alpha(1-\delta)-1)}\underline{q}\right]$.

Moreover, when $\frac{1}{2} - (1-\delta)\alpha > 0$, since buyer 1 has a non-negative payoff, $\frac{d\pi_1(q_1)}{dq_1} \geq 0$ for $p_1 \leq p_1^*$. The maximum price is reached when $\frac{q_1+\underline{q}}{2} = p_1^*$. Consequently, the upper bound of $\widehat{q_1}$ is $\frac{(2A-1)}{2\delta A^2 + (2A-1)(2\alpha(1-\delta)-1)}\underline{q}$, and $\widehat{q_1} \in \left[q_L, \frac{(2A-1)}{2\delta A^2 + (2A-1)(2\alpha(1-\delta)-1)}\underline{q}\right]$. Here, the equilibrium price p_1^* matches the value of the seller's disclosure threshold for p_1.

Case 2: The other case is $d^2(q, \{p_1^*\}) = \{p_1^*\}$ for some q. From Lemma 3, I have $d^2(q, \{p_1^*\}) = \{p_1^*\}$ for all q. Here, buyer 2 picks $p_2^* = A\widehat{q_2}$ where $\widehat{q_2} = \frac{A}{2A-1}\widehat{q_1}$, if p_1^* is disclosed to him. In this case, buyer 1's equilibrium offer satisfies $p_1^* = \delta\frac{A^2}{2A-1}\widehat{q_1} + (1-\delta)\alpha\widehat{q_1}$.

For any $p_1' > p_1^*$, the seller still has an incentive to disclose it when she rejects it, for the same reason as discussed in Case 1. When non-disclosure happens, buyer 2 believes that p_1' must be lower than p_1^*, and his belief of the cutoff type $\widehat{q_1}'$ must also be lower than the equilibrium cutoff type $\widehat{q_1}$.

If $\widehat{q_1}' = \widehat{q_1}$, buyer 2 maintains the same belief of the cutoff in period 1 when non-disclosure happens; in addition, he would offer the same price p_2^*. This is essentially the same as in Case 1 in which the equilibrium offer is not disclosed, and I will skip the detailed construction here. Note that in this situation, the seller would conceal offers made lower than the equilibrium offer. Here, when non-disclosure happens, buyer 2 could predict that the offer in period 1 were lower than the equilibrium offer p_1^*. It would not be optimal for him to offer a price equal to p_2^*.

If $\widehat{q_1}' < \widehat{q_1}$, I claim that $\widehat{q_1} = q_L$. Suppose there exists $\widehat{q_1} > q_L$. Essentially, every offer greater than $\delta\frac{A^2}{2A-1}\widehat{q_1}' + (1-\delta)\alpha\widehat{q_1}'$ would be disclosed. When $p_1 \geq \delta\frac{A^2}{2A-1}\widehat{q_1}' + (1-\delta)\alpha\widehat{q_1}'$,

$$\pi_1(q_1) = \frac{q_1 - \underline{q}}{\Delta q}\left(\frac{q_1 + \underline{q}}{2} - (1-\delta)\alpha q_1 - \delta A\frac{A}{2A-1}q_1\right)$$

Note that $\pi_1(q_1)$ is a quadratic function which is maximized at q_L. When $\widehat{q_1}' \geq q_L$, the seller has the incentive to deviate to offer $p_1 = \delta\frac{A^2}{2A-1}\widehat{q_1}' + (1-\delta)\alpha\widehat{q_1}'$. Here the cutoff in period 1 would be $\widehat{q_1}'$. As $\frac{d\pi_1(q_1)}{dq_1} \leq 0$ for $q_1 > q_L$, buyer 1 would achieve a higher payoff from deviation. When $\widehat{q_1}' < q_L$, buyer 1 has the incentive to deviate to $p_1 = \delta\frac{A^2}{2A-1}q_L + (1-\delta)\alpha q_L$. Therefore, the only supported cutoff is $\widehat{q_1} = q_L$.

In terms of the off-path belief $\widehat{q_1}'$ in the situation where p_1' is concealed, the only possible belief is $\widehat{q_1}' = \underline{q}$. As an offer will be concealed only when it is lower than p_1^*, the cutoff type $\widehat{q_1}'$ when there is non-disclosure is less than $\widehat{q_1}$. I have $\widehat{q_1}' = E(q_1|q < \widehat{q_1}) < \widehat{q_1}$. However, the seller has the incentive to reveal any offer that induces a cutoff type higher than $\widehat{q_1}'$. Only the offers inducing a cutoff type below $\widehat{q_1}'$ are concealed. Conditional on this, I still have $E\left(q_1|q<\widehat{q_1}'\right)<\widehat{q_1}'$. Thus, any offer inducing a cutoff type higher than $E\left(q_1|q<\widehat{q_1}'\right)$ will be revealed. I go over this reasoning process repeatedly, and eventually the only possible belief when there is non-disclosure is $\underline{q}$. This suggests that the seller essentially discloses every offer in this equilibrium, which is formally stated in Proposition 6.

Note: In the paper, I make the assumption about no randomization in disclosing the equilibrium offer for the purpose of equilibrium characterization. Removing this assumption does not salvage my result. Consider an equilibrium in which $d^2\left(q,p^1\right)=\emptyset$ with probability $\beta\left(q\right)<1$ if $p_1=p_1^*$ for some q. In this case, buyer 2 must offer the same price whether p_1^* is disclosed or not. In this equilibrium, let a type-q seller disclose p_1^* with probability $P\left(q,D\right)$. In period 2, when p_1^* is disclosed, buyer 2 forms a posterior belief of the seller's type with a density function $\phi_1\left(q\right)=\frac{\frac{1}{\bar{q}-\widehat{q_1}}P(q,D)}{\int_{\widehat{q_1}}^{\bar{q}}\frac{1}{\bar{q}-\widehat{q_1}}P(q,D)dq}$ and a distribution function ϕ_1. When p_1^* is concealed, buyer 2 forms a posterior belief of the seller's type with a density function $\phi_2\left(q\right)=\frac{\frac{1}{\bar{q}-\widehat{q_1}}(1-P(q,D))}{\int_{\widehat{q_1}}^{\bar{q}}\frac{1}{\bar{q}-\widehat{q_1}}(1-P(q,D))dq}$ and a distribution function Φ_2. When p_1^* is disclosed, buyer 2's maximization problem depends on the cutoff type q_2 in period 2, which can be written as $\max\int_{\widehat{q_1}}^{q_2}\left(q-Aq_2\right)\phi_1\left(q\right)dq$. The solution $\widehat{q_2}$ satisfies the condition that $\left(1-A\right)\widehat{q_2}\phi_1\left(\widehat{q_2}\right)=A\Phi_1\left(\widehat{q_2}\right)$. In other words, it is

$$\left(1-A\right)\widehat{q_2}P\left(D,\widehat{q_2}\right)=A\int_{\widehat{q_1}}^{\widehat{q_2}}P\left(D,q\right)dq \tag{B.14}$$

Buyer 2's price is $A\widehat{q_2}$. When p_1^* is concealed, buyer 2's maximization problem can be written as $\max\int_{\widehat{q_1}}^{q_2}\left(q-Aq_2\right)\phi_2\left(q\right)dq$. Since $\widehat{q_2}$ is also the solution, I have $\left(1-A\right)\widehat{q_2}\phi_2\left(\widehat{q_2}\right)=A\Phi_2\left(\widehat{q_2}\right)$. In other words,

$$\left(1-A\right)\widehat{q_2}\left(1-P\left(D,\widehat{q_2}\right)\right)=A\int_{\widehat{q_1}}^{\widehat{q_2}}\left(1-P\left(D,q\right)\right)dq \tag{B.15}$$

Combining equations (B.14) and (B.15) together, I derive $\widehat{q_2}=\frac{A}{2A-1}\widehat{q_1}$, which is the same as in Case 1. Using the same arguments as in Case 1, I find that the seller's equilibrium strategy includes disclosing any offer greater than p_1^* while concealing any offer less than p_1^*. In addition, any equilibrium, where $d^2(q,\{p_1^*\}) = \emptyset$ with probability $\beta(q) < 1$ for some q, can also be sustained as an equilibrium where $d^2(q,\{p_1^*\}) = \emptyset$ with probability 1 for all q, while other parts of the strategy profile remain unchanged. Thus the exact probabilities that the seller mixes between disclosing and concealing the equilibrium offer does not affect the equilibrium construction. In conclusion, for any equilibrium where some seller types randomize between disclosing and concealing p_1^*, there is an equilibrium generating the same equilibrium outcome in which all seller types conceal p_1^*. For simplicity and definiteness of the equilibrium characterization, I make the assumption that there is no randomization in disclosing the equilibrium offer. ■

Proof of Proposition 8. The expected welfare generated in trade is

$$\begin{aligned} W &= (1-\alpha)\left[\frac{\widehat{q}_1 - \underline{q}}{\Delta q}\cdot\frac{\widehat{q}_1+\underline{q}}{2} + \delta\frac{\widehat{q}_2-\widehat{q}_1}{\Delta q}\cdot\frac{\widehat{q}_2+\widehat{q}_1}{2}\right] \\ &= \frac{1-\alpha}{2}\left[\frac{\widehat{q}_1^{\,2}-\underline{q}^2}{\Delta q} + \delta\frac{\left(\frac{A}{2A-1}\right)^2\widehat{q}_1^{\,2}-\widehat{q}_1^{\,2}}{\Delta q}\right] \end{aligned}$$

W is an increasing function of $\widehat{q}_1$. Therefore, the maximum welfare level is reached at the upper bound of $\widehat{q}_1$, and the minimum welfare level is reached at the lower bound of $\widehat{q}_1$. I can easily calculated the buyers' surplus W_B and the seller's surplus W_S.

$$W_B = \frac{1}{2}\left[\frac{\widehat{q}_1^{\,2}-\underline{q}^2}{\Delta q} - p_1^*\frac{\widehat{q}_1-\underline{q}}{\Delta q}\right] + \delta\left[\frac{\left(\frac{A}{2A-1}\right)^2\widehat{q}_1^{\,2}-\widehat{q}_1^{\,2}}{\Delta q} - p_2^*\frac{\widehat{q}_2-\widehat{q}_1}{\Delta q}\right]$$

$$W_S = \frac{1}{2}\left[p_1^*\frac{\widehat{q}_1-\underline{q}}{\Delta q} - \alpha\frac{\widehat{q}_1^{\,2}-\underline{q}^2}{\Delta q}\right] + \delta\left[p_2^*\frac{\widehat{q}_2-\widehat{q}_1}{\Delta q} - \alpha\frac{\left(\frac{A}{2A-1}\right)^2\widehat{q}_1^{\,2}-\widehat{q}_1^{\,2}}{\Delta q}\right]$$

■

Proof of Proposition 9. See the proof of Proposition 6. ■

Proof of Proposition 10. [2] When the offers are unobservable, buyer 2's belief of the cutoff type

[2] The proof is very similar to the proof of Proposition 3 in [18]. Please refer to their proof for further information.

in period 1 is independent of the price history. In period 2, I have $p_2 = Aq_2$, where q_2 is the cutoff type accepting p_2. Let $K_2(\cdot)$ represent the cdf of the cutoff induced by the equilibrium price in period 2. Type q's expected payoff in period 2 is

$$V_2^C(q) = AqK_2(q) + \int_q^{\overline{q}} A\widetilde{q}dK_2(\widetilde{q})$$

Since the cutoff type q_1 satisfies the condition that $p_1 - \alpha q_1 = \delta\left(V_2^C(q_1) - \alpha q\right)$,

$$p_1 = \delta\left[\int_{q_1}^{\overline{q}} A\widetilde{q}dK_2(\widetilde{q}) + K_2(q_1)Aq_1\right] + (1-\delta)aq_1 \tag{B.16}$$

Let $K_1(\cdot)$ represent the cdf of the cutoff induced by the equilibrium price in period 1. Buyer 2's payoff is a function of the cutoff type q_2 and $K_1(\cdot)$. Specifically,

$$\pi_2(q_2;K_1) = \int_{\underline{q}}^{q_2}\int_{\underline{q}}^{c}(c - Aq_2)\frac{1}{1-\frac{\widetilde{q}-\underline{q}}{\overline{q}-\underline{q}}}dK_1(\widetilde{q})\,dc$$

The first order condition shows that

$$\begin{aligned}
(q_2 - Aq_2)\int_{\underline{q}}^{q_2}\frac{1}{\overline{q}-\widetilde{q}}dK_1(\widetilde{q}) &= A\int_{\underline{q}}^{q_2}\frac{q_2-\widetilde{q}}{\overline{q}-\widetilde{q}}dK_1(\widetilde{q}) = A\int_{\underline{q}}^{q_2}\left(1+\frac{q_2-\overline{q}}{\overline{q}-\widetilde{q}}\right)dK_1(\widetilde{q}) \Rightarrow \\
\int_{\underline{q}}^{q_2}\frac{1}{\overline{q}-\widetilde{q}}dK_1(\widetilde{q}) &= \frac{A}{(1-A)q_2 + A(\overline{q}-q_2)}
\end{aligned}$$

Moreover, buyer 1's payoff is a function of the cutoff type q_1. Here,

$$\pi_1(q_1;K_2) = \frac{q_1-\underline{q}}{\overline{q}-\underline{q}}\left(\frac{q_1+\underline{q}}{2} - p_1\right)$$

The first order condition is

$$\frac{d\pi_1(q_1;K_2)}{dq_1} = \frac{1}{\overline{q}-\underline{q}}\left(\frac{q_1+\underline{q}}{2} - p_1\right) + \frac{q_1-\underline{q}}{\overline{q}-\underline{q}}\left[\frac{1}{2} - \alpha(1-\delta) - \delta AK_2(q_1)\right] = 0$$

Considering a pure strategy equilibrium, I have $K_2(q_1) = 0$ and $p_1 = \delta p_2 + (1-\delta) a q_1$. If $\frac{1}{2} - \alpha(1-\delta) > 0$, the first order derivative satisfies $\frac{d\pi_1(q_1;K_2)}{dq_1} > 0$, as buyer 1's payoff is non-negative in equilibrium. Therefore, for any p_1^*, buyer 1 wants to deviate to a price higher than p_1^*. In this case, there is no pure strategy equilibrium. In addition, if $\frac{1}{2} - \alpha(1-\delta) \leq 0$, there exists a price that maximizes buyer 1's profit. After I solve the maximization problem, I derive the cutoffs that $q_1^{**} = \frac{(2A-1)(1-\delta)\alpha}{\delta A^2 + (2A-1)(2\alpha(1-\delta)-1)} \underline{q}$ and $q_2^{**} = \frac{A}{2A-1} q_1^{**}$.

In the next step, I will characterize the equilibrium when $\frac{1}{2} - \alpha(1-\delta) \geq 0$, i.e., $\delta > \delta^*$. I make the following claims.

Claim 1: When $\frac{1}{2} - \alpha(1-\delta) \geq 0$, if buyer 1 mixes continuously between prices that result in continuous cutoffs as an interval (a,b), buyer 1 must have a payoff of zero, and (a,b) is not in the support of $K_2(\cdot)$ in equilibrium.

If in equilibrium, buyer 1 mixes continuously between prices that result in cutoffs as an interval (a,b), then for any $q_1 \in (a,b)$, I have $\frac{q_1 + \underline{q}}{2} - p_1^{**} = \frac{B}{q_1 - \underline{q}}$, for some non-negative constant B and some equilibrium price p_1^{**}. The cutoff type accepting p_1^{**} satisfies equation (B.16). Taking derivatives of both sides, I have

$$\begin{aligned} \frac{1}{2} - \alpha(1-\delta) - A\delta K_2(q_1) &= -\frac{B}{\left(q_1 - \underline{q}\right)^2} \Rightarrow \\ \frac{1}{2} - \alpha(1-\delta) + \frac{B}{\left(q_1 - \underline{q}\right)^2} &= A\delta K_2(q_1) \end{aligned}$$

The left hand side of the equation is non-increasing in q_1. However, the right hand side is non-decreasing in q_1. Thus the only possible case is $B = 0$ and $\frac{1}{2} - \alpha(1-\delta) = A\delta K_2(q_1)$, which suggests that $K_2(q_1)$ is constant in interval (a,b). Therefore, $\frac{q_1 + \underline{q}}{2} = p_1^{**}$ holds. Also, (a,b) is not in the support of $K_2(\cdot)$, and $K_2(q_1) = \frac{\frac{1}{2} - \alpha(1-\delta)}{A\delta}$ for any $q_1 \in (a,b)$.

Claim 2: When $\frac{1}{2} - \alpha(1-\delta) \geq 0$, buyer 2 mixes between at most countably many prices.

Suppose that buyer 2 mixes continuously between prices that result in continuous cutoffs as an interval (a,b). For any $q_2 \in (a,b)$, I have $\frac{d\pi_2(q_2;K_1)}{dq_2} = 0$. This suggests $\int_{\underline{q}}^{q_2} \frac{1}{\overline{q}-\widetilde{q}} dK_1(\widetilde{q}) = \frac{A}{(1-A)q_2+A(\overline{q}-q_2)}$. However, the right hand side of the equation is increasing in q_2, which indicates that the left hand side of the equation is also increasing in q_2; thus, (a,b) is in the support of $K_1(\cdot)$. Nevertheless, this contradicts Claim 1. In conclusion, buyer 2 cannot mix between prices that result in cutoffs as an interval.

Suppose that q_1' is the smallest cutoff in the support of $K_1(\cdot)$. For a small enough ε,

$$\frac{d\pi_2(q_2;K_1)}{dq_2}|_{q_2=q_1'+\varepsilon} = \int_{q_1'}^{q_2} \frac{1}{\overline{q}-\widetilde{q}} dK_1(\widetilde{q}) \left((1-A)q_1' - (2A-1)\varepsilon\right) > 0$$

Buyer 2 does not choose any price that results in a cutoff equal to or below q_1'. From the claims established earlier, the support of $K_2(\cdot)$ is discrete, and $p_1(\cdot)$ is piece-wise linear and continuous. At any cutoff that is induced with positive probability in period 2, $\pi_1(q_1;K_2)$ has a kink. Similarly, at every cutoff that is induced with positive probability in period 1, $\pi_2(q_2;K_1)$ has a kink. Due to similar arguments found in the proof of Proposition 3 in [18], π_1 is a parabola open to the left, and $\underline{q}$ is in the support of $K_1(\cdot)$. Since $\pi_1\left(\underline{q};K_2\right) = 0$, I have $K_2(q) = 0$ for q close enough to $\underline{q}$.

Suppose that there exists an interval $\left[\underline{q}_2, \overline{q}_2\right]$ such that the induced cutoffs in period 1 are in this interval. From Claim 1, I know that $\left(\underline{q}_2, \overline{q}_2\right)$ is not in the support of $K_2(\cdot)$. Also, $\pi_2(q_2;K_1)$ has kinks at $\underline{q}_2$ and $\overline{q}_2$; thus, $\underline{q}_2$ and $\overline{q}_2$ must be in the support of $K_2(\cdot)$. Buyer 2's payoff π_2 is a piecewise quadratic parabola opening from below. Moreover, buyer 1 can only mix between $\{\underline{q}\} \cup \left[\underline{q}_2, \overline{q}_2\right]$. There is no other disjoint interval in the support of $K_1(\cdot)$ due to the quadratic property, as in the proof of Proposition 3 in [18]. In this case, buyer 2 randomizes between $\underline{q}_2$ and $\overline{q}_2$; in addition, $K_2\left(\underline{q}_2\right) = \frac{1-2\alpha(1-\delta)}{2\delta A}$ from previous analysis. Moreover, as $\pi_1\left(\underline{q}_2;K_2\right) = \pi_1(\overline{q}_2;K_2) = 0$, I can solve $\overline{q}_2 = \frac{1}{2\delta A+2\alpha(1-\delta)-1}\underline{q}$. As $\frac{d\pi_2(q_2;K_1)}{dq_2}|_{q_2=\underline{q}_2} = 0$, I solve $\underline{q}_2 = \frac{A}{2A-1}\underline{q}$. In addition, $\frac{d\pi_2(q_2;K_1)}{dq_2}|_{q_2=\overline{q}_2} = 0$. I get $\int_{\underline{q}}^{\overline{q}_2} \frac{1}{\overline{q}-\widetilde{q}} dK_1(\widetilde{q}) = \frac{A}{(1-A)\overline{q}_2+A(\overline{q}-\overline{q}_2)}$. As $\int_{\underline{q}}^{\overline{q}_2} \frac{1}{\overline{q}-\widetilde{q}} dK_1(\widetilde{q}) \leq K_1\left(\underline{q}\right)\frac{1}{\overline{q}-\underline{q}} + \left(1-K_1\left(\underline{q}\right)\right)\frac{1}{\overline{q}-\overline{q}_2}$, I further derive $\frac{K_1\left(\underline{q}\right)}{\Delta q} \leq \frac{(1-A)\overline{q}_2}{[(1-A)\overline{q}_2+A(\overline{q}-\overline{q}_2)]\left(\overline{q}_2-\underline{q}\right)}$.

Moreover, I also know that $\pi_2\left(\underline{q}_2;K_1\right)=\pi_2\left(\overline{q}_2;K_1\right)$. Here,

$$\begin{aligned}
\pi_2\left(\underline{q}_2;K_1\right) &= \int_{\underline{q}}^{\underline{q}_2}\int_{\underline{q}}^{c}\left(c-A\underline{q}_2\right)\frac{1}{\overline{q}-\widetilde{q}}dK_1\left(\widetilde{q}\right)dc \\
&= \frac{K_1\left(\underline{q}\right)}{2\Delta q}\left[\underline{q}_2^2-\underline{q}^2-2A\underline{q}_2\left(\underline{q}_2-\underline{q}\right)\right] \\
&= \frac{K_1\left(\underline{q}\right)}{2\left(2A-1\right)\Delta q}\left(1-A\right)^2\underline{q}^2
\end{aligned}$$

I can further simplify $\pi_2\left(\overline{q}_2;K_1\right)$ as

$$\begin{aligned}
\pi_2\left(\overline{q}_2;K_1\right) &= \int_{\underline{q}}^{\overline{q}_2}\int_{\underline{q}}^{c}\left(c-A\overline{q}_2\right)\frac{1}{\overline{q}-\widetilde{q}}dK_1\left(\widetilde{q}\right)dc \\
&= \int_{\underline{q}}^{\overline{q}_2}\left(\frac{1}{2}\overline{q}_2^2-\frac{1}{2}\widetilde{q}^2-A\overline{q}_2\left(\overline{q}_2-\widetilde{q}\right)\right)\frac{1}{\overline{q}-\widetilde{q}}dcdK_1\left(\widetilde{q}\right) \\
&= \frac{1}{2}\overline{q}+\frac{1}{2}\int_{\underline{q}}^{\overline{q}_2}\widetilde{q}dK_1\left(\widetilde{q}\right)-\frac{\overline{q}^2-\overline{q}_2^2}{2}\int_{\underline{q}}^{q_2}\frac{1}{\overline{q}-\widetilde{q}}dK_1\left(\widetilde{q}\right)-A\overline{q}_2 \\
&\quad +A\overline{q}_2\left(\overline{q}-\overline{q}_2\right)\int_{\underline{q}}^{q_2}\frac{1}{\overline{q}-\widetilde{q}}dK_1\left(\widetilde{q}\right)
\end{aligned}$$

This gives me another expression about the expected trading type in period 1.

$$\int_{\underline{q}}^{\overline{q}_2}\widetilde{q}dK_1\left(\widetilde{q}\right)=\frac{K_1\left(\underline{q}\right)}{\left(2A-1\right)\Delta q}\left(1-A\right)^2\underline{q}^2+2A\overline{q}_2-\overline{q}+A\left(\overline{q}-\overline{q}_2\right)\frac{\overline{q}-\left(2A-1\right)\overline{q}_2}{\left(1-A\right)\overline{q}_2+A\left(\overline{q}-\overline{q}_2\right)} \tag{B.17}$$

■

Proof of Proposition 11. I will prove (1) and (2) together first. In equation (B.17) in the proof of Proposition 10, I find the expected trading type in period 1 in the BMU model, which is

$$\int_{\underline{q}}^{\overline{q}_2}\widetilde{q}dK_1\left(\widetilde{q}\right)=\frac{K_1\left(\underline{q}\right)}{\left(2A-1\right)\Delta q}\left(1-A\right)^2\underline{q}^2+2A\overline{q}_2-\overline{q}+A\left(\overline{q}-\overline{q}_2\right)\frac{\overline{q}-\left(2A-1\right)\overline{q}_2}{\left(1-A\right)\overline{q}_2+A\left(\overline{q}-\overline{q}_2\right)}$$

Here I have $\frac{K_1\left(\underline{q}\right)}{\Delta q}\leq\frac{\left(1-A\right)\overline{q}_2}{\left[\left(1-A\right)\overline{q}_2+A\left(\overline{q}-\overline{q}_2\right)\right]\left(\overline{q}_2-\underline{q}\right)}$. Furthermore, I compare the expected trading type in the BMU model with the trading type in the welfare-maximizing equilibrium in my model. In period 1, the comparison is given below, where $\widehat{q_1} = \frac{(2A-1)}{2\delta A^2+(2A-1)(2\alpha(1-\delta)-1)}\underline{q}$. For simplicity, I will refer to the welfare-maximizing equilibrium of my model as the WM equilibrium from now on.

$$\begin{aligned}
\int_{\underline{q}}^{\overline{q}_2} \widetilde{q} dK_1(\widetilde{q}) - \widehat{q_1} &= \frac{K_1(\underline{q})}{(2A-1)\Delta q}(1-A)^2 \underline{q}^2 + 2A\overline{q}_2 - \overline{q} + A(\overline{q}-\overline{q}_2)\frac{\overline{q}-(2A-1)\overline{q}_2}{(1-A)\overline{q}_2 + A(\overline{q}-\overline{q}_2)} \\
&\quad - \frac{(2A-1)}{2\delta A^2 + (2A-1)(2\alpha(1-\delta)-1)}\underline{q} \\
&= (2A-1)(1-A)\overline{q}_2\left(\frac{(\overline{q}-\overline{q}_2)}{(1-A)\overline{q}_2 + A(\overline{q}-\overline{q}_2)} - \frac{2(\delta A + 2\alpha(1-\delta)-1)}{2\delta A^2 + (2A-1)(2\alpha(1-\delta)-1)}\right) \\
&\quad + \frac{K_1(\underline{q})}{(2A-1)\Delta q}(1-A)^2 \underline{q}^2 \\
&\leq (2A-1)(1-A)\overline{q}_2\left(\frac{(\overline{q}-\overline{q}_2)}{(1-A)\overline{q}_2 + A(\overline{q}-\overline{q}_2)} - \frac{2(\delta A + 2\alpha(1-\delta)-1)}{2\delta A^2 + (2A-1)(2\alpha(1-\delta)-1)}\right) \\
&\quad + \frac{(1-A)^3\overline{q}_2}{[(1-A)\overline{q}_2 + A(\overline{q}-\overline{q}_2)](2A-1)(\overline{q}_2-\underline{q})}\underline{q}^2 \qquad \text{(B.18)}
\end{aligned}$$

The sign of the right hand side of inequality (B.18) depends on the parameters. Under some assumptions, I can show that $\int_{\underline{q}}^{\overline{q}_2} \widetilde{q} dK_1(\widetilde{q}) < \widehat{q_1}$. One possible assumption needs $\overline{q}$ to be relatively small. I suppose $\frac{(\overline{q}-\overline{q}_2)}{(1-A)\overline{q}_2+A(\overline{q}-\overline{q}_2)} \leq \frac{3}{2}\frac{\delta A+2\alpha(1-\delta)-1}{2\delta A^2+(2A-1)(2\alpha(1-\delta)-1)}$ for now. The above inequality can be further simplified to

$$\int_{\underline{q}}^{\overline{q}_2} \widetilde{q} dK_1(\widetilde{q}) - \widehat{q_1} \leq \frac{1-A}{2A-1}\left[\frac{(1-A)}{(\overline{q}_2-\underline{q})}\underline{q}^2 - \frac{(2A-1)^2}{2}\overline{q}_2\frac{\delta A + 2\alpha(1-\delta)-1}{2\delta A^2 + (2A-1)(2\alpha(1-\delta)-1)}\right] \qquad \text{(B.19)}$$

The right hand side of the inequality eventually becomes strictly negative as A gradually approaches 1. For any set of parameters α and δ, there exists a $A^{**}(\delta,\alpha)$ such that for $A > A^{**}$, $\int_{\underline{q}}^{\overline{q}_2} \widetilde{q} dK_1(\widetilde{q}) < \widehat{q_1}$. Let me denote $M = \frac{(\overline{q}-\overline{q}_2)}{(1-A)\overline{q}_2+A(\overline{q}-\overline{q}_2)} - \frac{3}{2}\frac{\delta A+2\alpha(1-\delta)-1}{2\delta A^2+(2A-1)(2\alpha(1-\delta)-1)}$. I have $\frac{dM}{d\overline{q}} \geq 0$. For any set of parameters δ, α, A, and $\underline{q}$, there exists a $\overline{q}^*(\delta,\alpha,A,\underline{q})$ such that when $\overline{q} \leq \overline{q}^*$, the condition that $\frac{(\overline{q}-\overline{q}_2)}{(1-A)\overline{q}_2+A(\overline{q}-\overline{q}_2)} \leq \frac{3}{2}\frac{\delta A+2\alpha(1-\delta)-1}{2\delta A^2+(2A-1)(2\alpha(1-\delta)-1)}$ is guaranteed. Another possible assumption is $\alpha(1-\delta)$ being sufficiently large. As $\alpha(1-\delta) \to \frac{1}{2}$, I have $\int_{\underline{q}}^{\overline{q}_2} \widetilde{q} dK_1(\widetilde{q}) - \widehat{q_1} \to \frac{(1-A)^2}{(2A-1)^2}\left[\frac{K_1(\underline{q})}{\Delta q}\underline{q}^2 - (2A-1)^2\overline{q}_2^2\frac{1}{A(A\overline{q}-(2A-1)\overline{q}_2)}\right]$. Since $(2A-1)\overline{q}_2 \geq \underline{q}$, the sign of $\int_{\underline{q}}^{\overline{q}_2} \widetilde{q} dK_1(\widetilde{q}) - \widehat{q}_1$ can be determined. Here,

$$\frac{K_1(\underline{q})}{\Delta q}\underline{q}^2 - (2A-1)^2 \overline{q}_2^2 \frac{1}{A(A\overline{q} - (2A-1)\overline{q}_2)} \leq \underline{q}^2 \left(\frac{K_1(\underline{q})}{\Delta q} - \frac{1}{A(A\overline{q} - \underline{q})} \right) < 0$$

In this case, I also have $\int_{\underline{q}}^{\overline{q}_2} \widetilde{q} dK_1(\widetilde{q}) < \widehat{q}_1$.

Moreover, in the proof of Proposition 10, I show that buyer 1 in the BMU model makes a profit of zero. This suggest that buyer 1's price is equal to the average value, i.e., $p_1^{**} = \frac{\underline{q}+q_1^*}{2}$, for any equilibrium cutoff $q_1^* \in \left[\underline{q}_2, \overline{q}_2\right]$. When the cutoff type is the lowest type, buyer 1's offer will not exceed $\underline{q}$. This suggests that buyer 1's expected payoff satisfies

$$\int_{\underline{q}}^{\overline{q}_2} p_1(\widetilde{q}) dK_1(\widetilde{q}) \leq \int_{\underline{q}}^{\overline{q}_2} \frac{\widetilde{q} + \underline{q}}{2} = \frac{\underline{q}}{2} + \frac{1}{2}\int_{\underline{q}}^{\overline{q}_2} \widetilde{q} dK_1(\widetilde{q})$$

In the WM equilibrium of the optional disclosure model, buyer 1 also makes a payoff of zero, i.e., $p_1^* = \frac{\underline{q}+\widehat{q}_1}{2}$. Thus I derive the inequality below.

$$\int_{\underline{q}}^{\overline{q}_2} p_1(\widetilde{q}) dK_1(\widetilde{q}) \leq \frac{\underline{q}}{2} + \frac{1}{2}\int_{\underline{q}}^{\overline{q}_2} \widetilde{q} dK_1(\widetilde{q}) \leq \frac{\underline{q}}{2} + \frac{\widehat{q}_1}{2}$$

The trading price in period 1 in the WM equilibrium is strictly higher than the expected price in period 1 in any equilibrium of the BMU model. Next, I compare the expected trading type in period 2. In the BMU model, buyer 2 mixes between the prices that induce $\underline{q}_2$ and $\overline{q}_2$ with probabilities $K_2\left(\underline{q}_2\right) = \frac{1-2\alpha(1-\delta)}{2\delta A}$ and $K_2(\overline{q}_2) = 1 - K_2\left(\underline{q}_2\right)$, where $\underline{q}_2 = \frac{A}{2A-1}\underline{q}$ and $\overline{q}_2 = \frac{1}{2\delta A + 2\alpha(1-\delta) - 1}\underline{q}$.

$$\begin{aligned} \int_{\underline{q}_2}^{\overline{q}_2} \widetilde{q} dK_2(\widetilde{q}) - \widehat{q}_2 &= K_2\left(\underline{q}_2\right)\underline{q}_2 + \left(1 - K_2\left(\underline{q}_2\right)\right)\overline{q}_2 - \frac{A}{2A-1}\widehat{q}_1 \\ &= \frac{(1-2\alpha(1-\delta))\left[2\delta A(1-A)^2 + (2A-1)(1-A-2A(1-\delta)(1-\alpha))\right]}{2\delta A(2A-1)(2\delta A^2 + (2A-1)(2\alpha(1-\delta)-1))} \end{aligned}$$

The above expression becomes negative as A goes to 1. There exists a threshold $A^*(\delta, \alpha)$ such that, for $A > A^*$, $\int_{\underline{q}_2}^{\overline{q}_2} \widetilde{q} dK_2(\widetilde{q}) - \widehat{q}_2 < 0$. Actually, when $A > \delta + \alpha(1-\delta)$, $\int_{\underline{q}_2}^{\overline{q}_2} \widetilde{q} dK_2(\widetilde{q}) - \widehat{q}_2 < 0$

holds naturally for any $\alpha > \frac{1}{2}$ and $\delta > \delta^*$.

In terms of the equilibrium price in period 2, I know that $p_2^{**} = Aq_2^*$. This indicates that $\int_{\underline{q}_2}^{\overline{q}_2} p_2(\widetilde{q})\, dK_2(\widetilde{q}) < p_2(\widehat{q}_2)$. In summary, I establish the result that the trading prices in the WM equilibrium of my model are strictly higher than those in any equilibrium of the BMU model. This indicates that all seller types are exante better off with the disclosure option. Moreover, buyer 1 ends up with a payoff of zero in the WM equilibrium of my model, as well as in any equilibrium of the BMU model. For buyer 2, the surplus is $\frac{\widehat{q}_2-\widehat{q}_1}{\overline{q}-\widehat{q}_1}\left(\frac{\widehat{q}_2+\widehat{q}_1}{2} - A\widehat{q}_2\right) = \frac{(1-A)^2\widehat{q}_1{}^2}{2(2A-1)(\overline{q}-\widehat{q}_1)}$ in the WM equilibrium, while it is $\frac{(1-A)^2\underline{q}^2}{2(2A-1)\left(\overline{q}-\underline{q}\right)}$ in any equilibrium of the BMU model. Buyer 2 is better off in the WM equilibrium of my model as well. Therefore, the WM equilibrium ex-ante Pareto dominates any equilibrium of the BMU model.

The third statement is straightforward to see. Since the welfare-minimizing equilibrium of the optional disclosure model has exactly the same two-period cutoffs as the unique equilibrium of the BMO model. Any equilibrium of the optional disclosure model yields a weakly higher social surplus than the equilibrium of the BMO model. ■

B.2 Proofs and Supplementary Materials for Chapter 2.6.1

In this subsection, I provide supplementary materials for Section 2.6.1. Proposition 19 describes the situation in which there is a small variation in the item's quality. In this case, all buyers will simply offer $\alpha\overline{q}$, which is accepted by all seller types.

Proposition 19 *When $\delta > \delta^*$ and $\underline{q} \geq (2\alpha - 1)\overline{q}$, every buyer submits the offer $\alpha\overline{q}$, and all seller types accept the offer $\alpha\overline{q}$ in equilibrium.*

Proof of Proposition 19. Suppose that buyer 1 offers $p_1' \leq \alpha\overline{q}$, which serves types up to q_1'. Here, $p_1' \geq \alpha q_1'$. Buyer 1's payoff satisfies

$$\pi_1 = (q_1' - \underline{q})\left(\frac{q_1' + \underline{q}}{2} - p_1'\right) \leq (q_1' - \underline{q})\left(\frac{q_1' + \underline{q}}{2} - \alpha q_1'\right)$$

Here, $\left(q_1' - \underline{q}\right)\left(\frac{q_1'+\underline{q}}{2} - \alpha q_1'\right)$ is the upper bound of π_1. The derivative of $\left(q_1' - \underline{q}\right)\left(\frac{q_1'+\underline{q}}{2} - \alpha q_1'\right)$ with respect to q_1' is $(1-2\alpha)\, q_1' + \alpha \underline{q}$, which is positive when $\underline{q}$ is large enough, i.e., $\underline{q} \geq \frac{2\alpha-1}{\alpha}\overline{q}$. The upper bound reaches the maximum when $q_1' = \overline{q}$. When $p_1' = \alpha\overline{q}$, all seller types accept the offer p_1', and $q_1' = \overline{q}$. In this case, the maximum of π_1 is also reached. Every buyer will offer $\alpha\overline{q}$ in equilibrium.

Moreover, if $\underline{q} < \frac{2\alpha-1}{\alpha}\overline{q}$, suppose that $q_1' > \frac{2\alpha-1}{\alpha}\overline{q}$ for now. Then buyer 2 will make the offer $\alpha\overline{q}$. Since $p_1' - \alpha q_1' = \delta\left(\alpha\overline{q} - q_1'\right)$, buyer 1's payoff is

$$\pi_1 = \left(q_1' - \underline{q}\right)\left(\frac{q_1'+\underline{q}}{2} - \delta\alpha\overline{q} - (1-\delta)\, q_1'\right)$$

This is increasing in q_1', and buyer 1 still offers the price $\alpha\overline{q}$. The other possible case is $q_1' \leq \frac{2\alpha-1}{\alpha}\overline{q}$. The derivative of $\left(q_1' - \underline{q}\right)\left(\frac{q_1'+\underline{q}}{2} - \alpha q_1'\right)$ with respect to q_1' is still positive, as

$$(1-2\alpha)\, q_1' + \alpha\underline{q} \geq (2\alpha-1)\left(\alpha\overline{q} - q_1'\right) \geq 0$$

The upper bound reaches the maximum when $q_1' = \frac{2\alpha-1}{\alpha}\overline{q}$, and buyer 1's payoff satisfies

$$\pi_1 \leq \left(\frac{2\alpha-1}{\alpha}\overline{q} - \underline{q}\right)\left(\frac{\frac{2\alpha-1}{\alpha}\overline{q}+\underline{q}}{2} - \alpha\frac{2\alpha-1}{\alpha}\overline{q}\right)$$

However,

$$\left(\frac{2\alpha-1}{\alpha}\overline{q} - \underline{q}\right)\left(\frac{\frac{2\alpha-1}{\alpha}\overline{q}+\underline{q}}{2} - \alpha\frac{2\alpha-1}{\alpha}\overline{q}\right) < \left(\overline{q} - \underline{q}\right)\left(\frac{\overline{q}+\underline{q}}{2} - \alpha\overline{q}\right) \tag{B.20}$$

When $q_1' \leq \frac{2\alpha-1}{\alpha}\overline{q}$, inequality (B.20) shows that choosing $p_1 = \alpha\overline{q}$ yields a strictly higher payoff for buyer 1. Buyer 1 is better off by picking $\alpha\overline{q}$. A similar proof applies to all future buyers; thus, all of them will offer $\alpha\overline{q}$. If the variation of the product quality is small, every buyer will serve all seller types by offering $\alpha\overline{q}$. Here, having the disclosure option will not affect the equilibrium outcome. ■

Lemma 8 *In any equilibrium, when $\delta > \delta^*$ and $\underline{q} < (2\alpha - 1)\overline{q}$, a buyer plays a pure strategy when the cutoff type accepting his offer is strictly lower than $\overline{q}$. The buyer plays a mixed strategy when he serves the type-$\overline{q}$ seller with positive probability. In this case, the buyer randomizes between two offers, one of which is $\alpha\overline{q}$.*

Proof of Lemma 8. In the infinite-horizon game, similar to the two-period model, the seller's continuation value is non-decreasing in her type. The proof of the first part is very similar to the proof of Lemma 2, and I will skip it here. The only possible situation that a buyer randomizes is when he trades with type $\overline{q}$ with positive probability. In this case, the buyer randomizes between only two offers; the higher one is $\alpha\overline{q}$. Otherwise, there are at least two offers resulting in different cutoff types strictly below $\overline{q}$, which would create the disclosure incentive for the seller. Similar to the proof of Lemma 2, the higher offer of the two would be disclosed; thus, this buyer would have no incentive to randomize. ■

Lemma 8 states that a buyer's equilibrium strategy does not include randomizing unless the highest type accepts his offer with positive probability.

Proof of Lemma 5. This is essentially proving $\widehat{p_{t,t+1}}\left(q, p^{t-1}\right) = \widehat{p_{t,t+k}}\left(q, p^{t+k-1}\right)$ for any $k > 1$. I will prove it by contradiction. Consider an arbitrary ST equilibrium. In that ST equilibrium, if the seller rejects p_t, she discloses the offer p_t to buyer $t+1$ when $p_t > \widehat{p_{t,t+1}}\left(q, p^t\right)$, and she discloses p_t to buyer $t+2$ when $p_t > \widehat{p_{t,t+2}}\left(q, p^{t+1}\right)$.

Suppose $\widehat{p_{t,t+1}}\left(q, p^t\right) < \widehat{p_{t,t+2}}\left(q, p^{t+1}\right)$ holds in the equilibrium of some subgame. The equilibrium offer p_t^* in the subgame must match the value of $\widehat{p_{t,t+1}}\left(q, p^t\right)$ on the equilibrium path. I consider a situation in which buyer t raises his price slightly higher to $p_t^* + \varepsilon$, and $p_t^* + \varepsilon < \widehat{p_{t,t+2}}\left(q, p^{t+1}\right)$ holds. In this case, the offer $p_t^* + \varepsilon$ will be disclosed to buyer $t+1$, but not to buyer $t+2$. Note that the equilibrium offer p_{t+1}^* in the subgame must match the value of $\widehat{p_{t+1,t+2}}\left(q, p^{t+1}\right)$ on the equilibrium path. Here, buyer $t+1$ would not want to raise his price. For any belief of the cutoff type accepting $p_t^* + \varepsilon$ that buyer $t+1$ has, if buyer $t+1$ raised his price, he would be worse off as buyer $t+2$ would perceive this situation as an deviation from buyer $t+1$ rather than from buyer t. Buyer $t+2$ would pick a more aggressive price. Here, buyer $t+2$ would not observe the deviation from buyer t. He would only observe buyer $t+1$'s deviation and act aggressively. From the equilibrium definition, buyer $t+1$ must be worse off if he deviates to a higher price. Thus, he would still offer the original price p^*_{t+1}. Nevertheless, buyer t would achieve a higher payoff by deviating to $p^*_t+\varepsilon$, which cannot happen in equilibrium. Therefore, $\widehat{p_{t,t+1}} \geq \widehat{p_{t,t+2}}$ must hold in the equilibrium of every subgame. If $\widehat{p_{t,t+1}} > \widehat{p_{t,t+2}}$ holds in equilibrium , this suggests that the equilibrium offer p^*_t matches the value of $\widehat{p_{t,t+1}}$ on the equilibrium path. The offer p^*_t will be disclosed to buyer $t+2$, which violates the definition of the ST equilibrium. Moreover, if for some price history, $\widehat{p_{t,t+1}} > \widehat{p_{t,t+2}}$ holds. There is an offer $\widehat{p_{t,t+1}}-\varepsilon$ that will be disclosed to buyer $t+2$ but not to buyer $t+1$. This offer must be lower than the equilibrium offer. Otherwise, it would be disclosed to buyer $t+1$. However, it will not not rational for the seller to disclose a price that is lower than the equilibrium price to future buyers as shown in the two-period model. Hence, I prove that $\widehat{p_{t,t+1}} = \widehat{p_{t,t+2}}$. By induction, I can show that $\widehat{p_{t,t+1}} = \widehat{p_{t,t+k}}$ for $k>1$. ■

Proof of Theorem 2. I claim that for any other equilibrium $(d^t, p^*_t, \widehat{q_t})$, if $p^*_t \in d^{t+k}\left(q, p^{t+k-1}\right)$, $p_t \in d^{t+k}\left(q, p^{t+k-1}\right)$ for any $p_t > p^*_t$. Similar to the proof of Proposition 7, here, concealing an offer will induce a lower belief of the cutoff type than the corresponding equilibrium belief. If an offer higher than the equilibrium price is revealed, it will lead the next buyer to form a higher belief of the cutoff type than the corresponding equilibrium belief. This can be shown by contradiction. Suppose buyer $t+1$ would form a belief of the cutoff type in period t lower than the corresponding equilibrium belief, when some offer p_t higher than p^*_t were revealed. Here, I can construct a deviation of buyer t. If buyer t were to make the offer p_t, regardless of the seller's disclosure rule, it would lead buyer $t+1$ to form a belief of the cutoff type lower than the equilibrium belief. Given a lower belief, buyer $t+1$ must offer a price p_{t+1} weakly lower than his equilibrium price p^*_{t+1}. Let $\Delta\widehat{q_t}$ denote the change in the cutoff type $\widehat{q_t}$, and $\varepsilon = p_t - p^*_t$.

$$p_t = \delta p_{t+1} + \alpha(1-\delta)(\widehat{q_t} + \Delta\widehat{q_t})$$

In this case, buyer t would receive a higher payoff than his equilibrium payoff, as the higher seller types would accept p_t. The gain is more than $\frac{\varepsilon}{2\alpha(1-\delta)}$ while the cost is ε. Given $\delta > \delta^*$, buyer t is better off when he deviates to p_t. Here, I show a deviation. Thus the seller's strategy must include revealing p_t to buyer $t+1$ when $p_t > p^*_t$. Using similar arguments established in proving Lemma 5, I find that p_t will be revealed to all future buyers given $p_t > p^*_t$.

From now on, I will construct the ST equilibrium. I modify the seller's strategy as $p_m \in d^t\left(q, p^{t-1}\right)$ if $p_m > \widehat{p_m}\left(q, p^{t-1}\right)$ for $m \leq t-1$, and $p_m \notin d^t\left(q, p^{t-1}\right)$ otherwise. The disclosure thresholds satisfy the conditions that $\widehat{p_1}\left(q, p^0\right) = p^*_1$, and $\widehat{p_t}\left(q, p^{t-1}\right) = p^*_t$ if $p_m \leq \widehat{p_m}\left(q, p^{m-1}\right)$ for all $m \leq t-1$. If there exists $m \leq t-1$ such that $p_m > \widehat{p_m}$, the seller reports p_m to buyer t; in addition, the disclosure threshold for the offer in period t will be raised to $\widehat{p_t}\left(q, p^{t-1}\right) = p^*_t + \varepsilon$ where $\varepsilon = p_m - \widehat{p_m}$.

In addition to this modification, I let the buyers' strategies and the seller's acceptance rule remain the same as in the original equilibrium. Buyers' beliefs are derived from the Bayes' rule. Under this modification, every buyer still maintains the same equilibrium belief of the cutoff type, since the seller's acceptance rule is unchanged. In equilibrium, the value of the seller's disclosure threshold $\widehat{p_t}$ matches the equilibrium price p^*_t, and no equilibrium offer is disclosed on the equilibrium path.

I can show that it is still an equilibrium. Conditional on the seller's strategy, buyer t has no incentive to reduce his price. Moreover, he also has no incentive to increase his price. If buyer t increases his price to $p^*_t + \varepsilon'$, this will be reported to buyers in the future, which is the same as in the original equilibrium. If $p^*_t + \varepsilon'$ gives buyer t a higher surplus in the modified equilibrium, it should also give buyer t a higher surplus in the original equilibrium. Therefore, buyer t must have no incentive to deviate. Moreover, as buyers' offers are unchanged, the seller must have no incentive to deviate from the current acceptance rule. In conclusion, it is indeed an ST equilibrium, and the equilibrium prices and cutoff types remain unchanged. ■

Proof of Proposition 12. Consider an ST equilibrium with a strategy profile $(d^t, p_t, \widehat{q}_t)$:

(1) Suppose that there exists a period s and some n ($n > 1$) such that the seller picks a disclosure threshold $\alpha(2\alpha-1)^n\overline{q}$, and will reveal any offer above the threshold in all future periods. This indicates that every offer above $\alpha(2\alpha-1)^n\overline{q}$ will be observable to all future buyers after period s. The equilibrium cutoff type in period s is denoted by $\widehat{q}_s$. For now, I suppose $\widehat{q}_s \geq (2\alpha-1)^{n+1}\overline{q}$. I first discuss the case where $\widehat{q}_s > (2\alpha-1)^n\overline{q}$. Any offer above $\alpha\widehat{q}_s$ will be observable from period s onwards. I make the following statement.

Claim: If $\widehat{q}_s \in ((2\alpha-1)^k\overline{q}, (2\alpha-1)^{k-1}\overline{q}]$, buyer $s+1$ offers $\alpha(2\alpha-1)^{k-1}\overline{q}$ in equilibrium. The seller type up to $(2\alpha-1)^{k-1}\overline{q}$ accepts the offer. Every future buyer will make an offer no larger than $\alpha(2\alpha-1)^{k-1}\overline{q}$.

I will prove this claim by induction.

Step 1: Suppose $\widehat{q}_s \in ((2\alpha-1)\overline{q}, \overline{q}]$. Buyer $s+1$ offers $\alpha\overline{q}$ in equilibrium. All seller types accept this offer. Every future buyer will make an offer no larger than $\alpha\overline{q}$.

Similar to the proof of Proposition 19, I can show that offering $\alpha\overline{q}$ maximizes buyer $s+1$'s surplus. Every future buyer will make an offer equal to $\alpha\overline{q}$. The first step is complete.

Suppose that the statement is true for $k-1$. Now I will prove it for k.

Step 2: Suppose $\widehat{q}_s \in ((2\alpha-1)^k\overline{q}, (2\alpha-1)^{k-1}\overline{q}]$, and buyer $s+1$ offers $\alpha(2\alpha-1)^{k-1}\overline{q}$ in equilibrium. The seller type up to $(2\alpha-1)^{k-1}\overline{q}$ accepts the offer. Every future buyer will make an offer no larger than $\alpha(2\alpha-1)^{k-1}\overline{q}$.

This argument is proved by contradiction. Suppose this is not true, and there exists a type-q seller ($q > (2\alpha-1)^{k-1}\overline{q}$) accepting buyer $s+1$'s offer p_{s+1}. I also have $q \leq (2\alpha-1)^{k-2}\overline{q}$, as buyer $s+1$ must make a non-negative payoff. The price offer will be $\alpha(2\alpha-1)^{k-2}\overline{q}$ in the next period. This suggests that the price p_{s+1} satisfies the following condition.

$$p_{s+1} - \alpha q \geq \delta \left(\alpha (2\alpha - 1)^{k-2} \overline{q} - \alpha q \right) \Rightarrow p_{s+1} \geq \delta \alpha (2\alpha - 1)^{k-2} \overline{q} + (1 - \delta) \alpha q$$

However,

$$\begin{aligned} \delta\alpha(2\alpha-1)^{k-2}\overline{q}+(1-\delta)\alpha q-\frac{\widehat{q_s}+q}{2} &\geq \delta\alpha(2\alpha-1)^{k-2}\overline{q}+(1-\delta)\alpha q-\frac{(2\alpha-1)^{k-1}\overline{q}+q}{2} \\ &= \frac{1}{2}\left((2\alpha-1)^{k-2}\overline{q}-q\right)(1-2\alpha(1-\delta))>0 \end{aligned}$$

In this case, $p_{s+1} > \frac{\widehat{q_s}+q}{2}$, and buyer $s+1$ would make a negative surplus, which cannot happen in equilibrium.

Moreover, I will prove that none of the future offers will exceed $\alpha (2\alpha - 1)^{k-1} \overline{q}$. Suppose there is an equilibrium in which buyers in the future will make offers greater than $\alpha (2\alpha - 1)^{k-1} \overline{q}$. I let p^M be the upper bound of these offers. There is a p such that $p - \alpha (2\alpha - 1)^{k-1} \overline{q} = \delta \left(p^M - (2\alpha - 1)^{k-1} \overline{q} \right)$. This indicates that any offer greater than p will be accepted by the seller type up to $\alpha (2\alpha - 1)^{k-1} \overline{q}$. No equilibrium offer will exceed p, which contradicts the definition of p^M. All future offers will not exceed $\alpha (2\alpha - 1)^{k-1} \overline{q}$. This complete the proof the claim.

This shows one possible equilibrium construction of the infinite-horizon game. Notice that I assume $\widehat{q_s} > (2\alpha - 1)^n \overline{q}$, and the disclosure threshold is $\alpha (2\alpha - 1)^n \overline{q}$. In this case, the equilibrium offers will be disclosed after period s. Although this construction does not fall into the category of the ST equilibrium, it is crucial for my analysis.

When the cutoff type $\widehat{q_s} \leq (2\alpha - 1)^n \overline{q}$, from the previous discussion, I know that any future price will not exceed $\alpha (2\alpha - 1)^n \overline{q}$. Under the condition that $\widehat{q_s} \geq (2\alpha - 1)^{n+1} \overline{q}$, buyer $s+1$ will offer $p_{s+1} = \alpha (2\alpha - 1)^n \overline{q}$ which will induce the cutoff type $(2\alpha - 1)^n \overline{q}$. All future offers will not exceed $\alpha (2\alpha - 1)^n \overline{q}$. When $\widehat{q_s} < (2\alpha - 1)^{n+1} \overline{q}$, offering $\alpha (2\alpha - 1)^n \overline{q}$ will give buyer $s+1$ a negative surplus. However, offering any other price below $\alpha (2\alpha - 1)^n \overline{q}$ will create an incentive for buyer $s+1$ to deviate. This cannot happen in equilibrium.

I can construct one type of the ST equilibrium in the following way. There exists a $n^{\star}$ such that, for some $\widehat{q_{n^{\star}}} = (2\alpha-1)^{n}\overline{q}$, the seller will disclose any offer above $\alpha\widehat{q_{n^{\star}}}$ in every $t \geq n^{\star}$. All future prices will not exceed $\alpha\widehat{q_{n^{\star}}}$. The cutoff type will remain unchanged in every $t \geq n^{\star}$, i.e., $\widehat{q_{n^{\star}}} = \widehat{q_{n^{\star}+1}} = \cdots$. The value of the disclosure threshold $\widehat{p_t}\left(q, p^{t-1}\right)$ coincides with the equilibrium price p_t^{*} when $t \leq n^{\star}$.

The same technique can be applied to construct the equilibrium in which there is a threshold $\widehat{p}$ and n ($n \leq 1$) so that $\alpha(2\alpha-1)^{n+1}\overline{q} < \widehat{p} < \alpha(2\alpha-1)^{n}\overline{q}$. In this equilibrium, every seller type will disclose any offer above $\widehat{p}$ in all future periods after some period s. This type of equilibrium is similar to the type of equilibrium I have discussed previously. The construction is almost the same. Similarly, to construct the ST equilibrium, I need the cutoff type $\widehat{q_s} \geq (2\alpha-1)^{n+1}\overline{q}$. In period $t > s$, the equilibrium offer will be $\alpha(2\alpha-1)^{n}\overline{q}$, which will be above the disclosure threshold. In this case, some equilibrium offers will be disclosed, and this equilibrium does not fall into the category of the ST equilibrium.

(2) Suppose that there is no such threshold $\widehat{p} < \alpha(2\alpha-1)\overline{q}$ that every seller type will disclose any offer above $\widehat{p}$ in all future periods after some period s. In this case, the value of the equilibrium disclosure thresholds must be weakly increasing across periods, and so do the cutoff types. Otherwise, suppose that there exists $\widehat{p_k}\left(q, p^{k-1}\right) > \widehat{p_{k+1}}\left(q, p^{k}\right)$. For some small ε, this suggests that the offer $\widehat{p_k} - \varepsilon$ made in period k will be concealed to buyer $k+2$, but $\widehat{p_k} - 2\varepsilon$ made in period $k+1$ will be disclosed. However, disclosing $\widehat{p_k} - \varepsilon$ made by buyer k will lead buyer $k+2$ to update his belief to a much higher cutoff type than disclosing $\widehat{p_k} - 2\varepsilon$ made by buyer $k+1$. From sequential rationality, the value of the disclosure thresholds must be weakly increasing, i.e., $\widehat{p_k}\left(q, p^{k-1}\right) \leq \widehat{p_{k+1}}\left(q, p^{k}\right)$.

In this category of equilibrium, type $\overline{q}$ is eventually served. Let $n^{\star\star}$ denote the first period that type $\overline{q}$ is served with positive probability. I claim that when serving type $\overline{q}$, a buyer randomizes between 2 offers, one of which is $\alpha\overline{q}$. Suppose there is an equilibrium in which $\alpha\overline{q}$ is offered with probability 1 in period $n^{\star\star}$. The cutoff type in period $n^{\star\star}-1$ is denoted by $\widehat{q_{n^{\star\star}-1}}$, and $\widehat{q_{n^{\star\star}-1}} < \overline{q}$. The cutoff type satisfies $p_{n^{\star\star}} - \alpha\widehat{q_{n^{\star\star}-1}} = \delta\left(\alpha\overline{q} - \alpha\widehat{q_{n^{\star\star}-1}}\right)$. Buyer $n^{\star\star}-1$'s surplus function is

$$\frac{\widehat{q_{n^{\star\star}-1}}-\widehat{q_{n^{\star\star}-2}}}{\overline{q}-\widehat{q_{n^{\star\star}-2}}}\left(\frac{\widehat{q_{n^{\star\star}-1}}+\widehat{q_{n^{\star\star}-2}}}{2}-\delta\alpha\overline{q}-\alpha(1-\delta)\widehat{q_{n^{\star\star}-1}}\right)$$

This function is increasing in $\widehat{q_{n^{\star\star}-1}}$. For any $\widehat{q_{n^{\star\star}-1}}<\overline{q}$, buyer $n^{\star\star}-1$ would have an incentive to raise his price and serve the higher types. Therefore, in equilibrium, $\alpha\overline{q}$ must be offered with probability strictly less than 1. From Lemma 8, I know that buyer $n^{\star\star}$ must randomize between two offers, i.e., $p^H=\alpha\overline{q}$ and $p^L_{n^{\star\star}}$.

In addition, I will prove that $\widehat{q_{n^{\star\star}-1}}=(2\alpha-1)\overline{q}$. If $\widehat{q_{n^{\star\star}-1}}<(2\alpha-1)\overline{q}$, buyer $n^{\star\star}$ will have a negative payoff when offering $\alpha\overline{q}$, which cannot be possible. However, if $\widehat{q_{n^{\star\star}-1}}>(2\alpha-1)\overline{q}$, offering $\alpha\overline{q}$ will yield a positive payoff. To make buyer $n^{\star\star}$ willing to randomize, offering $p^L_{n^{\star\star}}$ must also yield a positive payoff. I use $q'_{n^{\star\star}}$ ($q'_{n^{\star\star}}>\widehat{q_{n^{\star\star}-1}}$) to denote the cutoff type accepting $p^L_{n^{\star\star}}$. Also, I have $\left(q'_{n^{\star\star}}-\widehat{q_{n^{\star\star}-1}}\right)\left(\frac{q'_{n^{\star\star}}+\widehat{q_{n^{\star\star}-1}}}{2}-p^L_{n^{\star\star}}\right)=\left(\overline{q}-\widehat{q_{n^{\star\star}-1}}\right)\left(\frac{\overline{q}+\widehat{q_{n^{\star\star}-1}}}{2}-\alpha\overline{q}\right)$. The payoff of offering $p^L_{n^{\star\star}}$ satisfies

$$\pi_{n^{\star\star}}=\left(q'_{n^{\star\star}}-\widehat{q_{n^{\star\star}-1}}\right)\left(\frac{q'_{n^{\star\star}}+\widehat{q_{n^{\star\star}-1}}}{2}-p^L_{n^{\star\star}}\right)\leq\left(q'_{n^{\star\star}}-\widehat{q_{n^{\star\star}-1}}\right)\left(\frac{q'_{n^{\star\star}}+\widehat{q_{n^{\star\star}-1}}}{2}-\alpha q'_{n^{\star\star}}\right)$$

Similar to what I have done in proving Proposition 19, when $\widehat{q_{n^{\star\star}-1}}\geq\frac{2\alpha-1}{\alpha}\overline{q}$ is satisfied, $\left(q'_{n^{\star\star}}-\widehat{q_{n^{\star\star}-1}}\right)\left(\frac{q'_{n^{\star\star}}+\widehat{q_{n^{\star\star}-1}}}{2}-\alpha q'_{n^{\star\star}}\right)$ is increasing in $q'_{n^{\star\star}}$, and the maximum is reached at $q'_{n^{\star\star}}=\overline{q}$. In this case, buyer $n^{\star\star}$ cannot randomize between the two offers $p^H=\alpha\overline{q}$ and $p^L_{n^{\star\star}}$. If $\widehat{q_{n^{\star\star}-1}}<\frac{2\alpha-1}{\alpha}\overline{q}$ and $\widehat{q_{n^{\star\star}}}\geq\frac{2\alpha-1}{\alpha}\overline{q}$, for buyer $n^{\star\star}+1$, there is no offer that gives him the same payoff as the offer $\alpha\overline{q}$. Therefore, there will be no randomization in period $n^{\star\star}+1$, which presents a contradiction. The last possible case is $\widehat{q_{n^{\star\star}}}<\frac{2\alpha-1}{\alpha}\overline{q}$. However, as I have done in proving Proposition 19, $\alpha\overline{q}$ generates a strictly higher payoff than any lower $p^L_{n^{\star\star}}$; thus, buyer $n^{\star\star}$ cannot randomize. This shows another contradiction.

In conclusion, $\widehat{q_{n^{\star\star}-1}}=(2\alpha-1)\overline{q}$. In this case, buyer $n^{\star\star}$ ends up with a payoff of zero, and the offer $p^L_{n^{\star\star}}$ is rejected by any type $q\geq(2\alpha-1)\overline{q}$. This also suggests that I must have stationary cutoff types in all future periods. For any $t \geq n^{\star\star}$, buyer t's lower offer cannot exceed the disclosure threshold $\widehat{p_t}$, and the seller's continuation value is also weakly higher than $\widehat{p_t}$. As $\widehat{p_t}$ is non-decreasing, I have stationarity in the continuation value. In period $n^{\star\star}$, the continuation value of the cutoff type $(2\alpha-1)\overline{q}$ is equal to the equilibrium offer in period $n^{**}-1$, which is $p_{n^{\star\star}-1}$. Any buyer t ($t \geq n^{\star\star}$) will randomize between the two offers $\alpha\overline{q}$ and p_t^L where $p_t^L \leq p_{n^{\star\star}-1} \leq \frac{\widehat{q_{n^{\star\star}-1}}+\widehat{q_{n^{\star\star}-2}}}{2}$. The probability of offering $\alpha\overline{q}$ is λ such that the following condition holds.

$$p_{n^{\star\star}-1}-\alpha\widehat{q_{n^{\star\star}-1}}=\delta\left(\lambda\alpha\overline{q}+(1-\lambda)\,p_{n^{\star\star}-1}-\alpha\widehat{q_{n^{\star\star}-1}}\right)$$

When $t \leq n^{\star\star}-1$, the value of the disclosure threshold $\widehat{p_t}\left(q,p^{t-1}\right)$ matches the equilibrium price p_t^* on the equilibrium path. Moreover, $\widehat{p_t}\left(q,p^{t-1}\right)$ matches $p_{n^{\star\star}-1}$ when $t > n^{\star\star}-1$. ■

Appendix C

Proofs for Chapter 3

In Appendix C, I provide detailed proofs for most of the propositions and lemmas in Chapter 3.

C.1 Proofs for Chapter 3.3

When contracting with a present-biased agent, the seller's profit maximization problem is

$$\max_{\{t_s,x_s\}}\left\{\mu\left[t_1\left(\overline{\theta}\right)-\frac{1}{2}x_1^2\left(\overline{\theta}\right)+\delta E\left(\pi_2\left(\theta|\overline{\theta}\right)\right)\right]+(1-\mu)\left(t_1\left(\underline{\theta}\right)-\frac{1}{2}x_1^2\left(\underline{\theta}\right)+\delta E\left(\pi_2\left(\theta|\underline{\theta}\right)\right)\right)\right\}$$

subject to the IC and IR constraints for all θ_1 and θ_1', θ_2 and θ_2', θ_3 and θ_3', where

$$
\begin{aligned}
\text{IC1} \quad : \quad & \theta_1 x_1(\theta_1) - t_1(\theta_1) + \beta\delta E_{\theta_2}(\theta_2 x_2(\theta_2|\theta_1) - t_2(\theta_2|\theta_1)|\theta_1) \\
& + \beta\delta^2 E_{\theta_2,\theta_3}(\theta_3 x_3(\theta_3|\theta_1,\theta_2) - t_3(\theta_3|\theta_1,\theta_2)|\theta_1) \\
\geq \quad & \theta_1 x_1(\theta_1') - t_1(\theta_1') + \beta\delta E_{\theta_2}(\theta_2 x_2(\theta_2|\theta_1') - t_2(\theta_2|\theta_1')|\theta_1) \\
& + \beta\delta^2 E_{\theta_2,\theta_3}(\theta_3 x_3(\theta_3|\theta_1',\theta_2) - t_3(\theta_3|\theta_1',\theta_2)|\theta_1) \\
\text{IC2} \quad : \quad & \theta_2 x_2(\theta_2|\theta_1) - t_2(\theta_2|\theta_1) + \beta\delta E_{\theta_3}(\theta_3 x_3(\theta_3|\theta_1,\theta_2) - t_3(\theta_3|\theta_1,\theta_2)|\theta_2) \\
\geq \quad & \theta_2 x_2(\theta_2'|\theta_1) - t_2(\theta_2'|\theta_1) + \beta\delta E_{\theta_3}(\theta_3 x_3(\theta_3|\theta_1,\theta_2') - t_3(\theta_3|\theta_1,\theta_2')|\theta_2) \\
\text{IC3} \quad : \quad & \theta_3 x_3(\theta_3|\theta_1,\theta_2) - t_3(\theta_3|\theta_1,\theta_2) \geq \theta_3 x_3(\theta_3'|\theta_1,\theta_2) - t_3(\theta_3'|\theta_1,\theta_2) \\
\text{IR1} \quad : \quad & 0 \leq \theta_1 x_1(\theta_1) - t_1(\theta_1) + \beta\delta E_{\theta_2}(\theta_2 x_2(\theta_2|\theta_1) - t_2(\theta_2|\theta_1)|\theta_1) \\
& + \beta\delta^2 E_{\theta_2,\theta_3}(\theta_3 x_3(\theta_3|\theta_1,\theta_2) - t_3(\theta_3|\theta_1,\theta_2)|\theta_1) \\
\text{IR2} \quad : \quad & \theta_2 x_2(\theta_2|\theta_1) - t_2(\theta_2|\theta_1) + \beta\delta E_{\theta_3}(\theta_3 x_3(\theta_3|\theta_1,\theta_2) - t_3(\theta_3|\theta_1,\theta_2)|\theta_2) \geq 0 \\
\text{IR3} \quad : \quad & \theta_3 x_3(\theta_3|\theta_1,\theta_2) - t_3(\theta_3|\theta_1,\theta_2) \geq 0
\end{aligned}
$$

Proof of Proposition 13. The proof can be found in [5], and it is omitted here. ■

Proof of Lemma 6. Lemma 6 can be easily proved using the IR and the IC constraints characterized above. The proof is standard in the mechanism design literature, and it is omitted here. ■

Proof of Proposition 14. Applying the results in Lemma 6, I can characterize the optimal payment scheme, which satisfies

$$
\begin{aligned}
t_3\left(\underline{\theta}|\theta_1,\theta_2\right) &= \underline{\theta}x_3\left(\underline{\theta}|\theta_1,\theta_2\right) \\
t_3\left(\overline{\theta}|\theta_1,\theta_2\right) &= \overline{\theta}x_3\left(\overline{\theta}|\theta_1,\theta_2\right)-\Delta\theta x_3\left(\underline{\theta}|\theta_1,\theta_2\right) \\
t_2\left(\underline{\theta}|\theta_1\right) &= \underline{\theta}x_2\left(\underline{\theta}|\theta_1\right)+\beta\delta\left(1-r\right)\Delta\theta x_3\left(\underline{\theta}|\theta_1,\underline{\theta}\right) \\
t_2\left(\overline{\theta}|\theta_1\right) &= \overline{\theta}x_2\left(\overline{\theta}|\theta_1\right)-\Delta\theta x_2\left(\underline{\theta}|\theta_1\right)+\beta\delta\left(1-r\right)\Delta\theta x_3\left(\underline{\theta}|\theta_1,\underline{\theta}\right)+\alpha\beta\delta\Delta\theta\left(x_3\left(\underline{\theta}|\theta_1,\overline{\theta}\right)-x_3\left(\underline{\theta}|\theta_1,\underline{\theta}\right)\right) \\
t_1\left(\underline{\theta}\right) &= \underline{\theta}x_1\left(\underline{\theta}\right)+\left(1-\gamma\right)\beta\delta\Delta\theta x_2\left(\underline{\theta}|\underline{\theta}\right)+\left(1-\gamma\right)\beta^2\delta^2\left(\alpha+\gamma-1\right)\Delta\theta x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right) \\
&\quad+\left(1-\gamma\right)\beta\left(1-\beta\right)\alpha\delta^2\Delta\theta x_3\left(\underline{\theta}|\underline{\theta},\overline{\theta}\right)+\left(1-\gamma\right)\beta\left(1-\beta\right)\gamma\delta^2\Delta\theta x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right) \\
t_1\left(\overline{\theta}\right) &= \overline{\theta}x_1\left(\overline{\theta}\right)-\Delta\theta x_1\left(\underline{\theta}\right)+\left(1-\gamma\right)\beta\delta\Delta\theta x_2\left(\underline{\theta}|\underline{\theta}\right)+\left(1-\gamma\right)\beta^2\delta^2\left(\alpha+\gamma-1\right)\Delta\theta x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right) \\
&\quad+\left(1-\gamma\right)\beta\left(1-\beta\right)\alpha\delta^2\Delta\theta x_3\left(\underline{\theta}|\underline{\theta},\overline{\theta}\right)+\left(1-\gamma\right)\beta\left(1-\beta\right)\gamma\delta^2\Delta\theta x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right) \\
&\quad+\alpha\beta\delta\Delta\theta\left(x_2\left(\underline{\theta}|\overline{\theta}\right)-x_2\left(\underline{\theta}|\underline{\theta}\right)\right)+\alpha\beta^2\delta^2\left(\alpha+\gamma-1\right)\Delta\theta\left(x_3\left(\underline{\theta}|\overline{\theta},\underline{\theta}\right)-x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right)\right) \\
&\quad+\beta\delta^2\left(1-\beta\right)\alpha^2\Delta\theta\left(x_3\left(\underline{\theta}|\overline{\theta},\overline{\theta}\right)-x_3\left(\underline{\theta}|\underline{\theta},\overline{\theta}\right)\right) \\
&\quad+\beta\delta^2\left(1-\beta\right)\left(1-\alpha\right)\left(1-\gamma\right)\Delta\theta\left(x_3\left(\underline{\theta}|\overline{\theta},\underline{\theta}\right)-x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right)\right)
\end{aligned}
$$

The principal's profit function can be transformed into

$$
\begin{aligned}
E\pi \quad = \quad & \mu\left(\overline{\theta}-\frac{1}{2}x_1\left(\overline{\theta}\right)\right)x_1\left(\overline{\theta}\right)+\left[(1-\mu)\left(\underline{\theta}-\frac{1}{2}x_1\left(\underline{\theta}\right)\right)-\mu\Delta\theta\right]x_1\left(\underline{\theta}\right) \\
& +\mu\alpha\delta\left(\overline{\theta}-\frac{1}{2}x_2\left(\overline{\theta}|\overline{\theta}\right)\right)x_2\left(\overline{\theta}|\overline{\theta}\right)+(1-\mu)(1-\gamma)\delta\left(\overline{\theta}-\frac{1}{2}x_2\left(\overline{\theta}|\underline{\theta}\right)\right)x_2\left(\overline{\theta}|\underline{\theta}\right) \\
& +\mu\delta\left[(1-\alpha)\left(\underline{\theta}-\frac{1}{2}x_2\left(\underline{\theta}|\overline{\theta}\right)\right)-\alpha(1-\beta)\Delta\theta\right]x_2\left(\underline{\theta}|\overline{\theta}\right) \\
& +\delta\left[(1-\mu)\gamma\left(\underline{\theta}-\frac{1}{2}x_2\left(\underline{\theta}|\underline{\theta}\right)\right)-\mu(\alpha+\gamma-1)\beta\Delta\theta-(1-\mu)(1-\gamma)(1-\beta)\Delta\theta\right]x_2\left(\underline{\theta}|\underline{\theta}\right) \\
& +\mu\alpha^2\delta^2\left(\overline{\theta}-\frac{1}{2}x_3\left(\overline{\theta}|\overline{\theta},\overline{\theta}\right)\right)x_3\left(\overline{\theta}|\overline{\theta},\overline{\theta}\right)+\mu(1-\alpha)(1-\gamma)\delta^2\left(\overline{\theta}-\frac{1}{2}x_3\left(\overline{\theta}|\overline{\theta},\underline{\theta}\right)\right)x_3\left(\overline{\theta}|\overline{\theta},\underline{\theta}\right) \\
& +(1-\mu)(1-\gamma)\alpha\delta^2\left(\overline{\theta}-\frac{1}{2}x_3\left(\overline{\theta}|\underline{\theta},\overline{\theta}\right)\right)x_3\left(\overline{\theta}|\underline{\theta},\overline{\theta}\right) \\
& +(1-\mu)\gamma(1-\gamma)\delta^2\left(\overline{\theta}-\frac{1}{2}x_3\left(\overline{\theta}|\underline{\theta},\underline{\theta}\right)\right)x_3\left(\overline{\theta}|\underline{\theta},\underline{\theta}\right) \\
& +\mu\alpha\delta^2\left[(1-\alpha)\left(\underline{\theta}-\frac{1}{2}x_3\left(\underline{\theta}|\overline{\theta},\overline{\theta}\right)\right)-\alpha(1-\beta)^2\Delta\theta\right]x_3\left(\underline{\theta}|\overline{\theta},\overline{\theta}\right) \\
& +\mu\delta^2\left[(1-\alpha)\gamma\left(\underline{\theta}-\frac{1}{2}x_3\left(\underline{\theta}|\overline{\theta},\underline{\theta}\right)\right)-(1-\beta)\left(\alpha^2\beta+(1-\gamma)(1-\beta-\alpha)\right)\Delta\theta\right]x_3\left(\underline{\theta}|\overline{\theta},\underline{\theta}\right) \\
& +\delta^2\left[(1-\mu)(1-\gamma)(1-\alpha)\left(\underline{\theta}-\frac{1}{2}x_3\left(\underline{\theta}|\underline{\theta},\overline{\theta}\right)\right)-\mu\alpha(\alpha+\gamma-1)\beta(1-\beta)\Delta\theta\right]x_3\left(\underline{\theta}|\underline{\theta},\overline{\theta}\right) \\
& -\alpha\delta^2(1-\mu)(1-\gamma)(1-\beta)^2\Delta\theta x_3\left(\underline{\theta}|\underline{\theta},\overline{\theta}\right) \\
& +\delta^2\left[(1-\mu)\gamma^2\left(\underline{\theta}-\frac{1}{2}x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right)\right)-\mu\beta(\alpha+\gamma-1)(\alpha\beta+\gamma-1)\Delta\theta\right]x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right) \\
& -\delta^2\left[(1-\mu)(1-\gamma)(\alpha+\gamma-1)\beta(1-\beta)\Delta\theta+(1-\mu)\gamma(1-\gamma)(1-\beta)^2\Delta\theta\right]x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right)
\end{aligned}
$$

So the optimal allocations are

$$\begin{aligned}
x_1\left(\overline{\theta}\right) &= x_2\left(\overline{\theta}|\theta_1\right) = x_3\left(\overline{\theta}|\theta_1,\theta_2\right) = \overline{\theta}, \\
x_1\left(\underline{\theta}\right) &= \underline{\theta} - \frac{\mu}{1-\mu}\Delta\theta \\
x_2\left(\underline{\theta}|\overline{\theta}\right) &= \underline{\theta} - \frac{\alpha}{1-\alpha}\left(1-\beta\right)\Delta\theta \\
x_2\left(\underline{\theta}|\underline{\theta}\right) &= \underline{\theta} - \beta\frac{\mu\left(\alpha+\gamma-1\right)}{\left(1-\mu\right)\gamma}\Delta\theta - \left(1-\beta\right)\frac{1-\gamma}{\gamma}\Delta\theta \\
x_3\left(\underline{\theta}|\overline{\theta},\overline{\theta}\right) &= \underline{\theta} - \left(1-\beta\right)^2\frac{\alpha}{1-\alpha}\Delta\theta \\
x_3\left(\underline{\theta}|\overline{\theta},\underline{\theta}\right) &= \underline{\theta} - \beta\left(1-\beta\right)\frac{\alpha\left(\alpha+\gamma-1\right)}{\left(1-\alpha\right)\gamma}\Delta\theta - \left(1-\beta\right)^2\frac{1-\gamma}{\gamma}\Delta\theta \\
x_3\left(\underline{\theta}|\underline{\theta},\overline{\theta}\right) &= \underline{\theta} - \beta\left(1-\beta\right)\frac{\mu\alpha\left(\alpha+\gamma-1\right)}{\left(1-\mu\right)\left(1-\gamma\right)\left(1-\alpha\right)}\Delta\theta - \left(1-\beta\right)^2\frac{\alpha}{1-\alpha}\Delta\theta \\
x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right) &= \underline{\theta} - \beta^2\frac{\mu\left(\alpha+\gamma-1\right)^2}{\left(1-\mu\right)\gamma^2}\Delta\theta + \beta\left(1-\beta\right)\frac{\mu\left(1-\gamma\right)\left(\alpha+\gamma-1\right)}{\left(1-\mu\right)\gamma^2}\Delta\theta \\
&\quad -\beta\left(1-\beta\right)\frac{\left(\alpha+\gamma-1\right)\left(1-\gamma\right)}{\gamma^2}\Delta\theta - \left(1-\beta\right)^2\frac{1-\gamma}{\gamma}\Delta\theta
\end{aligned}$$

■

Proof of Proposition 15. Using the results in Lemma 6, I only need to focus on the IR constraint for the high type in period 1 and the three IC constraints for the low type. All the constraints can be transformed into functions of β. I first prove that these constraints hold as strict inequalities under $\beta = 1$. For the constraints in period 1, when $\beta = 1$,

$$\begin{aligned}
IR_1\left(\overline{\theta}\right) &: \overline{\theta}x_1\left(\overline{\theta}\right) - t_1\left(\overline{\theta}\right) + \delta E\left(\widehat{U}_2\left(\theta_2|\overline{\theta}\right)|\overline{\theta}\right) \\
&= \Delta\theta\sum_{j=0}^{2}\delta^j\left(\alpha+\gamma-1\right)^j x_{1+j}\left(\underline{\theta}|\underline{\theta},\ldots,\underline{\theta}\right) > 0 \\
IC_1\left(\underline{\theta}\right) &: \underline{\theta}x_1\left(\underline{\theta}\right) - t_1\left(\underline{\theta}\right) + \delta E\left(\widehat{U}_2\left(\theta_2|\underline{\theta}\right)|\underline{\theta}\right) - \left[\underline{\theta}x_1\left(\overline{\theta}\right) - t_1\left(\overline{\theta}\right) + \delta E\left(\widehat{U}_2\left(\theta_2|\overline{\theta}\right)|\underline{\theta}\right)\right] \\
&= \Delta\theta^2\left(\frac{1}{1-\mu}\right) + \Delta\theta^2\delta\frac{\left(\alpha+\gamma-1\right)^2}{\gamma}\left(\frac{\mu}{1-\mu}\right) + \Delta\theta^2\delta^2\frac{\left(\alpha+\gamma-1\right)^4}{\gamma^2}\left(\frac{\mu}{1-\mu}\right) > 0
\end{aligned}$$

where

$$\widehat{U}_{s+1}\left(\theta_{s+1}|h_s,\theta_s\right) = \theta_{s+1}x_{s+1}\left(\theta_{s+1}|h_s,\theta_s\right) - t_{s+1}\left(\theta_{s+1}|h_s,\theta_s\right) + \delta E\left(\widehat{U}_{s+1}\left(\theta_{s+2}|h_{s+1},\theta_{s+1}\right)|\theta_{s+1}\right)$$

Similarly, the IC constraints in period 2 and period 3 satisfy the following conditions.

$$\begin{aligned}
IC_2\left(\underline{\theta}|\theta_1\right) &: \quad \underline{\theta}x_2\left(\underline{\theta}|\theta_1\right)-t_2\left(\underline{\theta}|\theta_1\right)+\delta E\left(\widehat{U}_3\left(\theta_3|\theta_1,\underline{\theta}\right)\right)-\left[\underline{\theta}x_2\left(\overline{\theta}|\theta_1\right)-t_2\left(\overline{\theta}|\theta_1\right)+\delta E\left(\widehat{U}_3\left(\theta_3|\theta_1,\overline{\theta}\right)\right)\right] \\
&= \quad \Delta\theta\left(x_2\left(\overline{\theta}|\theta_1\right)-x_2\left(\underline{\theta}|\theta_1\right)\right)+\Delta\theta\delta\left(\alpha+\gamma-1\right)\left(x_3\left(\underline{\theta}|\theta_1,\overline{\theta}\right)-x_3\left(\underline{\theta}|\theta_1,\underline{\theta}\right)\right)>0 \\
IC_3\left(\underline{\theta}|\theta_1,\theta_2\right) &: \quad \underline{\theta}x_3\left(\underline{\theta}|\theta_1,\theta_2\right)-t_3\left(\underline{\theta}|\theta_1,\theta_2\right)-\left[\underline{\theta}x_3\left(\overline{\theta}|\theta_1,\theta_2\right)-t_3\left(\overline{\theta}|\theta_1,\theta_2\right)\right] \\
&= \quad \Delta\theta\left(x_3\left(\overline{\theta}|\theta_1,\theta_2\right)-x_3\left(\underline{\theta}|\theta_1,\theta_2\right)\right)>0
\end{aligned}$$

As the constraints are continuous functions of β, there exists a threshold $\underline{\beta}$ such that for $\beta \geq \underline{\beta}$, the allocations listed before are indeed optimal.

■

Proof of Proposition 16. The seller's profit function is continuous and differentiable in β. Using the Envelop theorem, I have the following expression.

$$\begin{aligned}
\frac{dE\pi}{d\beta} &= \mu\alpha\delta\Delta\theta x_2\left(\underline{\theta}|\overline{\theta}\right)+\left[\left(1-\mu\right)\left(1-\gamma\right)-\mu\left(\alpha+\gamma-1\right)\right]\delta\Delta\theta x_2\left(\underline{\theta}|\underline{\theta}\right) \\
&+2\mu\alpha^2\delta^2\left(1-\beta\right)\Delta\theta x_3\left(\underline{\theta}|\overline{\theta},\overline{\theta}\right)-\mu\alpha\left(\alpha+\gamma-1\right)\delta^2\left(1-2\beta\right)\Delta\theta x_3\left(\underline{\theta}|\underline{\theta},\overline{\theta}\right) \\
&+\mu\delta^2\left(\alpha^2\beta+\left(1-\gamma\right)\left(1-\beta-\alpha\right)-\left(1-\beta\right)\left(\alpha^2+\gamma-1\right)\right)\Delta\theta x_3\left(\underline{\theta}|\overline{\theta},\underline{\theta}\right) \\
&-\delta^2\left(\mu\left(\alpha+\gamma-1\right)\left(2\alpha\beta+\gamma-1\right)-2\left(1-\mu\right)\gamma\left(1-\gamma\right)\left(1-\beta\right)\right)\Delta\theta x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right) \\
&-\delta^2\left(1-\mu\right)\left(1-\gamma\right)\left(\alpha+\gamma-1\right)\left(1-2\beta\right)\Delta\theta x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right)
\end{aligned}$$

When $\beta=1$, I have

$$\begin{aligned}
\frac{dE\pi}{d\beta}|_{\beta=1} &= \mu\alpha\delta\Delta\theta\overline{\theta}+\left[\left(1-\mu\right)\left(1-\gamma\right)-\mu\left(\alpha+\gamma-1\right)\right]\delta\Delta\theta x_2\left(\underline{\theta}|\underline{\theta}\right) \\
&+2\mu\alpha\left(\alpha+\gamma-1\right)\delta^2\Delta\theta\overline{\theta}-\mu\left(\alpha+\gamma-1\right)\left(2\alpha+\gamma-1\right)\delta^2\Delta\theta x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right) \\
&+\delta^2\left(1-\mu\right)\left(1-\gamma\right)\left(\alpha+\gamma-1\right)\Delta\theta x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right)
\end{aligned}$$

Since $x_2\left(\underline{\theta}|\underline{\theta}\right)<\overline{\theta}$ and $x_3\left(\underline{\theta}|\underline{\theta},\underline{\theta}\right)<\overline{\theta}$, $\frac{dE\pi}{d\beta}>0$ at $\beta=1$. As $\frac{dE\pi}{d\beta}$ is a continuous function of β, when β is close to 1, $\frac{dE\pi}{d\beta}$ is positive. Therefore, there exists a β' such that the seller's profit function is increasing in β when $\beta \geq \beta'$. ■

C.2 Proofs and Supplementary Materials for Chapter 3.4

Lemma 9 *The incentive compatibility constraint in period* 3 *is satisfied if and only if (i)* $x_3(\theta_3|\theta_1,\theta_2)$ *is non-decreasing in* θ_3*; (ii)* $V_3(\theta_3|\theta_1,\theta_2) = V_3(\underline{\theta}|\theta_1,\theta_2) + \int_{\underline{\theta}}^{\theta_3} x_3(s|\theta_1,\theta_2)\,ds$.

Proof of Lemma 9. This lemma results directly from the single crossing condition and the envelope condition. The proof is fairly standard, and it is omitted here. ■

Lemma 10 *The incentive compatibility constraint in period* 2 *is satisfied only if the following conditions hold.*

$$0 \leq (\theta_2-\theta_2')\left(x_2(\theta_2|\theta_1) - x_2(\theta_2'|\theta_1)\right) - \beta\delta\int_{\underline{\theta}}^{\overline{\theta}} \left[x_3(\theta_3|\theta_1,\theta_2) - x_3(\theta_3|\theta_1,\theta_2')\right]\left[F_3(\theta_3|\theta_2) - F_3(\theta_3|\theta_2')\right] d\theta_3$$

$$V_2'(\theta_2|\theta_1) = x_2(\theta_2|\theta_1) - \beta\delta\int_{\underline{\theta}}^{\overline{\theta}} x_3(\theta_3|\theta_1,\theta_2)\frac{\partial F_3(\theta_3|\theta_2)}{\partial\theta_2} d\theta_3$$

Proof of Lemma 10. Taking any θ_1 and θ_1', I have that

$$\begin{aligned}
V_2(\theta_2|\theta_1) &\geq V_2(\theta_2',\theta_2|\theta_1) = V_2(\theta_2'|\theta_1) + (\theta_2-\theta_2')x_2(\theta_2'|\theta_1) \\
&\quad + \beta\delta\int_{\underline{\theta}}^{\overline{\theta}} V_3(\theta_3|\theta_1,\theta_2')(f_3(\theta_3|\theta_2) - f_3(\theta_3|\theta_2'))d\theta_3 \\
V_2(\theta_2|\theta_1) - V_2(\theta_2'|\theta_1) &\geq (\theta_2-\theta_2')x_2(\theta_2'|\theta_1) + \beta\delta\int_{\underline{\theta}}^{\overline{\theta}} V_3(\theta_3|\theta_1,\theta_2')(f_3(\theta_3|\theta_2) - f_3(\theta_3|\theta_2'))d\theta_3
\end{aligned}$$

Similarly, I also have

$$V_2(\theta_2|\theta_1) - V_2(\theta_2'|\theta_1) \leq (\theta_2-\theta_2')x_2(\theta_2|\theta_1) + \beta\delta\int_{\underline{\theta}}^{\overline{\theta}} V_3(\theta_3|\theta_1,\theta_2)(f_3(\theta_3|\theta_2) - f_3(\theta_3|\theta_2'))d\theta_3$$

Furthermore,

$$\int_{\underline{\theta}}^{\overline{\theta}} V_3(\theta_3|\theta_1,\theta_2) f_3(\theta_3|\theta_2)\,d\theta_3 = V_3(\overline{\theta}|\theta_1,\theta_2') - \int_{\underline{\theta}}^{\overline{\theta}} x_3(\theta_3|\theta_1,\theta_2) F_3(\theta_3|\theta_2)\,d\theta_3$$

$$\Rightarrow \quad V_2(\theta_2|\theta_1) - V_2(\theta_2'|\theta_1) \leq (\theta_2 - \theta_2') x_2(\theta_2|\theta_1) - \beta\delta \left[\int_{\underline{\theta}}^{\overline{\theta}} x_3(\theta_3|\theta_1,\theta_2)(F_3(\theta_3|\theta_2) - F_3(\theta_3|\theta_2'))d\theta_2 \right]$$

$$\Rightarrow \quad V_2(\theta_2|\theta_1) - V_2(\theta_2'|\theta_1) \geq (\theta_2 - \theta_2') x_2(\theta_2'|\theta_1) - \beta\delta \left[\int_{\underline{\theta}}^{\overline{\theta}} x_3(\theta_3|\theta_1,\theta_2')(F_3(\theta_3|\theta_2) - F_3(\theta_3|\theta_2'))d\theta_2 \right]$$

Dividing both sides of the above inequalities by $\theta_2 - \theta_2'$ and taking limits, I get

$$V_2'(\theta_2|\theta_1) = x_2(\theta_2|\theta_1) - \beta\delta \int_{\underline{\theta}}^{\overline{\theta}} x_3(\theta_3|\theta_1,\theta_2) \frac{\partial F_3(\theta_3|\theta_2)}{\partial \theta_2} d\theta_3$$

Rearranging terms, I get

$$0 \leq (\theta_2 - \theta_2')(x_2(\theta_2|\theta_1) - x_2(\theta_2'|\theta_1)) - \beta\delta \int_{\underline{\theta}}^{\overline{\theta}} [x_3(\theta_3|\theta_1,\theta_2) - x_3(\theta_3|\theta_1,\theta_2')] [F_3(\theta_3|\theta_2) - F_3(\theta_3|\theta_2')] d\theta_3$$

■

Lemma 11 *The incentive compatibility constraint in period* 1 *is satisfied only if the following monotonicity condition is satisfied*

$$\begin{aligned}
0 \leq\ & (\theta_1 - \theta_1')(x_1(\theta_1) - x_1(\theta_1')) - \beta\delta \int_{\underline{\theta}}^{\overline{\theta}} [x_2(\theta_2|\theta_1) - x_2(\theta_2|\theta_1')] [F_2(\theta_2|\theta_1) - F_2(\theta_2|\theta_1')] d\theta_2 \\
& + \beta^2\delta^2 \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} [x_3(\theta_3|\theta_1,\theta_2) - x_3(\theta_3|\theta_1',\theta_2)] [F_2(\theta_2|\theta_1) - F_2(\theta_2|\theta_1')] \frac{\partial F_3(\theta_3|\theta_2)}{\partial \theta_2} d\theta_2 d\theta_3 \\
& + (\beta - \beta^2)\delta^2 \int_{\underline{\theta}}^{\overline{\theta}} [V_3(\underline{\theta}|\theta_1,\theta_2) - V_3(\underline{\theta}|\theta_1',\theta_2)] (f_2(\theta_2|\theta_1) - f_2(\theta_2|\theta_1')) d\theta_2 \\
& + (\beta - \beta^2)\delta^2 \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} [x_3(\theta_3|\theta_1,\theta_2) - x_3(\theta_3|\theta_1',\theta_2)] [1 - F_3(\theta_3|\theta_2)] (f_2(\theta_2|\theta_1) - f_2(\theta_2|\theta_1')) d\theta_2 d\theta_3
\end{aligned}$$

and the envelope condition is satisfied.

$$\begin{aligned}
V_1'(\theta_1) =\ & x_1(\theta_1) - \beta\delta \int_{\underline{\theta}}^{\overline{\theta}} x_2(\theta_2|\theta_1) \frac{\partial F_2(\theta_2|\theta_1)}{\partial \theta_1} d\theta_2 \\
& + \beta^2\delta^2 \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} x_3(\theta_3|\theta_1,\theta_2) \frac{\partial F_3(\theta_3|\theta_2)}{\partial \theta_2} \frac{\partial F_2(\theta_2|\theta_1)}{\partial \theta_1} d\theta_2 d\theta_3 \\
& + (\beta - \beta^2)\delta^2 \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} x_3(\theta_3|\theta_1,\theta_2)[1 - F_3(\theta_3|\theta_2)] \frac{\partial f_2(\theta_2|\theta_1)}{\partial \theta_1} d\theta_2 d\theta_3 \\
& + (\beta - \beta^2)\delta^2 \int_{\underline{\theta}}^{\overline{\theta}} V_3(\underline{\theta}|\theta_1,\theta_2) \frac{\partial f_2(\theta_2|\theta_1)}{\partial \theta_1} d\theta_2
\end{aligned}$$

Proof of Lemma 11. Take any θ_1 and θ_1', the IC constraint in period 1 shows that

$$\begin{aligned} V(\theta_1) &\geq V(\theta_1',\theta_1) = V(\theta_1') + (\theta_1-\theta_1')x_1(\theta_1') + \beta\delta\int_{\underline{\theta}}^{\overline{\theta}} V_2(\theta_2|\theta_1')(f_2(\theta_2|\theta_1) - f_2(\theta_2|\theta_1'))d\theta_2 \\ &\quad + (\beta-\beta^2)\delta^2 \int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} V_3(\theta_3|\theta_1',\theta_2)(f_2(\theta_2|\theta_1) - f_2(\theta_2|\theta_1'))f_3(\theta_3|\theta_2)\,d\theta_2 d\theta_3 \\ V(\theta_1) - V(\theta_1') &\geq (\theta_1-\theta_1')x_1(\theta_1') + \beta\delta\int_{\underline{\theta}}^{\overline{\theta}} V_2(\theta_2|\theta_1')(f_2(\theta_2|\theta_1) - f_2(\theta_2|\theta_1'))d\theta_2 \\ &\quad + (\beta-\beta^2)\delta^2 \int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} V_3(\theta_3|\theta_1',\theta_2)(f_2(\theta_2|\theta_1) - f_2(\theta_2|\theta_1'))f_3(\theta_3|\theta_2)\,d\theta_2 d\theta_3 \end{aligned}$$

Similarly, I also have

$$\begin{aligned} V(\theta_1) - V(\theta_1') &\leq (\theta_1-\theta_1')x_1(\theta_1) + \beta\delta\int_{\underline{\theta}}^{\overline{\theta}} V_2(\theta_2|\theta_1)(f_2(\theta_2|\theta_1) - f_2(\theta_2|\theta_1'))\,d\theta_2 \\ &\quad + (\beta-\beta^2)\delta^2 \int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} V_3(\theta_3|\theta_1,\theta_2)(f_2(\theta_2|\theta_1) - f_2(\theta_2|\theta_1'))f_3(\theta_3|\theta_2)\,d\theta_2 d\theta_3 \end{aligned}$$

Furthermore,

$$\begin{aligned} \int_{\underline{\theta}}^{\overline{\theta}} V_2(\theta_2|\theta_1') f_2(\theta_2|\theta_1)\,d\theta_2 &= V_2(\overline{\theta}|\theta_1') - \int_{\underline{\theta}}^{\overline{\theta}} x_2(\theta_2|\theta_1') F_2(\theta_2|\theta_1)\,d\theta_2 \\ &\quad + \beta\delta\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} x_3(\theta_3|\theta_1',\theta_2)\frac{\partial F_3(\theta_3|\theta_2)}{\partial\theta_2} F_2(\theta_2|\theta_1)\,d\theta_2 d\theta_3 \end{aligned}$$

$$\begin{aligned} \int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} V_3(\theta_3|\theta_1',\theta_2) f_2(\theta_2|\theta_1) f_3(\theta_3|\theta_2)\,d\theta_2 d\theta_3 &= \int_{\underline{\theta}}^{\overline{\theta}} V_3(\overline{\theta}|\theta_1',\theta_2) f_2(\theta_2|\theta_1)\,d\theta_2 \\ &\quad - \int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} x_3(\theta_3|\theta_1',\theta_2) f_2(\theta_2|\theta_1) F_3(\theta_3|\theta_2)\,d\theta_2 d\theta_3 \end{aligned}$$

$$\begin{aligned} V(\theta_1',\theta_1) &= V(\theta_1') + (\theta_1-\theta_1')x_1(\theta_1') - \beta\delta\int_{\underline{\theta}}^{\overline{\theta}} x_2(\theta_2|\theta_1')(F_2(\theta_2|\theta_1) - F_2(\theta_2|\theta_1'))\,d\theta_2 \\ &\quad + \beta^2\delta^2\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} x_3(\theta_3|\theta_1',\theta_2)\frac{\partial F_3(\theta_3|\theta_2)}{\partial\theta_2}(F_2(\theta_2|\theta_1) - F_2(\theta_2|\theta_1'))\,d\theta_2 d\theta_3 \\ &\quad + (\beta-\beta^2)\delta^2\int_{\underline{\theta}}^{\overline{\theta}} V_3(\overline{\theta}|\theta_1',\theta_2)(f_2(\theta_2|\theta_1) - f_2(\theta_2|\theta_1'))\,d\theta_2 \\ &\quad - (\beta-\beta^2)\delta^2\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} x_3(\theta_3|\theta_1',\theta_2)(f_2(\theta_2|\theta_1) - f_2(\theta_2|\theta_1'))F_3(\theta_3|\theta_2)\,d\theta_2 d\theta_3 \end{aligned}$$

Following a similar technique established in proving Lemma 10, I can derive the envelope condition that

$$\begin{aligned} V_1'(\theta_1) &= x_1(\theta_1) - \beta\delta \int_{\underline{\theta}}^{\overline{\theta}} x_2(\theta_2|\theta_1) \frac{\partial F_2(\theta_2|\theta_1)}{\partial \theta_1} d\theta_2 \\ &+ \beta^2\delta^2 \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} x_3(\theta_3|\theta_1,\theta_2) \frac{\partial F_3(\theta_3|\theta_2)}{\partial \theta_2} \frac{\partial F_2(\theta_2|\theta_1)}{\partial \theta_1} d\theta_2 d\theta_3 \\ &+ \left(\beta - \beta^2\right)\delta^2 \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} x_3(\theta_3|\theta_1,\theta_2)\left[1 - F_3(\theta_3|\theta_2)\right] \frac{\partial f_2(\theta_2|\theta_1)}{\partial \theta_1} d\theta_2 d\theta_3 \\ &+ \left(\beta - \beta^2\right)\delta^2 \int_{\underline{\theta}}^{\overline{\theta}} V_3(\underline{\theta}|\theta_1,\theta_2) \frac{\partial f_2(\theta_2|\theta_1)}{\partial \theta_1} d\theta_2 \end{aligned}$$

and the monotonicity condition.

$$\begin{aligned} 0 &\leq \left(\theta_1 - \theta_1'\right)\left(x_1(\theta_1) - x_1\left(\theta_1'\right)\right) - \beta\delta \int_{\underline{\theta}}^{\overline{\theta}} \left[x_2(\theta_2|\theta_1) - x_2\left(\theta_2|\theta_1'\right)\right]\left[F_2(\theta_2|\theta_1) - F_2\left(\theta_2|\theta_1'\right)\right] d\theta_2 \\ &+ \beta^2\delta^2 \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \left[x_3(\theta_3|\theta_1,\theta_2) - x_3\left(\theta_3|\theta_1',\theta_2\right)\right]\left[F_2(\theta_2|\theta_1) - F_2\left(\theta_2|\theta_1'\right)\right] \frac{\partial F_3(\theta_3|\theta_2)}{\partial \theta_2} d\theta_2 d\theta_3 \\ &+ \left(\beta - \beta^2\right)\delta^2 \int_{\underline{\theta}}^{\overline{\theta}} \left[V_3(\underline{\theta}|\theta_1,\theta_2) - V_3\left(\underline{\theta}|\theta_1',\theta_2\right)\right]\left(f_2(\theta_2|\theta_1) - f_2\left(\theta_2|\theta_1'\right)\right) d\theta_2 \\ &+ \left(\beta - \beta^2\right)\delta^2 \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \left[x_3(\theta_3|\theta_1,\theta_2) - x_3\left(\theta_3|\theta_1',\theta_2\right)\right]\left[1 - F_3(\theta_3|\theta_2)\right]\left(f_2(\theta_2|\theta_1) - f_2\left(\theta_2|\theta_1'\right)\right) d\theta_2 d\theta_3 \end{aligned}$$

■

Proof of Lemma 7. I substitute the envelope conditions for all three periods from Lemma 9, 10 and 11 into the seller's payoff function, and the seller's surplus can be written as

$$\begin{aligned}
E\pi &= W\left(x_1,x_2,x_3,V_1\left(\underline{\theta}\right),V_2\left(\underline{\theta}|\theta_1\right),V_3\left(\underline{\theta}|\theta_1,\theta_2\right)\right) \\
&= E_{\theta_1}\left[t_1-\frac{1}{2}x_1^2\right]+\delta E_{\theta_1,\theta_2}\left[t_2-\frac{1}{2}x_2^2\right]+\delta^2 E_{\theta_1,\theta_2,\theta_3}\left[t_3-\frac{1}{2}x_3^2\right] \\
&= \int_{\underline{\theta}}^{\overline{\theta}}\left[\theta_1 x_1-\frac{1}{2}x_1^2-V_1\left(\theta_1\right)\right]f_1\left(\theta_1\right)d\theta_1+\delta\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}}\left[\beta\theta_2 x_2-\frac{1}{2}x_2^2+(1-\beta)t_2\right]f_2\left(\theta_2|\theta_1\right)f_1\left(\theta_1\right)d\theta_2 d\theta_1 \\
&\quad +\delta^2\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}}\left[\beta\theta_3 x_3-\frac{1}{2}x_3^2+(1-\beta)t_3\right]f_3\left(\theta_3|\theta_2\right)f_2\left(\theta_2|\theta_1\right)f_1\left(\theta_1\right)d\theta_3 d\theta_2 d\theta_1 \\
&= \int_{\underline{\theta}}^{\overline{\theta}}\left[\theta_1 x_1-\frac{1}{2}x_1^2-V_1\left(\theta_1\right)\right]f_1\left(\theta_1\right)d\theta_1+\delta\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}}\left[\theta_2 x_2-\frac{1}{2}x_2^2-(1-\beta)V_2\left(\theta_2|\theta_1\right)\right]f_2\left(\theta_2|\theta_1\right)f_1\left(\theta_1\right)d\theta_2 d\theta_1 \\
&\quad +\delta^2\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}}\left[\theta_3 x_3-\frac{1}{2}x_3^2-(1-\beta)^2 V_3\left(\theta_3|\theta_1,\theta_2\right)\right]f_3\left(\theta_3|\theta_2\right)f_2\left(\theta_2|\theta_1\right)f_1\left(\theta_1\right)d\theta_3 d\theta_2 d\theta_1 \\
&= \int_{\underline{\theta}}^{\overline{\theta}}J_1\left(\theta_1\right)f_1\left(\theta_1\right)d\theta_1+\delta\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}}J_2\left(\theta_1,\theta_2\right)f_2\left(\theta_2|\theta_1\right)f_1\left(\theta_1\right)d\theta_2 d\theta_1 \\
&\quad +\delta^2\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}}J_3\left(\theta_1,\theta_2,\theta_3\right)f_3\left(\theta_3|\theta_2\right)f_2\left(\theta_2|\theta_1\right)f_1\left(\theta_1\right)d\theta_3 d\theta_2 d\theta_1-R\left(\underline{\theta}\right)
\end{aligned}$$

where

$$\begin{aligned}
J_1\left(\theta_1\right) &= \theta_1 x_1-\frac{1}{2}x_1^2-\frac{1-F_1\left(\theta_1\right)}{f_1\left(\theta_1\right)}x_1 \\
J_2\left(\theta_1,\theta_2\right) &= \theta_2 x_2-\frac{1}{2}x_2^2-\beta\left(-\frac{\partial F_2\left(\theta_2|\theta_1\right)/\partial\theta_1}{f_2\left(\theta_2|\theta_1\right)}\right)\frac{1-F_1\left(\theta_1\right)}{f_1\left(\theta_1\right)}x_2-(1-\beta)\frac{1-F_2\left(\theta_2|\theta_1\right)}{f_2\left(\theta_2|\theta_1\right)}x_2 \\
J_3\left(\theta_1,\theta_2,\theta_3\right) &= \theta_3 x_3-\frac{1}{2}x_3^2-\beta^2\frac{\partial F_3\left(\theta_3|\theta_2\right)/\partial\theta_2}{f_3\left(\theta_3|\theta_2\right)}\frac{\partial F_2\left(\theta_2|\theta_1\right)/\partial\theta_1}{f_2\left(\theta_2|\theta_1\right)}\frac{1-F_1\left(\theta_1\right)}{f_1\left(\theta_1\right)}x_3 \\
&\quad -\left(\beta-\beta^2\right)\frac{\partial f_2\left(\theta_2|\theta_1\right)/\partial\theta_1}{f_2\left(\theta_2|\theta_1\right)}\frac{1-F_3\left(\theta_3|\theta_2\right)}{f_3\left(\theta_3|\theta_2\right)}\frac{1-F_1\left(\theta_1\right)}{f_1\left(\theta_1\right)}x_3 \\
&\quad -\left(\beta-\beta^2\right)\left(-\frac{\partial F_3\left(\theta_3|\theta_2\right)/\partial\theta_2}{f_3\left(\theta_3|\theta_2\right)}\right)\frac{1-F_2\left(\theta_2|\theta_1\right)}{f_2\left(\theta_2|\theta_1\right)}x_3-(1-\beta)^2\frac{1-F_3\left(\theta_3|\theta_2\right)}{f_3\left(\theta_3|\theta_2\right)}x_3 \\
R\left(\underline{\theta}\right) &= \beta(1-\beta)\delta^2\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}}V_3\left(\underline{\theta}|\theta_1,\theta_2\right)\frac{\partial f_2\left(\theta_2|\theta_1\right)}{\partial\theta_1}\left(1-F_1\left(\theta_1\right)\right)d\theta_2 d\theta_1+V_1\left(\underline{\theta}\right) \\
&\quad +(1-\beta)\delta\int_{\underline{\theta}}^{\overline{\theta}}V_2\left(\underline{\theta}|\theta_1\right)f_1\left(\theta_1\right)d\theta_1 \\
&\quad +(1-\beta)^2\delta^2\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}}V_3\left(\underline{\theta}|\theta_1,\theta_2\right)f_2\left(\theta_2|\theta_1\right)f_1\left(\theta_1\right)d\theta_2 d\theta_1
\end{aligned}$$

■

Proof of Proposition 17. In period 2,

$$\begin{aligned}
V_2\left(\theta_2',\theta_2|\theta_1\right) &= \theta_2 x_2\left(\theta_2'|\theta_1\right) - t_2\left(\theta_2'|\theta_1\right) + \beta\delta\int_{\underline{\theta}}^{\overline{\theta}}\left[\theta_3 x_3\left(\theta_3|\theta_1,\theta_2'\right) - t_3\left(\theta_3|\theta_1,\theta_2'\right)\right] f_3\left(\theta_3|\theta_2\right) d\theta_3 \\
\frac{\partial V_2\left(\theta_2',\theta_2|\theta_1\right)}{\partial\theta_2} &= x_2\left(\theta_2'|\theta_1\right) - \beta\delta\int_{\underline{\theta}}^{\overline{\theta}} x_3\left(\theta_3|\theta_1,\theta_2'\right)\frac{\partial F_3\left(\theta_3|\theta_2\right)}{\partial\theta_2}d\theta_3 \\
V_2'\left(\theta_2|\theta_1\right) &= x_2\left(\theta_2|\theta_1\right) - \beta\delta\int_{\underline{\theta}}^{\overline{\theta}} x_3\left(\theta_3|\theta_1,\theta_2\right)\frac{\partial F_3\left(\theta_3|\theta_2\right)}{\partial\theta_2}d\theta_3
\end{aligned}$$

I further derive that

$$\frac{\partial V_2\left(\theta_2',\theta_2|\theta_1\right)}{\partial\theta_2} - V_2'\left(\theta_2|\theta_1\right) = x_2\left(\theta_2'|\theta_1\right) - x_2\left(\theta_2|\theta_1\right) - \beta\delta\int_{\underline{\theta}}^{\overline{\theta}}\left(x_3\left(\theta_3|\theta_1,\theta_2'\right) - x_3\left(\theta_3|\theta_1,\theta_2\right)\right)\frac{\partial F_3\left(\theta_3|\theta_2\right)}{\partial\theta_2}d\theta_3$$

The FSD condition indicates that $\frac{\partial F_3(\theta_3|\theta_2)}{\partial\theta_2} \leq 0$. Given any $\theta_2 \geq \theta_2'$, as both $x_2^*\left(\theta_1,\theta_2\right)$ and $x_3^*\left(\theta_1,\theta_2,\theta_3\right)$ are non-decreasing in θ_2, I have $\frac{\partial V_2\left(\theta_2',\theta_2|\theta_1\right)}{\partial\theta_2} - V_2'\left(\theta_2|\theta_1\right) \leq 0$ for any θ_1. I can write

$$\begin{aligned}
V_2^*\left(\theta_2',\theta_2|\theta_1\right) &= V_2^*\left(\theta_2'|\theta_1\right) + \int_{\theta_2'}^{\theta_2}\frac{\partial V_2\left(\theta_2',s|\theta_1\right)}{\partial\theta_2}ds \\
V_2^*\left(\theta_2|\theta_1\right) &= V_2^*\left(\theta_2'|\theta_1\right) + \int_{\theta_2'}^{\theta_2} V_2'\left(s|\theta_1\right)ds \geq V_2^*\left(\theta_2',\theta_2|\theta_1\right)
\end{aligned}$$

If $\theta_2 < \theta_2'$, I have $\frac{\partial V_2\left(\theta_2',\theta_2|\theta_1\right)}{\partial\theta_2} - V_2'\left(\theta_2|\theta_1\right) \geq 0$. By integration, I also have $V_2^*\left(\theta_2|\theta_1\right) \geq V_2^*\left(\theta_2',\theta_2|\theta_1\right)$. In addition,

$$\begin{aligned}
V_1'\left(\theta_1\right) &= x_1\left(\theta_1\right) - \beta\delta\int_{\underline{\theta}}^{\overline{\theta}} x_2\left(\theta_2|\theta_1\right)\frac{\partial F_2\left(\theta_2|\theta_1\right)}{\partial\theta_1}d\theta_2 \\
&\quad + \beta^2\delta^2\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} x_3\left(\theta_3|\theta_1,\theta_2\right)\frac{\partial F_3\left(\theta_3|\theta_2\right)}{\partial\theta_2}\frac{\partial F_2\left(\theta_2|\theta_1\right)}{\partial\theta_1}d\theta_2 d\theta_3 \\
&\quad + \left(\beta-\beta^2\right)\delta^2\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} x_3\left(\theta_3|\theta_1,\theta_2\right)\left[1-F_3\left(\theta_3|\theta_2\right)\right]\frac{\partial f_2\left(\theta_2|\theta_1\right)}{\partial\theta_1}d\theta_2 d\theta_3 \\
\frac{\partial V\left(\theta_1',\theta_1\right)}{\partial\theta_1} &= x_1\left(\theta_1'\right) - \beta\delta\int_{\underline{\theta}}^{\overline{\theta}} x_2\left(\theta_2|\theta_1'\right)\frac{\partial F_2\left(\theta_2|\theta_1\right)}{\partial\theta_1}d\theta_2 \\
&\quad + \beta^2\delta^2\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} x_3\left(\theta_3|\theta_1',\theta_2\right)\frac{\partial F_3\left(\theta_3|\theta_2\right)}{\partial\theta_2}\frac{\partial F_2\left(\theta_2|\theta_1\right)}{\partial\theta_1}d\theta_2 d\theta_3 \\
&\quad + \left(\beta-\beta^2\right)\delta^2\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} x_3\left(\theta_3|\theta_1',\theta_2\right)\left[1-F_3\left(\theta_3 \mid \theta_2\right)\right]\frac{\partial f_2\left(\theta_2|\theta_1\right)}{\partial\theta_1}d\theta_2 d\theta_3
\end{aligned}$$

Given any $\theta_1 \geq \theta_1'$, I have $x_1(\theta_1') \leq x_1(\theta_1)$, $x_2(\theta_2|\theta_1') \leq x_2(\theta_2|\theta_1)$ and $x_3(\theta_3|\theta_1', \theta_2) \leq x_3(\theta_3|\theta_1, \theta_2)$. Under the assumptions made in Section 3.4.1, I have $\beta \frac{\partial F_3(\theta_3|\theta_2)/\partial\theta_2}{f_3(\theta_3|\theta_2)} \frac{\partial F_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)} + (1-\beta) \frac{\partial f_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)} \frac{1-F_3(\theta_3|\theta_2)}{f_3(\theta_3|\theta_2)} \geq 0$, $\frac{\partial F_2(\theta_2|\theta_1)}{\partial\theta_1} \leq 0$ and $\frac{\partial F_3(\theta_3|\theta_2)}{\partial\theta_2} \leq 0$.

$$\begin{aligned}
V_1'(\theta_1) - \frac{\partial V(\theta_1', \theta_1)}{\partial\theta_1} &= x_1(\theta_1) - x_1(\theta_1') - \beta\delta \int_{\underline{\theta}}^{\overline{\theta}} (x_2(\theta_2|\theta_1) - x_2(\theta_2|\theta_1')) \frac{\partial F_2(\theta_2|\theta_1)}{\partial\theta_1} d\theta_2 \\
&\quad + \beta^2\delta^2 \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} (x_3(\theta_3|\theta_1, \theta_2) - x_3(\theta_3|\theta_1', \theta_2)) \frac{\partial F_3(\theta_3|\theta_2)}{\partial\theta_2} \frac{\partial F_2(\theta_2|\theta_1)}{\partial\theta_1} d\theta_2 d\theta_3 \\
&\quad + \beta\delta^2 \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} (x_3(\theta_3|\theta_1, \theta_2) - x_3(\theta_3|\theta_1', \theta_2)) [1 - F_3(\theta_3|\theta_2)] \frac{\partial f_2(\theta_2|\theta_1)}{\partial\theta_1} d\theta_2 d\theta_3 \\
&\quad - \beta^2\delta^2 \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} (x_3(\theta_3|\theta_1, \theta_2) - x_3(\theta_3|\theta_1', \theta_2)) [1 - F_3(\theta_3|\theta_2)] \frac{\partial f_2(\theta_2|\theta_1)}{\partial\theta_1} d\theta_2 d\theta_3 \\
&\geq 0
\end{aligned}$$

I further concludes that $V_1(\theta_1) \geq V(\theta_1', \theta_1)$. If $\theta_1 < \theta_1'$, then $V_1'(\theta_1) \leq \frac{\partial V(\theta_1', \theta_1)}{\partial\theta_1}$ and $V_1(\theta_1) \geq V(\theta_1', \theta_1)$. As $V_1'(\theta_1) \geq 0$ and $V_2'(\theta_2|\theta_1) \geq 0$, the IR constraints in the first two periods are satisfied as well. According to the conditions in Proposition 17, the buyer has an incentive to report his true type in all periods; the above mechanism is indeed optimal.

■

Proof of Proposition 18. The seller's profit maximization problem is essentially:

$$\max_{\{t_s, x_s\}} \mu \left[t_1(\overline{\theta}|h_1) - \frac{1}{2} x_1^2(\overline{\theta}|h_1) + \delta E(\pi_2(\theta_2|h_1, \overline{\theta})) \right] + (1-\mu) \left(t_1(\underline{\theta}|h_1) - \frac{1}{2} x_1^2(\underline{\theta}|h_1) + \delta E(\pi_2(\theta_2|h_1, \underline{\theta})) \right)$$

subject to the IC and IR constraints for all θ_s and θ_s', where

$$\begin{aligned}
\text{ICs} \quad : \quad & \theta_s x_s(\theta_s|h_s) - t_s(\theta_s|h_s) + \beta\delta E\left(\widehat{U}_{s+1}(\theta_{s+1}|h_s, \theta_s)|\theta_s\right) \\
\geq \quad & \theta_s x_s(\theta_s'|h_s) - t_s(\theta_s'|h_s) + \beta\delta E\left(\widehat{U}_{s+1}(\theta_{s+1}|h_s, \theta_s')|\theta_s\right) \\
\text{IRs} \quad : \quad & \theta_s x_s(\theta_s|h_s) - t_s(\theta_s|h_s) + \beta\delta E\left(\widehat{U}_{s+1}(\theta_{s+1}|h_s, \theta_s)|\theta_s\right) \geq 0
\end{aligned}$$

and

$$\pi_2(\theta_2|h_1,\theta_1) = t_2(\theta_2|h_1,\theta_1) - \frac{1}{2}x_2^2(\theta_2|h_1,\theta_1) + \delta E(\pi_3(\theta_3|h_2,\theta_2)|\theta_2)$$

$$\widehat{U}_{s+1}(\theta_{s+1}|h_s,\theta_s) = \theta_{s+1}x_{s+1}(\theta_{s+1}|h_s,\theta_s) - t_{s+1}(\theta_{s+1}|h_s,\theta_s) + \delta E\left(\widehat{U}_{s+1}(\theta_{s+2}|h_{s+1},\theta_{s+1})|\theta_{s+1}\right)$$

Using the $IR_s(\underline{\theta})$ and $IC_s(\overline{\theta})$, I characterize the utility of type $\overline{\theta}$ in period s.

$$\begin{aligned}
U_s(\overline{\theta}|h_s) &= \overline{\theta}x_s(\overline{\theta}|h_s) - t_s(\overline{\theta}|h_s) + \beta\delta E\left(\widehat{U}_{s+1}(\theta_{s+1}|h_s,\overline{\theta})|\overline{\theta}\right) \\
&= \overline{\theta}x_s(\underline{\theta}|h_s) - t_1(\underline{\theta}|h_s) + \beta\delta E\left(\widehat{U}_{s+1}(\theta_{s+1}|h_s,\underline{\theta})|\overline{\theta}\right) \\
&= \Delta\theta\sum_{j=0}^{T-s}\beta^j\delta^j(\alpha+\gamma-1)^j x_{t+j}(\underline{\theta}|h_s,\underline{\theta},\dots,\underline{\theta}) \\
&\quad +\beta(1-\beta)\delta^2(\alpha+\gamma-1)\left[E\left(\widehat{U}_{s+2}(\theta_{s+2}|h_s,\underline{\theta},\overline{\theta})|\overline{\theta}\right) - E\left(\widehat{U}_{s+2}(\theta_{s+2}|h_s,\underline{\theta},\underline{\theta})|\underline{\theta}\right)\right] + \dots \\
&\quad +\beta^{T-2}(1-\beta)\delta^{T-1}(\alpha+\gamma-1)^{T-2}E\left(\widehat{U}_T(\theta_T|h_s,\underline{\theta},\dots,\underline{\theta},\overline{\theta})|\overline{\theta}\right) \\
&\quad -\beta^{T-2}(1-\beta)\delta^{T-1}(\alpha+\gamma-1)^{T-2}E\left(\widehat{U}_T(\theta_T|h_s,\underline{\theta},\dots,\underline{\theta},\underline{\theta})|\underline{\theta}\right)
\end{aligned}$$

From $IR_s\left(\underline{\theta}\right)$, the low type receives a payoff of zero at any time, and the expected profit of the seller can be transformed into

$$\begin{aligned}
E\pi_1\left(\theta_1|h_1\right) &= \mu\left[t_1\left(\overline{\theta}\right)-\frac{1}{2}x_1^2\left(\overline{\theta}\right)\right]+\left(1-\mu\right)\left(t_1\left(\underline{\theta}\right)-\frac{1}{2}x_1^2\left(\underline{\theta}\right)\right)\\
&\quad -\mu\Delta\theta\sum_{j=0}^{T-1}\beta^j\delta^j\left(\alpha+\gamma-1\right)^j x_{1+j}\left(\underline{\theta}|\underline{\theta},\ldots,\underline{\theta}\right)\\
&\quad -\mu\alpha\left(1-\beta\right)\delta\Delta\theta\sum_{j=0}^{T-2}\beta^j\delta^j\left(\alpha+\gamma-1\right)^j x_{2+j}\left(\underline{\theta}|\overline{\theta},\underline{\theta},\ldots,\underline{\theta}\right)\\
&\quad -\left(1-\mu\right)\left(1-\gamma\right)\left(1-\beta\right)\delta\Delta\theta\sum_{j=0}^{T-2}\beta^j\delta^j\left(\alpha+\gamma-1\right)^j x_{2+j}\left(\underline{\theta}|\underline{\theta},\underline{\theta},\ldots,\underline{\theta}\right)-\ldots\\
&\quad -\mu\alpha^{T-1}\left(1-\beta\right)^{T-1}\delta^{T-1}\Delta\theta x_T\left(\underline{\theta}|\overline{\theta},\ldots,\overline{\theta},\underline{\theta}\right)-\ldots\\
&\quad -\left(1-\mu\right)\left(1-\gamma\right)\gamma^{T-2}\delta^{T-1}\Delta\theta x_T\left(\underline{\theta}|\underline{\theta},\ldots,\underline{\theta},\underline{\theta}\right)\\
&\quad -\mu\alpha\left(\alpha+\gamma-1\right)\beta\left(1-\beta\right)\delta^2\Delta\theta\sum_{j=0}^{T-3}\beta^j\delta^j\left(\alpha+\gamma-1\right)^j x_{3+j}\left(\underline{\theta}|\underline{\theta},\overline{\theta},\underline{\theta},\ldots,\underline{\theta}\right)\\
&\quad +\mu\left(1-\gamma\right)\left(\alpha+\gamma-1\right)\beta\left(1-\beta\right)\delta^2\Delta\theta\sum_{j=0}^{T-3}\beta^j\delta^j\left(\alpha+\gamma-1\right)^j x_{3+j}\left(\underline{\theta}|\underline{\theta},\underline{\theta},\underline{\theta},\ldots,\underline{\theta}\right)-\ldots\\
&\quad -\mu\alpha\left(\alpha+\gamma-1\right)\beta\left(1-\beta\right)\delta^{T-1}\left(\beta\left(\alpha+\gamma-1\right)-\left(1-\gamma\right)\left(1-\beta\right)\right)^{T-3}\Delta\theta x_T\left(\underline{\theta}|\underline{\theta},\ldots,\underline{\theta},\overline{\theta},\underline{\theta}\right)\\
&\quad +\mu\left(1-\gamma\right)\left(\alpha+\gamma-1\right)\beta\left(1-\beta\right)\delta^{T-1}\left(\beta\left(\alpha+\gamma-1\right)+\gamma(1-\beta)\right)^{T-3}\Delta\theta x_T\left(\underline{\theta}|\underline{\theta},\ldots,\underline{\theta},\underline{\theta},\underline{\theta}\right)
\end{aligned}$$

The first-order optimal allocations can be summarized as,

$$\begin{aligned}
x_s\left(\overline{\theta}|h_s\right) &= \overline{\theta}\\
x_s\left(\underline{\theta}|h_s\right) &= \underline{\theta}-k_1\Delta\theta\sum_{j=0}^{N-1}\beta^{j+1}\left(1-\beta\right)^{s-j-2}\sum_{\substack{\sigma_1,\ldots\sigma_j\in\{2,\ldots,N\}\\ \sigma_1<\sigma_2<\cdots<\sigma_j}}p_{1,\sigma_1}p_{\sigma_1,\sigma_2}\cdots p_{\sigma_{j-1},\sigma_j}p_{\sigma_j,N+1}\\
&\quad -k_2\Delta\theta\sum_{j=0}^{N-2}\beta^{j+1}\left(1-\beta\right)^{s-j-2}\sum_{\substack{\sigma_1,\ldots\sigma_j\in\{3,\ldots,N\}\\ \sigma_1<\sigma_2<\cdots<\sigma_j}}p_{2,\sigma_1}p_{\sigma_1,\sigma_2}\cdots p_{\sigma_{j-1},\sigma_j}p_{\sigma_j,N+1}\\
&\quad -\cdots-k_N\Delta\theta\beta\left(1-\beta\right)^{s-2}p_{N,N+1}-k_{N+1}\left(1-\beta\right)^{s-1}\\
&= \underline{\theta}-\Delta\theta\sum_{n=1}^{N}k_n\sum_{j=0}^{N-n}\beta^{j+1}\left(1-\beta\right)^{s-j-2}\sum_{\substack{\sigma_1,\ldots\sigma_j\in\{n+1,\ldots,N\}\\ \sigma_1<\sigma_2<\cdots<\sigma_j}}p_{n,\sigma_1}p_{\sigma_1,\sigma_2}\cdots p_{\sigma_{j-1},\sigma_j}p_{\sigma_j,N+1}\\
&\quad -k_{N+1}\Delta\theta\left(1-\beta\right)^{s-1}\\
&= \underline{\theta}-\sum_{n=1}^{N}k_nA_n-k_{N+1}\left(1-\beta\right)^{s-1}
\end{aligned}$$

where

$$A_n = \sum_{j=0}^{N-n} \beta^{j+1} (1-\beta)^{s-j-2} \sum_{\substack{\sigma_1,\ldots\sigma_j \in \{n+1,\ldots,N\} \\ \sigma_1<\sigma_2<\cdots<\sigma_j}} p_{n,\sigma_1} p_{\sigma_1,\sigma_2} \cdots p_{\sigma_{j-1},\sigma_j} p_{\sigma_j,N+1}$$

$$k_1 = \begin{cases} \frac{\alpha}{1-\alpha}, & \text{if } \theta_1 = \overline{\theta} \\ \frac{\mu}{1-\mu}, & \text{if } \theta_1 = \underline{\theta} \end{cases}$$

$$k_n = \begin{cases} \frac{\alpha}{1-\alpha}, & \text{if } \theta_{m_n-1} = \overline{\theta} \\ \frac{1-\gamma}{\gamma},, & \text{if } \theta_{m_n-1} = \underline{\theta} \end{cases}$$

$$p_{n,n+i} = \begin{cases} \frac{(\alpha+\gamma-1)\alpha}{(1-\gamma)(1-\alpha)}, & \text{if } \theta_{m_n+1} = \overline{\theta}, i=1 \\ \frac{\alpha+\gamma-1}{\gamma}, & \text{if } \theta_{m_n+1} = \underline{\theta}, i=1 \\ \frac{\alpha+\gamma-1}{1-\gamma} k_{n+i}, & \text{if } \theta_{m_n+1} = \overline{\theta}, i>1 \\ -\frac{\alpha+\gamma-1}{\gamma} k_{n+i}, & \text{if } \theta_{m_n+1} = \underline{\theta}, i>1 \end{cases}$$

Similar to the sufficiency conditions in the three-period model, all the additional constraints can be transformed into functions of β. I can show that when $\beta = 1$, the $IR_s\left(\overline{\theta}\right)$ and $IC_s\left(\underline{\theta}\right)$ hold with strict inequalities under the first-order optimal allocations.

$$\begin{aligned} IR_s\left(\overline{\theta}\right) \quad &: \quad \overline{\theta} x_s\left(\overline{\theta}|h_s\right) - t_s\left(\overline{\theta}|h_s\right) + \delta E\left(\widehat{U}_{s+1}\left(\theta_{s+1}|h_s,\overline{\theta}\right)\right) \\ &= \quad \Delta\theta \sum_{j=0}^{T-s} \delta^j (\alpha+\gamma-1)^j x_{s+j}\left(\underline{\theta}|h_s,\underline{\theta},\ldots,\underline{\theta}\right) > 0 \\ IC_s\left(\underline{\theta}\right) \quad &: \quad \underline{\theta} x_s\left(\underline{\theta}|h_s\right) - t_s\left(\underline{\theta}|h_s\right) + \delta E\left(\widehat{U}_{s+1}\left(\theta_{s+1}|h_s,\underline{\theta}\right)\right) - \left[\underline{\theta} x_s\left(\overline{\theta}|h_s\right) - t_s\left(\overline{\theta}|h_s\right) + \delta E\left(\widehat{U}_{s+1}\left(\theta_{s+1}|h_s,\overline{\theta}\right)\right)\right] \\ &= \quad \Delta\theta\left(x_s\left(\overline{\theta}|h_s\right) - x_s\left(\underline{\theta}|h_s\right)\right) + \Delta\theta\delta(\alpha+\gamma-1)\left(x_{s+1}\left(\underline{\theta}|h_s,\overline{\theta}\right) - x_{s+1}\left(\underline{\theta}|h_s,\underline{\theta}\right)\right) \\ &\quad + \ldots + \Delta\theta\delta^{T-s}(\alpha+\gamma-1)^{T-s}\left(x_T\left(\underline{\theta}|h_s,\overline{\theta},\underline{\theta},\ldots,\underline{\theta}\right) - x_T\left(\underline{\theta}|h_s,\underline{\theta},\underline{\theta},\ldots,\underline{\theta}\right)\right) > 0 \end{aligned}$$

At $\beta = 1$, $IR_s\left(\overline{\theta}\right)$ and $IC_s\left(\underline{\theta}\right)$ are strictly positive. Furthermore, the allocations for the present-biased agent converge to the allocations for the time consistent agent, as β approaches 1. Since $IR_s\left(\overline{\theta}\right)$ and $IC_s\left(\underline{\theta}\right)$ are continuous functions of β, there exists a $\underline{\beta}_T$ such that if $\beta \geq \underline{\beta}_T$, the allocations presented in Proposition 18 are indeed optimal.

■

Bibliography

[1] Viral V Acharya, Thomas F Cooley, Matthew P Richardson, and Ingo Walter. *Regulating Wall Street: The Dodd-Frank Act and the New Architecture of Global Finance*, volume 608. John Wiley & Sons, 2010.

[2] Gian Luigi Albano and Alessandro Lizzeri. Strategic certification and provision of quality. *International Economic Review*, 42(1):267–283, 2001.

[3] Adam Ashcraft, Paul Goldsmith-pinkham, and James Vickery. Mbs ratings and the mortgage credit boom. Technical report, Federal Reserve Bank of New York, 2010.

[4] David P Baron and David Besanko. Regulation and information in a continuing relationship. *Information Economics and Policy*, 1(3):267–302, 1984.

[5] Marco Battaglini. Long-term contracting with markovian consumers. *American Economic Review*, 95(3):637–658, 2005.

[6] Marco Battaglini and Rohit Lamba. Optimal dynamic contracting: the first-order approach and beyond. *CEPR Discussion Papers*, 2015.

[7] Raphael Boleslavsky and Kyungmin Kim. Bayesian persuasion and moral hazard. *Working paper*, 2017.

[8] Patrick Bolton, Xavier Freixas, and Joel Shapiro. The credit ratings game. *Journal of Finance*, 67(1):85–111, 2012.

[9] Jess Cornaggia and Kimberly J Cornaggia. Estimating the costs of issuer-paid credit ratings. *Review of Financial Studies*, 26(9):2229–2269, 2013.

[10] Pascal Courty and Hao Li. Sequential screening. *Review of Economic Studies*, 67(4):697–717, 2000.

[11] Brendan Daley and Brett Green. Waiting for news in the market for lemons. *Econometrica*, 80(4):1433–1504, 2012.

[12] Jacques De Larosière, Leszek Balcerowicz, Otmar Issing, Rainer Masera, Callum McCarthy, Lars Nyberg, José Pérez, and Onno Ruding. Report of the high-level group on financial supervision in the eu. *European Commission. Brussels*, 2009.

[13] Stefano DellaVigna and Ulrike Malmendier. Contract design and self-control: Theory and evidence. *Quarterly Journal of Economics*, 119(2):353–402, 2004.

[14] Kfir Eliaz and Ran Spiegler. Contracting with diversely naive agents. *Review of Economic Studies*, 73(3):689–714, 2006.

[15] Susanna Esteban and Eiichi Miyagawa. Optimal menu of menus with self-control preferences. *NAJ Economics*, 2006. mimeo.

[16] Susanna Esteban and Eiichi Miyagawa. Temptation, self-control, and competitive nonlinear pricing. *Economics Letters*, 90(3):348–355, 2006.

[17] Erik R Fasten and Dirk Hofmann. Two-sided certification. *Working paper*, 2010.

[18] William Fuchs, Aniko Öry, and Andrzej Skrzypacz. Transparency and distressed sales under asymmetric information. *Theoretical Economics*, 11(3):1103–1144, 2016.

[19] Sanford J Grossman. The informational role of warranties and private disclosure about product quality. *Journal of Law and Economics*, 24(3):461–483, 1981.

[20] Faruk Gul and Wolfgang Pesendorfer. Temptation and self-control. *Econometrica*, 69(6):1403–1435, 2001.

[21] Barney Hartman-Glaser, Tomasz Piskorski, and Alexei Tchistyi. Optimal securitization with moral hazard. *Journal of Financial Economics*, 104(1):186–202, 2012.

[22] Jie Jack He, Jun Qj Qian, and Philip E Strahan. Are all ratings created equal? the impact of issuer size on the pricing of mortgage-backed securities. *Journal of Finance*, 67(6):2097–2137, 2012.

[23] Johannes Hörner and Nicolas Vieille. Public vs. private offers in the market for lemons. *Econometrica*, 77(1):29–69, 2009.

[24] John Xuefeng Jiang, Mary Harris Stanford, and Yuan Xie. Does it matter who pays for bond ratings? historical evidence. *Journal of Financial Economics*, 105(3):607–621, 2012.

[25] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

[26] Anil K Kashyap and Natalia Kovrijnykh. Who should pay for credit ratings and how?

Review of Financial Studies, 29(2):420–456, 2015.

[27] Ayça Kaya and Qingmin Liu. Transparency and price formation. *Theoretical Economics*, 10(2):341–383, 2015.

[28] Kyungmin Kim. Information about sellers past behavior in the market for lemons. *Journal of Economic Theory*, 169:365–399, 2017.

[29] Ilan Kremer and Andrzej Skrzypacz. Dynamic signaling and market breakdown. *Journal of Economic Theory*, 133(1):58–82, 2007.

[30] David Laibson. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 112(2):443–478, 1997.

[31] Alessandro Lizzeri. Information revelation and certification intermediaries. *RAND Journal of Economics*, 30(2):214–231, 1999.

[32] Joseph R Mason and Josh Rosner. Where did the risk go? how misapplied bond ratings cause mortgage backed securities and collateralized debt obligation market disruptions. *Working paper*, 2007.

[33] Jerome Mathis, James McAndrews, and Jean-Charles Rochet. Rating the raters: Are reputation concerns powerful enough to discipline rating agencies? *Journal of Monetary Economics*, 56(5):657–674, 2009.

[34] Atif Mian and Amir Sufi. The consequences of mortgage credit expansion: Evidence from the us mortgage default crisis. *Quarterly Journal of Economics*, 124(4):1449–1496, 2009.

[35] Paul R Milgrom. Good news and bad news: Representation theorems and applications. *Bell Journal of Economics*, 12(2):380–391, 1981.

[36] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[37] Georg Noldeke and Eric Van Damme. Signalling in a dynamic labour market. *Review of Economic Studies*, 57(1):1–23, 1990.

[38] Ted O'Donoghue and Matthew Rabin. Doing it now or later. *American Economic Review*, 89(1):103–124, 1999.

[39] Ted O'Donoghue and Matthew Rabin. Incentives for procrastinators. *Quarterly Journal of Economics*, 114(3):769–816, 1999.

[40] Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: A myerso-

nian approach. *Econometrica*, 82(2):601–653, 2014.

[41] Eloic Peyrache and Lucia Quesada. Intermediaries, credibility and incentives to collude. *Journal of Economics & Management Strategy*, 20(4):1099–1133, 2011.

[42] Edmund S Phelps and Robert A Pollak. On second-best national saving game-equilibrium growth. *Review of Economic Studies*, 35(2):185–199, 1968.

[43] David Rodina. Information design and career concerns. Technical report, Northwestern University, 2016.

[44] David Rodina and John Farragut. Inducing effort through grades. Technical report, Northwestern University, 2016.

[45] Vasiliki Skreta and Laura Veldkamp. Ratings shopping and asset complexity: A theory of ratings inflation. *Journal of Monetary Economics*, 56(5):678–695, 2009.

[46] Konrad Stahl and Roland Strausz. Certification and market transparency. *Review of Economic Studies*, 84(4):1842–1868, 2017.

[47] Roland Strausz. Honest certification and the threat of capture. *International Journal of Industrial Organization*, 23(1):45–62, 2005.

[48] Robert Henry Strotz. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, 23(3):165–180, 1955.

[49] Jeroen M Swinkels. Education signalling with preemptive offers. *Review of Economic Studies*, 66(4):949–970, 1999.

[50] Lawrence J White. Markets: The credit rating agencies. *Journal of Economic Perspectives*, 24(2):211–226, 2010.

[51] Han Xia and Günter Strobl. The issuer-pays rating model and ratings inflation: Evidence from corporate credit ratings. *Working paper*, 2012.