

INFORMATION AND THE PERSISTENCE OF PRIVATE-ORDER CONTRACTUAL ENFORCEMENT INSTITUTIONS: AN EXPERIMENTAL ANALYSIS

TOM WILKENING^{*,†}

May 15, 2015

We study an experimental market where sellers are prone to moral hazard and where existing public-order contract enforcement institutions are only partially effective at preventing sellers from renegeing on their contractual obligations. A private-order contract enforcement institution exists which can mediate trade and fully resolve the moral hazard problem. However, it does so at a significant cost. The existence of the private-order market institution gives rise to two market equilibria — mediated and unmediated — that vary in the use of the private-order institution. Theory suggests that the mediated equilibrium may make public and private market signals uninformative and inhibit learning. In a dynamic experimental environment, we study whether this potential information externality can limit efficient adaptation away from the private-order institution when market conditions improve. Consistent with theory, we find inefficient persistence after the mediated equilibrium is established, but efficient adaptation in other situations. Providing individuals in the mediated equilibrium with the information that was eliminated by the private-order CEI enables them to adapt. These results suggest that information externalities inherent in private-order contract enforcement institutions can have a profound impact on their adaptability and long-run efficiency.

Keywords: Persistence of Institutions, Equilibrium Selection, Information, Experimental Economics.

JEL Classification Codes: C92, D02, D40, D82

*I am grateful to Abhijit Banerjee, Ernst Fehr, and Sergei Izmalkov, for their thoughtful and detailed comments on earlier drafts. I am indebted to Georg Hartmann and Holger Herz for their help with translation and managing the laboratory. Daron Acemoglu, Jim Berry, David Cooper, Nisvan Erkal, Greg Fischer, Simon Loertscher, Nikos Nikiforakis, Chris Smith, Rob Townsend, Patrick Warren, and various seminar participants provided useful comments. All programs were written in Z-Tree (www.iew.unizh.ch/ztree/). I gratefully acknowledge the financial support of the George and Obie Schultz Fund Research Grant, The University of Melbourne Faculty of Business and Economics, and the ARC Discovery Early Career Research Award DE140101014.

[†]Department of Economics, The University of Melbourne, FBE 328, VIC 3010 Australia. E-mail: Tom.Wilkening@unimelb.edu.au

1 Introduction

The ability of individuals to engage in voluntary exchange is a key driver of economic efficiency. In all but the simplest environments, however, such exchange is jeopardized by enforcement problems that arise due to an unavoidable separation between the transfer of payments and the return of agreed upon service. This temporal delay – inherent in financial transactions, labor relationships, agency relationships, and the exchange of experience goods – give rise to what Greif (2000) terms the “fundamental problem of exchange,” whereby the second mover in a relationship has incentive to renege on their contractual obligations (Greif, 2005).

A myriad of contract-enforcement institutions (CEI) have arisen to induce individuals to commit to their contractual obligations. Public-order CEIs – the legal and regulatory rules that are imposed and enforced by the state through sanction – provide a basic level of legal protection on which to base contracting.¹ Balanced by a need to apply across a variety of settings and limited by the capabilities of the state, public-order CEIs are typically not tailored to a particular market or environment, and often lack the nuance necessary to guarantee fully efficient trade. Private-order CEIs, such as stock exchanges, credit rating companies, business associations, accreditation associations, banks, certifiers, and information intermediaries, often form to provide supplemental enforcement capabilities, but at additional costs.

Microeconomics typically takes the presence of well-functioning CEIs as given or assumes that institutions have evolved that maximize social efficiency. However, it is clear from both the historical literature (e.g. Wittfogel, 1957; Yang Li, 2002; Goetzmann and Köll, 2005) and the development literature (Fafchamps, 2004) that private-order CEIs do not always evolve and adapt efficiently. Private-order CEIs do not exist in many markets where they are likely to be beneficial. Further, CEIs which have developed in response to circumstances at one point in time persist even when they later become inefficient (Gerschenkron, 1962; Greif, 2002).

This paper explores one potential reason why private-order CEIs may not adjust optimally to changes in market conditions. We consider an information externality embedded in the services provided by private-order CEIs that may limit the ability for market participants to observe changes in their environment and adjust their institutions efficiently. In solving the fundamental problem of exchange, private-order CEIs eliminate the incentives of sellers who use the service to renege on their promises. As utilization of the private-order CEI increases, however, information about what sellers would do in the counterfactual case where only public-order CEIs support trade is lost. We hypothesize that this information externality can have a profound influence on the persistence of private-order CEIs since market participants cannot observe changes to their environment and coordinate a response.

We consider an environment where sellers can produce high-quality and/or low-quality experi-

¹The taxonomy of public-order CEIs and private order CEIs is discussed in Menger (1883) and Greif (2005). Institutional Economics makes an additional distinction between organic and designed institutions. As discussed by Greif (2000), designed institutions appear to be most important for large and dynamic economies that can gain from impersonal exchange. In the experiments that follow, we are interested in understanding trade in a setting where identity is anonymous and thus consider CEIs that would typically be classified as being designed.

ence goods that the buyer cannot differentiate between at the point of sale. Sellers are heterogeneous in their production cost and, without the use of the private-order CEI, face different incentives to renege on their contractual obligations based on their production costs and the expected fees charged by a public-order CEI, which punishes sellers who are detected exchanging low-quality units. Buyers and sellers in the market have access to a private-order CEI who can perfectly distinguish between high and low-quality goods and guarantee quality. However, as this certification firm is an information intermediary with some market power, it does so at a substantial cost.

We first show that when the public-order CEI is only partially effective at resolving the moral hazard problem and induces only low-cost sellers to produce high-quality units, two different types of rational expectations equilibria may exist, which vary in the use of the private-order CEI and differ in terms of efficiency and in the information embedded in publicly observable market signals and the outcomes of private trade. We refer to these equilibria as unmediated and mediated reflecting the use of the private-order CEI in mediating trades.

In the unmediated equilibrium, the cost of certifying a unit exceeds the difference in price between certified and uncertified goods. As a result, no party uses the certification firm. High- and low-quality products are traded within a single market and moderate-cost sellers renege on their contractual obligations and produce low-quality units. As there is uncertainty regarding the quality of the product and only low-cost sellers produce high-quality units, the market price contains information about the expected proportion of low-cost sellers. An exogenous decrease in the number of low-cost sellers therefore leads to an observable decrease in the market price. This decline in price can lead to an arbitrage opportunity for low-cost and medium-cost sellers by adopting certification and provides a natural channel by which market participants may endogenously adopt the private-order CEI.

In the mediated equilibrium the private-order CEI is utilized by both low-cost and medium-cost sellers. Consequently, their actions no longer reveal their types and market prices provide no new information. Hence, there is no direct way for individuals to share the information that is necessary to adapt away from utilizing the private-order CEI.

We demonstrate that – relative to the unmediated equilibrium – the mediated equilibrium is more informative about the quality of the good at the point of sale, but less informative about the underlying distribution of seller types. This may hinder learning when market conditions change. Thus the use of private-order CEIs can have a profound impact on the aggregation of information in environments where there is uncertainty and where the market environment varies over time.

Given the strong theoretical difference in the informativeness of markets where the private-order CEI is used and where it is not used, a natural conjecture is that this asymmetry in information may influence the adoption and persistence of the private-order CEI. To explore this idea, laboratory experiments are used to study equilibrium selection and the persistence of the private-order CEI in an environment where the underlying population of sellers changes over time. Subjects initially trade in one of two environments – Safe and Hazardous – which vary in the composition of sellers in the market. In the Safe environment, the proportion of low-cost sellers in the market is large,

thus favouring the formation of an unmediated equilibrium. In the Hazardous environment, low-cost sellers are replaced with medium-cost sellers who will renege on their contractual obligations without the private-order CEI, leading to substantial amounts of moral hazard and a predisposition toward a mediated equilibrium. Subjects who begin in the Safe environment are switched to the Hazardous environment midway through the experiment. Likewise subjects who begin in the Hazardous environment are switched to the Safe environment.

Consistent with the model's predictions, individuals who begin in the Safe environment establish an unmediated equilibrium and then adapt to the mediated equilibrium in response to a change in the underlying environment. Subjects who begin in the Hazardous environment form a mediated equilibrium and remain in this equilibrium when the environment is changed to Safe. Looking at individual decision making, we find evidence of learning in the unmediated equilibrium both through an individual's personal purchase experiences and through his or her observation of other buyers' trades. By contrast, there is little evidence of learning in markets where the mediated equilibrium has formed.

While the data from our experiment supports the hypothesis that information externalities are impacting the usage pattern of the private-order CEI, an alternative interpretation of our results is that the persistence of the private-order CEI is simply due to the fact that there is always a rational expectations equilibrium where the private-order CEI is used. To distinguish between information externalities and the inability of individuals to coordinate away from Pareto dominated market equilibrium, we run additional sessions where we provide information about the proportion of low-cost sellers in the environment by showing subjects the results of a bonus game where they are asked to guess the proportion of low-cost sellers.

Subjects in this information treatment begin with the Hazardous environment and quickly form the mediated equilibrium where the private-order CEI is used. Consistent with the information story, however, a small subset of buyers and low-cost sellers eventually trade without the intermediary when the environment is changed to Safe. A partially-mediated equilibrium forms in all sessions that does not differ significantly from the equilibria observed by groups that began in the Safe environment. Our results thus suggest that at least a part of the persistence in the private-order CEI is due to the information externality that is inherent in its use.

We also run an additional extreme test of the information channel by conducting additional sessions that begin in a variant of the Hazardous environment and slowly switch to a Fully Safe environment where all sellers are low-cost. In this Fully Safe environment, the public-order CEI fully eliminates moral hazard and the private-order CEI serves no purpose. In sessions without information, 66.1% of trades utilized the private-order CEI in the last 6 periods. Further, 2 of the 6 sessions remain in the mediated equilibrium to the end of the session and utilize the private-order CEI for almost all trades. In sessions where information is provided through the bonus game, only 20.6% of trades continued to use the private-order technology and all sessions converge to one of the unmediated equilibria.

Taken together, these results provide evidence that adopted institutions can have a substantial

impact on the ability of individuals to learn due to an information externality that is inherent in their use. This information externality opens a channel by which long-term inefficient institutions can persist even under conditions where market forces efficiently select the optimal market institutions in the short run.

This paper is related to the “institutions-as-equilibrium” literature, which seeks to understand how the interactions of agents might lead to institutions being adopted and how adoption of institutions leads individuals to act in a manner that perpetuates the institution. The institutions-as-equilibrium literature (e.g. Schotter, 1981; Greif, 1994, 1998; Calvert, 1995; Aoki, 2001; Dixit, 2004; Kingston and Caballero, 2011; Greif and Kingston, 2011) seeks to understand how institutions might be self-enforcing either by (1) confirming beliefs about the types or actions of others through observed outcomes, (2) inter-temporally transmitting equilibrium beliefs to newcomers, or (3) reinforcing actions through coordination (Greif, 2006).

We show that private-order CEI may be self-enforcing in all three ways: The information externality embedded in the private-order CEI eliminates information that individuals in the economy could obtain through private trade and garbles signals that might be used by newcomers to learn the state of nature. Adoption of the private-order CEI also makes trading without the institution more risky making coordination away from the institution more difficult.

An alternative explanation of institutional persistence, studied in other institutional contexts by North (1981), Brainard and Verdier (1994), Coate and Morris (1999), Acemoglu and Robinson (2008), and Acemoglu and Robinson (2006) is that private-order CEIs may exploit their political and market power to maintain their market influence. This channel is supported by empirical evidence provided by Hoffman, Postel-Vinay, and Rosenthal (2000) that suggests private-order CEIs form from an initial informational advantage which enables them to gain from exchanging this information with others. As information is typically associated with high fixed costs and low variable costs, CEIs inherently have the features of natural monopolies.

By running experiments, we are able to control for market power by introducing a private-order CEI with exogenously set fees that do not adjust to market conditions. Experiments also allow for the study of equilibrium selection in a replicable environment where there is exogenous control of supply, demand, information, and the number of equilibria. This allows for an experimental study of market dynamics with minimal assumptions about the strategies of agents.

The institutions-as-equilibrium literature builds on the conventions literature (e.g. Foster and Young, 1990; Young, 1993; Kandori, 1992) where there has been a long tradition of using experimental economics to understand equilibrium selection and learning. Closest to the experiments studied in this paper is the research program initiated by David Cooper and John Kagel who study the entry limit pricing game of Milgrom and Roberts (1982). In Cooper, Garvin, and Kagel (1997a,b), a simplified version of the entry price game is introduced where either a separating equilibrium or both a pooling and separating equilibrium exists. It is found that individuals without experience begin near the myopic optima, attempt to pool, and then separate if no unmediated equilibrium exists. Individuals with experience are heavily influenced by the history of their past play. While

adaptive learning based on fictitious play can rationalize the data in Cooper, Garvin, and Kagel (1997a,b), subsequent work in Cooper and Kagel (2003, 2005, 2008) suggests that some individuals in the environment are sophisticated and can predict the response of myopic agents to changes in the environment. The proportion of sophisticated individuals grows with experience and context.

Our paper contributes to the literature by showing how information externalities embedded in a private-order CEI can influence equilibrium selection. To our knowledge, these information externalities are not analysed in prior experiments and provide a potentially important channel by which institutions can become persistent. This paper is also related to Brandts and Holt (1992), Cooper, DeJong, Forsythe, and Ross (1990), Van Huyck, Battalio, and Beil (1993), Cachon and Camerer (1996), and Blume and Ortmann (2007) where pre-play actions and communication can lead to coordination on Pareto efficient equilibria, though the mechanism by which initial actions impact future actions is different.

While this paper is the first to analyse information externalities in private-order CEIs, there is work that studies the impact of information externalities in the persistence of public-order institutions. Jehiel and Newman (2014) study an intergenerational environment in which contracts put into place today limit observation of potentially detrimental actions in the future. As principles are replaced, loopholes can arise. Warren and Wilkening (2012) study information externalities in regulation where adopting some regulatory policies limit what can be learned about the state of nature. Even with benevolent social planners, the information externality can lead to persistence in information suppressing institutions. Fernandez and Rodrik (1991) and Friedrich (2013) combine information externalities with voting and rent seeking to study policy persistence.

The current paper is also related to the literature on asymmetric information in equilibrium markets, especially the areas related to refinements and equilibrium selection. Closest in modeling spirit is Gale (1992) which uses a similar rational expectations equilibrium concept to study equilibrium selection in a general equilibrium framework with adverse selection. Whereas Gale and similar papers such as Rothschild and Stiglitz (1976), Riley (1979), and Hellwig (1987) attempt to develop selection criterion for a single equilibrium, this paper is interested in an environment where multiple stable equilibria exist.²

The paper is organized as follows. Section 2 builds the theoretical model and characterizes its rational expectations equilibria in terms of efficiency and information. Section 3 develops the experimental design. Section 4 reports the main experimental results and is divided into three parts. Section 4.1 looks at initial convergence of the experimental market in the Safe and Hazardous environments. Section 4.2 demonstrates the difference in adaptation between the unmediated equilibrium and the mediated equilibrium. Section 5 concludes.

²See also Nöldeke and Samuelson (1997) for a dynamic model in which both pooling and separating equilibria might be stable.

2 Theoretical Motivations

In this section we provide the theoretical motivations for our experiment. We begin by developing a rational expectations model of a market with heterogeneity in seller costs and costly certification. We show that in this market multiple stable equilibria exist that vary in the use of the certification technology. We then study the informational properties of these equilibria to understand how public and private signals might be useful to update beliefs about the underlying distribution of seller costs. We conclude by discussing how the lack of updating in the separating equilibria may lead to its overall persistence.

To emphasize intuition, we discuss the equilibria in relation to the parameters used in the actual experiment. A formal and more general construction of the rational expectations equilibrium is included in the appendix.

2.1 Primitives

Consider a world with experience goods of high (H) and low (L) quality which are referred to as “units”. There are $N = 15$ buyers indexed by $i \in \{1, \dots, 15\}$ divided into a finite number of types $b \in \mathcal{B}$. There are $M = 12$ sellers indexed by $j \in \{1, \dots, 12\}$ divided into three types $s \in \{G, C, B\}$ (Good, Conditional, and Bad). The number of buyers who are of type b is N_b . Likewise the number of sellers who are of type s is M_s . There is exactly one type- B seller (i.e. $M_B = 1$). The true proportion of type- G sellers and type- C sellers is \mathbf{g} and \mathbf{c} respectively.

Each buyer can consume a single high- or low-quality unit. Likewise, each seller can produce a single high- or low-quality unit. Initially we consider the case where there is only one type of buyer denoted by λ_0 . Buyers of type λ_0 have gross utilities for consuming the high and low quality good of $U^H = 200$ and $U^L = 100$ relative to a separable numéraire good, are risk and loss neutral, and receive zero utility if they do not trade. Thus the net utility of a buyer receiving a good of quality q at price P is simply $U^q - P$. Buyers of type λ_0 also have a common (though potentially incorrect) prior about the proportion of type- G sellers in the environment. Let $p(\hat{g})$ be the prior distribution regarding the proportion of good types in the economy, which has support over $g \in \{0, \frac{1}{M}, \frac{2}{M}, \dots, \frac{M-1}{M}\}$ and expected value $\mathbb{E}(\hat{g})$.

The quality of units being traded is initially unknown to buyers. However, sellers have access to a private-order CEI that can certify quality. Certification costs $T = 60$ and eliminates all uncertainty over the quality of the unit to the buyer. This certification cost is common knowledge and is paid by the seller when a trade occurs. Since $U^H > U^L$, certifying the low-quality unit can not increase its value and thus a certified low-quality unit will never be offered by a profit maximizing firm. Analysis is thus restricted to cases where all certified units are of high quality.

Seller who produce and exchange low-quality units pay no marginal cost. However, there exists a public-order CEI that is able to detect the sale of low-quality products with probability α and fine sellers an amount F . We assume that all sellers are risk and loss neutral. Thus, the expected cost of producing a low-quality unit is $C^L := \alpha F$, which we set to 50 in the experiment.

If a seller of type s produces and exchanges a high-quality unit, she pays a marginal cost C_s^H , where $C_B^H = 130$, $C_C^H = 80$, and $C_G^H = 30$. Types are defined such that

$$C_B^H > C_C^H > \alpha F > C_G^H \quad (1)$$

and

$$C_B^H > \alpha F + U^H - U^L - T > C_C^H. \quad (2)$$

Condition (1) implies that (low-cost) type- G sellers have no incentive to sell low-quality units based on the incentives generated by the public-order CEI. They will thus always produce high-quality units regardless of their certification decision. Condition (2) implies that (high-cost) type- B sellers will never have an incentive to produce certified units for any potential set of equilibrium prices and will always trade low-quality units in the uncertified market. Both conditions in combination imply that (moderate-cost) type- C sellers have an incentive to produce low-quality uncertified units in the uncertified market but may find it worthwhile to certify their goods if all type- G sellers certify theirs.

Note that $\alpha F < U^L$, which implies trade is always welfare improving, and that $C_B^H - \alpha F < U^H - U^L$, which implies that the social optimum occurs when all three seller types produce high-quality units. As type- B sellers always produce low-quality units, all equilibria are inefficient relative to the first best.

2.2 The Rational Expectation Equilibria

A rational expectations equilibria is one in which given a set of prices, (i) each seller offers a certified or uncertified product that maximizes their expected utility, (ii) each buyer chooses to buy (or not buy) a certified and uncertified object that maximizes their utility given correct predictions regarding the certification decision of each seller type, and (iii) the supply and demand for certified and uncertified items are equal. To determine the rational expectation equilibrium, it is easiest to think of certified and uncertified units being independent markets each with its own price. Let $m \in \mathcal{M} = \{\mathcal{C}, \mathcal{NC}, \emptyset\}$ be the set of markets, where \mathcal{C} is a market for high-quality certified units, \mathcal{NC} is a market of uncertified units, and \emptyset is a “market” without trades. Further, let $P^{\mathcal{C}}$ be the price for high-quality certified units and $P^{\mathcal{NC}}$ be the price of uncertified units of unknown quality. Finally, assume that $P^{\mathcal{NC}} \geq \alpha F$ so that the sellers, who are on the long end of the market, will always have an incentive to trade.³

For a given set of prices, a seller’s decision on which market to trade in is based on (i) the difference in prices between the markets $\Delta P = P^{\mathcal{C}} - P^{\mathcal{NC}}$ and (ii) the difference in cost between selling certified high-quality goods and producing a unit of either quality in the uncertified market. A seller will enter into the certified market if

$$\Delta P \geq T + \max(0, C_s^H - \alpha F). \quad (3)$$

³This assumption will always hold in equilibrium.

Under the parameters chosen for the experiment, condition (3) implies that type- G sellers will always certify their goods if $\Delta P > 60$, type- C sellers will always certify their goods if $\Delta P > 90$, and type- B sellers certify their goods if $\Delta P > 140$. This implies when $\Delta P < 60$, all three types of sellers will produce and trade uncertified units, while if $\Delta P \in (60, 90)$, type- G sellers will produce and trade certified units and the other types will produce and trade uncertified units. Likewise, when $\Delta P \in (90, 140)$, both type- G and type- C sellers will produce and trade certify units while only type- B sellers will produce and trade in the uncertified market.⁴

As the difference in prices fully informs buyers about the incentives of each seller type, a buyer's likelihood of receiving a high-quality unit in the uncertified market is dependent on the set of prices he observes. Let $\pi^H(\Delta P, \mathbb{E}(\hat{g}))$ be the likelihood that a buyer receives a high-quality unit given the observed prices and the expected number of type- G sellers he expects in the market. As buyers are on the long end of the market, the utility gained from buying a certified and uncertified unit must be exactly equal to zero and thus equal to each other. This requires that:

$$\pi^H(\Delta P, \mathbb{E}(\hat{g}))U^H + (1 - \pi^H(\Delta P, \mathbb{E}(\hat{g})))U^L - P^{\mathcal{N}C} = U^H - P^C = 0. \quad (4)$$

Note that the second equality here requires that $P^C = U^H$ and thus that the price of certified units (when traded) is pinned down by the assumption of excess demand.

For a high enough initial belief about the proportion of type- G sellers in the environment, two rational expectations equilibria exist which vary in the use of the certification technology. These equilibria are as follows:

- **Mediated Equilibrium:** $P^C = U^H = 200, P^{\mathcal{N}C} = U^L = 100$. Type- G and type- C sellers produce and sell certified high-quality units. Type- B sellers produce uncertified low-quality units. $M_G + M_C$ buyers buy in the certified market and M_B buyers buy in the uncertified market.
- **Unmediated Equilibrium:**⁵ $P^{\mathcal{N}C} = U^H - (1 - \mathbb{E}(\hat{g}))(U^H - U^L) = 100 + 100 * \mathbb{E}(\hat{g}), P^C = U^H = 200$. Type- G sellers produce uncertified high-quality units. Type- C and type- B sellers produce uncertified low-quality units. M buyers buy from the uncertified market.

The stability and existence of the two potential equilibria can best be seen by plotting the marginal buyer's belief about receiving a high-quality unit in the uncertified market as a function of ΔP and relating this belief to the buyer's demand over certified and uncertified goods. As seen by the line segment in Figure 1 marked "Type- G Sellers Do Not Use Private-Order CEI," when the difference in prices are smaller than T , all seller types trade in the uncertified market. As only the

⁴When a seller is indifferent between the two markets, the rational expectations model can assign them to either market in order to balance supply and demand. This case is discussed in an extension of the model where heterogeneity in buyer types leads to a partially-mediated equilibrium.

⁵Note that in the unmediated equilibrium, there are no sellers in the certified market and thus beliefs about the distribution of seller types in the certified market are arbitrary. While each set of beliefs could technically be considered a different rational expectations equilibrium, for exposition purposes they are classified as a single equilibrium since their price and quantity characteristics are the same.

type- G buyers will provide high-quality units in the uncertified market, the likelihood of receiving a high-quality unit in this market is equal to the proportion of type- G sellers expected in the market, $\mathbb{E}(\hat{g})$. When the difference in price is greater than ΔP , however, the type- G sellers certify their good and the probability of receiving a high-quality unit is zero.

Just as we can plot beliefs as a function of ΔP , a buyer's demand over certified and uncertified units can be mapped into the space of ΔP and $\pi^H(\Delta P, \mathbb{E}(\hat{g}))$. As can be seen by the regions marked "Buyer Demand Certified" and "Buyer Demand Uncertified," a buyer will demand certified units if the difference in price is small or his belief about receiving high-quality units in the uncertified market is pessimistic. As the buyers are in the long end of the market, the buyers must be indifferent between certified and uncertified units. These indifference points are shown by the downward sloping (blue) line separating the two regions.

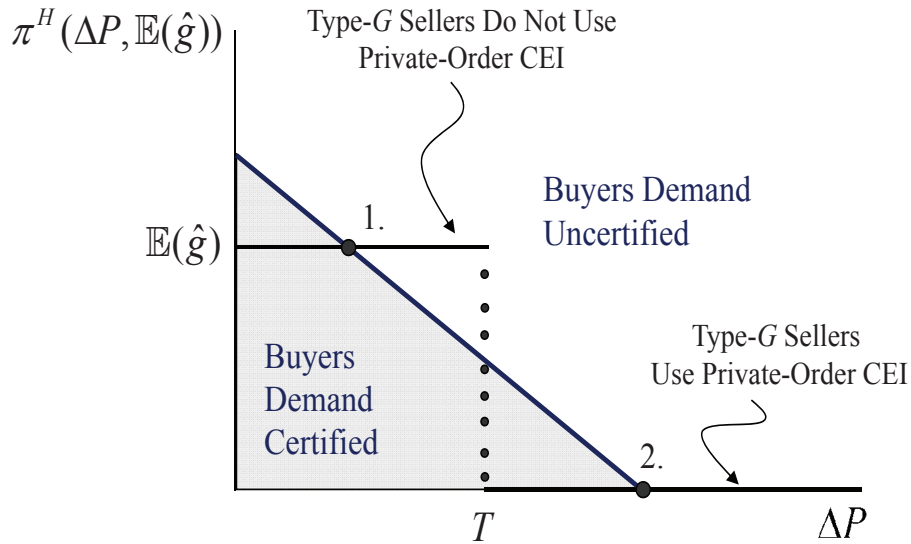


Figure 1: The two potential rational expectations equilibria as a function of the difference in price for a certified and uncertified good (ΔP) and the corresponding expected probability of receiving a high-quality unit in the uncertified market ($\pi^H(\Delta P, \mathbb{E}(\hat{g}))$). The black line represents the expected proportion of high-quality units in the uncertified market based on the optimal actions of the sellers and beliefs of the buyer. When the difference in price of certified units is less than the certification cost T , good sellers trade in the uncertified market leading to an expected proportion $\mathbb{E}(\hat{g})$ of high-quality units in the uncertified market. When the difference in price is greater than T , good sellers certify their goods, eliminating high-quality units from the uncertified market. As buyers are on the long end of the market, the utility gained from buying a certified and uncertified unit must be exactly equal to zero and thus equal to each other. This is represented by the downward sloping (blue) line. The unmediated equilibrium and mediated equilibrium occur at points (1.) and (2.) respectively. The partially-mediated equilibrium is unlikely to exist in the case of risk-neutral buyers due to units being traded in integer values.

Rational expectations equilibria exist in each location where the buyer's belief function intersects the buyer's indifference condition. If the proportion of type- G individuals is (believed to be) small

or the transaction cost T is small relative to the value of quality, the two lines intersect only once and only the mediated equilibrium exists. If the proportion of type- G sellers (is believed to be) large, however, a second equilibrium emerges in which no seller certifies and high and low-quality units are traded within a single market. It follows:

Proposition 1 *Existence:* *The mediated equilibrium always exists. The unmediated equilibrium exists if and only if $(1 - \mathbb{E}(\hat{g}))(U^H - U^L) \leq T$.*

Under the parameterizations chosen in the experiment, the unmediated equilibrium will exist if $\mathbb{E}(\hat{g}) \geq .4$.

2.3 Market Information

Having defined the mediated and unmediated equilibrium, we now return to the central question of information and the organization of markets. We begin in the most straight forward case where all buyers in the market are homogeneous and have the same prior $p(\hat{g})$ about the proportion of type- G sellers in the environment. Based on the market equilibrium, we determine what a new buyer could learn from observing the market price. In section 2.3.1 we allow for buyers to have heterogeneous beliefs about the distribution of good sellers in the market and ask whether these beliefs converge to the true value as a result of repeated trade.

Consider a period in which all buyers have the same (potentially incorrect) prior about the proportion of type- G sellers. If a new buyer enters the market and observes price and the volume of trades in each market, what can he deduce about the proportion of sellers who are good, conditional and bad?

In the mediated equilibrium, the prices $P^C = U^H$ and $P^{NC} = U^L$ only provide information about the demand function of buyers. Since only bad sellers trade in the non-certified market, the share of units traded in the uncertified market provides information on the proportion of sellers who are of type- B but provides no additional information about the relative proportion of type- G and type- C sellers. From the perspective of a new buyer who already knows that $M_b = 1$, the market primitives are uninformative.

By contrast, in the unmediated equilibrium, the price of uncertified goods, $P^{NC} = 100 + 100 * \mathbb{E}(\hat{g})$, and thus $\mathbb{E}(\hat{g}) = \frac{P^{NC} - 100}{100}$. Thus, given only the uncertified price and knowledge about U^H and U^L , a new buyer can determine $\mathbb{E}(\hat{g})$.

Proposition 2 *In an unmediated equilibrium with a common prior, price is a sufficient statistic for $\mathbb{E}(\hat{g})$. In the mediated equilibrium, no market signal generates information that can distinguish between type- G and type- C sellers.*

2.3.1 Heterogeneous Beliefs and Learning

The discussion above highlights the relationship between market organization and the informativeness of market primitives. However, it is based on the premise that individuals who are in the

market have a common prior. As this is precisely the information which is of interest in evaluating the existence of the unmediated equilibrium and the efficiency of both markets, it is of interest to determine under what conditions individuals can learn this distribution of values under repeated trade. We show that under the unmediated equilibrium, at least M buyers learn the proportion of type- G sellers even in cases where buyers are myopic. Further, since the pivotal buyer is fully informed over time, all buyers learn the distribution of types if they correctly incorporate information from market prices into their posterior. By contrast, we show that in the unmediated equilibrium no agent can distinguish between type- G and type- C consumers in the economy and thus beliefs regarding the proportion of these groups may be arbitrary.

To begin, let $p_t^i(\hat{g})$ be the prior distribution of buyer i at time t regarding the proportion of good types in the economy which has support over $g \in \{0, \frac{1}{M}, \frac{2}{M}, \dots, \frac{M-1}{M}\}$ and where the discrete distribution is single peaked.⁶ Further, define the type of an individual by his prior.

For a given price and allocation rule, a rational expectations equilibrium is *ex post* stable if no individual desires to change their allocation given the revelation of information from that allocation. As price is a required component of the allocation rule, and this price is pinned down by the value of the last buyer who is willing to trade, we require that each buyer must be willing to trade given the revelation that they are the pivotal buyer. In the unmediated equilibrium, this requires that for each buyer assigned an uncertified unit:

$$P^{NC} \leq U^L + \mathbb{E}(\hat{g}|P^{NC})(U^H - U^L). \quad (5)$$

Let P^* be the largest P^{NC} satisfying equation (5) for at least M buyers. Then, if $P^* \geq T + U^L$ an unmediated rational expectations equilibrium exists where M buyers trade at the price P^* .⁷

Consider the case where all buyers are myopic and do not take price into account. In this case, each of the M individuals who receive a unit discover its quality and update their beliefs from their private purchase experiences alone. As there are M individuals trading each period, there are at least M individuals who update their beliefs in a given period. As these individuals continue to get new information regarding the true valuation of the good, their priors converge to the true distribution over time.

Proposition 3 *Consider a sequence of periods in which the unmediated equilibrium occurs each period and individuals update their beliefs only from their private purchases. Then there exists at least M buyers such that*

$$p_t^i(\hat{g}) \xrightarrow{a.s.} \mathfrak{g}. \quad (6)$$

⁶Single peaked priors are not required for the convergence of beliefs but ensures that the willingness of individuals to buy in the uncertified market is decreasing in price when the unmediated equilibrium exists. We can think of these beliefs as arising from previous purchases of uncertified goods in the environment. In this way, the heterogeneous priors assumption can be thought of as a common prior with additional information coming from a random generating process of initial trades.

⁷As the demand function is now downward sloping and discrete, any price between P^* and the willingness to pay of the $(M + 1)$ th can be supported as an equilibrium. Choosing the price for which the last buyer is indifferent to trading ensures that this party knows with certainty that he is pivotal.

An individual who is updating optimally can discard any information which decreases the precision of his or her posterior. As such, the worst posterior an individual can have after each period is the myopic one where individuals use information only from their private signals. It follows that there exists at least M individuals who have accurate beliefs of \mathbf{g} over time. As P^* is pinned down by the value of the M th buyer, and his beliefs are accurate, $\mathbb{E}(\hat{g}|P^{NC}) \rightarrow \mathbf{g}$ and the trade price gives perfect information regarding the value of the good. Thus, over time, price is informative even in cases where individuals have different beliefs and heterogeneous priors.

By contrast, in the mediated equilibrium, individuals in the market for certified and uncertified goods learn no new information from their purchases since the qualities are guaranteed. Further, the market price carries no information about the priors of the buyers in each period of time. It follows that beliefs regarding the proportion of type- G sellers in the certifying equilibrium may be arbitrary and that there is no reason to expect convergence to true beliefs over time.

Proposition 4 *Consider a sequence of periods $t = 0, \dots, \infty$ in which the mediated equilibrium occurs each period and individuals update their beliefs optimally. Then for all i ,*

$$p_0^i(\hat{g}) = \dots = p_{t+1}^i(\hat{g}) = \dots = p_\infty^i(\hat{g}). \quad (7)$$

As can be seen from proposition 4, the mediated equilibrium eliminates all information that might be used to update beliefs. Thus, if a market reaches a mediated equilibrium, it is likely to get stuck in this market organization. Further, if there is an exogenous shift in the proportion of type- G and type- C sellers, buyer beliefs will remain unchanged.

2.3.2 Heterogeneity in Loss Preferences, Partially-Mediated Equilibria, and Public Information

In experimental settings, individuals typically exhibit heterogeneous levels of risk and loss aversion. Even if all individuals have common beliefs about the distribution of seller types, such heterogeneity can lead to partially-mediated equilibria which have slightly different informational properties than either the unmediated equilibrium or the mediated equilibrium. We characterize these equilibria and discuss their informational properties before moving on to the experimental design.

Consider an extension of the baseline model where buyers are loss averse and suffer additional disutility for trades that end in a loss.⁸ Let $\mathcal{B} = \{\lambda_1, \lambda_2, \dots, \lambda_N\}$ where λ_i is the idiosyncratic loss aversion parameter for buyer i with $\lambda_i \geq 1$ for $i \in \{1, 2, \dots, N\}$ and return to the baseline case where all individuals have a common prior $p(\hat{g})$. Without loss of generality, we order buyers according to their loss aversion parameter such that $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_N$ and continue to normalize the utility obtained from not trading to zero.

⁸The intuition developed here holds for heterogeneity due to risk aversion and most reference dependent utility models. We have chosen loss aversion due to its tractability and due to answers in the exit survey. In the exit survey we asked buyers, ‘‘How did you decide on the price you were willing to pay for an uncertified good?’’ 53% of respondents indicated that they were unwilling to take losses or factored in the potential for losses into their decisions.

In the unmediated equilibrium, the market price $P^{\mathcal{N}C} > U^L$ and thus there is a potential for losses. When a buyer receives a low quality unit in the uncertified market, his net utility is $-\lambda_i(P^{\mathcal{N}C} - U^L)$ which is decreasing in λ_i . Since buyers are heterogeneous in loss aversion, the aggregate demand curve for uncertified units becomes downward sloping and the uncertified price is pinned down by the loss aversion of the M^{th} buyer. If the M^{th} buyer is sufficiently loss averse, he may be unwilling to trade for uncertified units at a price where $\Delta P \geq T$. In this case, partially-mediated equilibria may form. Let S^C be the number of certified units in an equilibrium. Then for each $S^C < M_G$, a partially-mediated equilibrium may exist with the following properties:

- **Partially-mediated Equilibria:** $P^{\mathcal{N}C} = U^H - T$, $P^C = U^H$. Type- C and type- B sellers produce uncertified low-quality units. S^C type- G sellers produce certified high quality goods. $M_G - S^C$ type- G sellers produce uncertified high quality goods. Buyers $i \in \{1, \dots, M - S^C\}$ buy uncertified units. S^C other buyers buy certified units.

In the benchmark model, the partially-mediated equilibria were unlikely to occur because both type- G sellers and all buyers needed to be indifferent between trading in the certified and uncertified market. With heterogeneity in buyer preferences, however, partially-mediated equilibrium may be stable since the willingness to pay for uncertified units is decreasing in loss aversion leading to a downward sloping aggregate demand function.

In the partially-mediated equilibrium, since $P^{\mathcal{N}C} = U^H - T = 140$ and $P^C = U^H = 200$, price alone does not convey information about the proportion of type- G sellers. However, as only type- G sellers are in the certified market, an individual can use the size of the certified market to partially update his beliefs. The precision of the posterior distribution may be further improved if individuals receive additional information regarding the proportion of high-quality units observed in the uncertified market. Such public information is uninformative in the mediated equilibrium and may be uninformative in the unmediated equilibrium if all parties have common beliefs.

3 The Experiment

The theoretical model suggests that the adoption of certification by market participants can have a strong impact on the informativeness of public and private signals. In markets where the mediated equilibrium has formed, publicly observed prices are uninformative and individuals do not learn from their private purchases. By contrast, in the unmediated equilibrium, price provides information regarding $\mathbb{E}(\hat{g})$, while private purchases refine this expectation toward the true proportion of type- G sellers. Thus, the combination of private experiences and public information should allow all individuals to track changes in the level of risk in the uncertified market over time.

In the remaining sections of the paper, we study an experimental market in which the differences in the informativeness of public and private signals between the two equilibria are predicted to influence usage patterns of private-order CEIs adapts as the environment changes. The goal of our design is to begin trade in environments in which the mediated and unmediated equilibria reliably

form and then perturb the underlying distribution of sellers in a way that should be undetectable in the mediated equilibrium, but which makes this equilibrium highly inefficient. We study both the way in which markets respond to these perturbations as well as studying how individuals learn in each environment and equilibrium.

3.1 Valuations and Costs

Each session of the experiment consisted of 5 buyers and 6 sellers who interacted in a sequence of 24 market periods. Each market period consisted of two simultaneous exchanges — one with certification and one without — in which buyers and sellers could exchange high-quality “red” units and low-quality “blue” units.

In a given period, each of the six sellers had capacity to produce and sell a total of two units across both markets in any combination of high and low quality. As shown in Table 1, sellers could be assigned one of three possible cost functions for producing high- and low-quality units which, following the notation of section 2, we designate as G , C , and B (Good, Conditional, and Bad).

To avoid seller-side risk and loss aversion from playing a role, low-quality products had a fixed marginal cost of 50 points which reflects the enforcement capabilities of the public-order CEI in our original model. Type- G sellers had a lower cost for producing a high-quality unit than low-quality units, type- C sellers had a slightly higher cost for producing high-quality units than low-quality units, and type- B sellers had a very high cost for producing high-quality units.

Table 1: Seller Production Costs

	Uncertified Low Units	Uncertified High Units	Certified High Units
<i>Good</i>	50	30	90
<i>Conditional</i>	50	80	140
<i>Bad</i>	50	130	190

The cost of using the private-order CEI, called a “certification cost,” was 60 points. This cost was known to both buyers and sellers. If the difference in price between the certified and uncertified market grew larger than the certification cost, type- G sellers had an incentive to sell a high-quality unit in the certified market rather than a high-quality unit in the uncertified market. Likewise if the difference in price between the certified and uncertified market grew larger than 90, type- C sellers had an incentive to sell a high-quality unit in the certified market rather than a low-quality unit in the uncertified market.

Each of the five buyers could purchase a total of three units across both markets creating an aggregate demand of 15 units. Since sellers could produce a total of 12 units, each experimental period had excess demand. This excess demand was implemented to allow sellers to capture any residual surplus that existed in either of the two markets and to capture rents generated through certification.

Buyers and sellers were allowed to trade multiple units in order to increase the thickness of the market and to avoid using passive buyers who might cause noise in the experiment by trying

to participate. The supply and demand curves were constructed so that no seller or buyer could change the equilibrium price by more than 10 points by withholding their entire supply or demand from the market. This was small relative to the market prices which ranged from 100 to 200 points. Since no buyer or seller had market power, the mediated and unmediated equilibrium for the experimental environment are the same as the simplified model of section 2.1. The set of potential partial-unmediated equilibria is slightly smaller in the experimental environment since the loss aversion coefficient for multiple units is from the same individual buyer. However, the price and informational properties of these partially-mediated equilibria remain the same.

As shown in Table 2, each buyer’s demand schedule was downward sloping. This downward slope was implemented to generate some surplus for the buyers, which is shown by Holt, Langan, and Villamil (1986) to improve the speed of convergence in markets. Conditional on buying a unit, the valuation of both the high- and low-quality units declined for each unit purchased. Thus, if buyer 1 had purchased a low-quality unit and then purchased a high-quality unit, his valuation for the two units would have been 140 and 220 respectively. The demand functions of buyers four and five were staggered slightly to smooth the aggregate demand function.

Table 2: Buyer Valuations

	Buyers 1-3			Buyers 4-5			
	Unit 1	Unit 2	Unit 3	Unit 1	Unit 2	Unit 3	
High Quality	240	220	200	High Quality	230	210	190
Low Quality	140	120	100	Low Quality	130	110	90

Earnings from one period did not carry over into the following periods. After each trade, the type of unit purchased was revealed and a buyer’s earnings or losses from the transaction were added to or subtracted from his current cash. To avoid bankruptcy, buyers were given 100 points as an initial cash endowment in each period. If at any point during a period a buyer had negative earnings, his trading privileges for the period were revoked. This form of bankruptcy was infrequent, occurring only 8 times out of the 1728 unique buyer-period observations.

3.2 The Trading Mechanism

Trade was conducted through two computerized exchanges where both buyers and sellers were anonymous. The only distinguishable feature between the various seller offers and buyer bids were the public price and quality characteristics visible in the exchange.

Each exchange was conducted as a double auction.⁹ Departing slightly from the design developed by Smith (1964), subjects in this experiment were free to enter the bid and ask queues at any price. Subjects were also free to accept any offer from the opposite side of the market and were not

⁹A double auction mechanism is traditionally defined as one in which 1) both buyers and sellers can submit bids and asks to a centralized exchange, 2) trade occurs continuously over a fixed time interval, and 3) trade occurs any time a buyer’s bid is above a seller’s ask or a seller’s ask is below a buyer’s bid. Due to moral hazard and the potential that low prices are informative of low value, we do not automatically fill transactions but instead require the second party to manually accept the offered contract from the other side of the market.

bound to accept the lowest possible price. These changes allowed sellers some flexibility in their pricing strategies and allowed buyers a way to avoid offers that they believed to be of low quality.

In the uncertified market, a seller who posted an offer publicly submitted an asking price and secretly selected the quality of the offered unit. A buyer who bid in the uncertified market publicly submitted a bid price and a quality request. Quality requests in the uncertified market were not binding and a seller who filled a request had the option of supplying either a high-quality unit or a low-quality unit. Information about the actual quality of units traded in the uncertified market were private and revealed only to the buyer who purchased the unit.

In the certified market, the quality of the seller's offered unit was observable and quality requests by buyers were binding.¹⁰ If a seller transacted in the certified market, either by having an offer accepted or fulfilling a buyer's trade request, she was charged the certification fee of 60 points.

Each seller could have one certified offer and one uncertified offer open at one time. Likewise, each buyer could have one certified bid and one uncertified bid open at any given time. If a seller sold her last unit or a buyer exhausted his demand, all their remaining open contracts were automatically withdrawn from the market. Bids and offers could be changed or withdrawn at any time with no restriction on pricing.

In the first three periods of the experiment, each trading period lasted four minutes to allow for subjects to become accustomed to the interface. In the remaining periods, the trading period lasted two minutes.¹¹

3.3 Information

Information about seller costs and buyer valuations was private information. At the beginning of the experiment, sellers were shown the three possible cost functions that they might be assigned in the instructions and told that their cost schedule might change across periods. Sellers were not given information on the assignment of other sellers or on the demand schedule of the buyers. Buyers were given only their own demand schedule and were informed that some of the sellers might have a lower cost for producing high-quality units than low-quality units.

In each period, a history of trades from the current period was available in graph form for all subjects in the market. Certified trades were shown in the color of the actual unit traded while uncertified trades showed up as black lines. If a buyer purchased an uncertified unit in a period, he was privately informed about the quality of the unit at the time of sale.

After each trading period, both buyers and sellers participated in a bonus phase. The bonus phase elicited beliefs about the number of type- G sellers. Subjects were paid a bonus of 20 points in each round they were correct. The bonus phase served as a measure of beliefs regarding the likelihood of receiving a high-quality unit. In the main treatments, individuals received no feedback regarding the accuracy of their guess between rounds.

¹⁰Buyers were free to request certified low-quality units. In practice, this rarely occurred.

¹¹One might be concerned that two minutes was too short for each period. However, in practice the double auctions cleared quickly. Over all treatments and periods, 73.3% of periods had 12 units traded, 19.9% of periods had 11 units traded, 6.3% of periods had 10 units traded, and 0.5% of periods had 9 units (or less) traded.

Following the bonus game, subjects were given a summary sheet which varied by the information treatment. In half of the main sessions, individuals were only informed about the total number of units traded with and without certification. In the remaining sessions, individuals were informed in the information screen about the actual number of high- and low-quality units traded in the uncertified market. These information variants are referred to as the “Private” and “Public” Information treatments respectively and are discussed below. Information was given *ex post* rather than during the trading period to keep the trading environment as similar as possible across treatments.

3.4 Treatments

Experimental sessions were divided into four treatments which varied in the amount of public information available about past trades and in the degree of moral hazard (the number of type-*C* sellers). Half the treatments were conducted using the Public Information treatment discussed in section 3.3. As was noted in the theory section, the public revelation of units traded in the uncertified market should generate new information in the partially-mediated equilibria that might form if buyers are heterogeneous in their willingness to accept gambles or in their beliefs. While not explicitly modeled, we expected that the public information treatment would increase the number of buyers who are willing to trade uncertified units when the partially-mediated equilibrium forms. We predicted no effect in markets where the mediated equilibrium formed.

Treatments were next stratified into two environments — Safe (\mathcal{S}) and Hazardous (\mathcal{H}) — which varied in the number of sellers who were assigned to the three seller types. In the Safe environment, five of the sellers were of type *G* and one seller was of type *B*. In the Hazardous environment, one seller was of type *G*, four sellers were of type *C*, and one seller was of type *B*. The single type-*B* seller was included in both treatments in order to have both certified and uncertified prices available when the mediated equilibrium formed.

Table 3: Moral Hazard Environments

	Good	Conditional	Bad
Safe (\mathcal{S})	5	0	1
Hazardous (\mathcal{H})	1	4	1

In the sessions that began in a Safe environment, the environment was switched to the Hazardous environment at period 13 by assigning new cost charts to four of the sellers who were originally of type *G*. This process was reversed in the sessions beginning in the Hazardous environment. To distinguish between periods before and after the switch, *Pre* and *Post* superscripts are appended to the environment identifier.

As the Hazardous and Safe environments are our main treatment variable, it is useful to discuss their design. The goal of our design was to study the link between the utilization of the private-order CEI and information. Thus, we wanted to begin with two environments: one where the mediated equilibrium was likely to form and one where the unmediated or partially-mediated equilibria was

Table 4: Treatments

Treatment	Periods 1-12	Periods 13-24	Information	Identifiers
1	Safe	Hazardous	Private	$\mathcal{S}^{Pre}, \mathcal{H}^{Post}$
2	Safe	Hazardous	Public	$\mathcal{S}^{Pre}, \mathcal{H}^{Post}$
3	Hazardous	Safe	Private	$\mathcal{H}^{Pre}, \mathcal{S}^{Post}$
4	Hazardous	Safe	Public	$\mathcal{H}^{Pre}, \mathcal{S}^{Post}$

likely to form. To this end, the Hazardous environment was designed so that, under full information about the distribution of types, only the mediated equilibrium existed. Our prediction here was that individuals who started in this environment and traded in the uncertified market early would update their beliefs downward and drive the risk premium past the certification cost. This would lead to the formation of the certifying equilibrium in markets that started in the Hazardous environment. Note that the consistent formation of the certifying equilibrium hinges on the ability of buyers to update their beliefs (either through trade or market signals) from early periods and the assumption that the price of uncertified units would adjust downward as a function of these beliefs. In the Hazardous environment, the predicted mediated equilibrium had the following properties:

- **Mediated Equilibria for Hazardous Environment** : $P^C = 200, P^{NC} = 100$. Type-*G* sellers sell certified high-quality units for a surplus of 110 per unit. Type-*C* sellers sell certified high-quality units for a surplus of 60 per unit. Type-*B* sellers produce uncertified low-quality units for a surplus of 50 per unit.

The Safe environment was designed so that under the full information about the distribution of types, the mediated equilibrium was extremely unlikely to form or persist. Due to the non-strategic nature of the rational expectation equilibria used as a solution concept, the mediated equilibrium is always an admissible outcome as a full-information equilibrium outcome. Nonetheless, the Safe environment was designed so that under full information, if a single type-*G* seller switched to the uncertified market, a loss-neutral buyer who knew the proportion of agents in each market would be willing to pay $.5U^H + .5U^L$ for an uncertified good and U^H for a certified good. Since $U^H - U^L$ was 100 points across all units, the difference in willingness to pay for a certified and an uncertified unit was $.5(U^H - U^L) = 50$. This difference was less than the certification cost of 60 points. Thus under full information, a paired deviation from the mediated equilibrium by a seller and loss-neutral buyer could eliminate the mediated equilibrium.

The consistent formation of the unmediated and partially-mediated equilibria in the Safe environment hinges on individuals having a high enough initial belief that buyers are willing to trade uncertified units at high prices in early periods and that the distribution of loss aversion was such that at least some buyers were willing to trade uncertified units even if the underlying distribution was known. If all buyers were loss neutral, the unmediated and mediated equilibrium under the safe environment were as follows:

- **Unmediated Equilibrium for Safe Environment** : $P^{NC} = 183$. Type-*G* sellers produce

uncertified high-quality units for a surplus of 153 points per unit. Type-*B* sellers produce uncertified low-quality units for a surplus of 133 per unit. All trades occur in the uncertified market.

- **Mediated Equilibrium for Safe Environment** : $P^C = 200, P^{NC} = 100$. Type-*G* sellers sell certified high-quality units for a surplus of 110 per unit. Type-*B* sellers produce uncertified low-quality units for a surplus of 50 per unit.

To study adaptation at period 13, the Hazardous and Safe treatments were further designed so that all changes to the seller types would be in the reassignment of type-*G* and type-*C* types. As was shown in the theory section, these changes are not expected to be observable by buyers in the mediated equilibrium since market prices and private consumption are uninformative. Thus, under the auxiliary assumption that markets converge to this equilibrium in the Hazardous environment, theory would predict that buyers cannot observe the changes in seller types and that sellers cannot coordinate to the unmediated equilibrium. Vice versa, buyers trading uncertified goods in an unmediated or partially-mediated equilibrium are exposed to additional low-quality units when the Safe treatment is changed to Hazardous. If individuals respond to private and public signals, it is predicted that market price will be responsive when the Safe environment is changed to Hazardous.

As the mediated equilibrium is always an equilibrium regardless of environment, a final concern in the design of treatments is that sellers who have their type changed must wish to reveal this information to the broader environment and shift the market to a different market structure. Comparing the two equilibria, type-*G* sellers receives a surplus of 153 points in the unmediated equilibrium versus 110 points in the mediated equilibrium. The type-*B* seller receives a surplus of 133 points in the unmediated equilibrium versus 50 points in the mediated equilibrium. Thus, all sellers were better off in the unmediated equilibrium and had group incentives to coordinate to this equilibrium.¹² Equilibria were efficiency ranked in the Safe environment with the unmediated equilibrium being most efficient and the mediated equilibrium being the least efficient. As noted in Table 5, all possible equilibria were inefficient relative to the first best due to inefficient production by the type-*B* seller.

Table 5: Efficiency

	Perfect Information	Unmediated Equilibrium	Mediated Equilibrium
Safe	2100*	2060	1460
Hazardous	1700*	1100*	1060

*not supportable as an equilibrium

¹²While no conditional sellers existed in the Safe environment, this type of seller also would have preferred the unmediated equilibrium.

3.5 Protocol

Subjects in this experiment were drawn from a centralized database comprised of undergraduate students from The University of Zurich and UTH-Zurich. 12 sessions were run each composed of 11 subjects who remained in fixed groups and fixed roles over all 24 periods. Trades were conducted in points and converted to Swiss Francs at the end of the experiment at a conversion rate of 30 points to 1 Swiss franc. A session lasted on average 140 minutes and paid an average of 45 Swiss Francs (\$38 at the time of the experiment). The first 40 minutes of each session was devoted to an extensive set of written, oral, and computerized instructions which included a control quiz. All programs for this experiment were written in Z-Tree.¹³

A series of follow-up experiments were run in Melbourne in 2014 and 2015 where an additional 18 sessions were run each with 11 subjects. The protocols in these experiments were identical to the original experiments with an exchange rate of 30 points to 1 AUD.

After all 24 periods of the main experiment, aversion to gambles was measured via a series of lottery choices similar to those used in Holt and Laury (2002). Subjects made a series of decisions between a guaranteed return of 90 points and a 50-50 gamble between earning 0 and x , where x varied between 60 and 360 in increments of 30. Individuals were considered averse to gambles if they rejected the 50/50 gamble with high payment of 210. Interpreted as risk aversion with initial wealth of zero, this corresponds to a $\sigma = .19$ in a CRRA utility function of the form $u(x) = \frac{x^{1-\sigma}}{1-\sigma}$. Interpreted as loss aversion with the earnings from the safe gamble used as the reference point, this corresponds to a loss aversion $\lambda = 1.333$.¹⁴

4 Experimental Results

Empirical analysis is taken in three steps. We first establish that absent a pre-existing market structure, the efficient unmediated equilibrium or a partially-mediated equilibria forms in the Safe environment, while the mediated equilibrium forms under the Hazardous environment. This would suggest that absent an established market structure, buyers and sellers select the most efficient equilibria starting from an uninformed prior. We then turn to our main question of how the utilization of the private-order CEI influences the response of buyers and sellers to exogenous changes to the number of type- G and type- C sellers in the environment. We then explore patterns of individual learning both before and after the mediated equilibrium has formed. Follow up experiments are discussed in section 4.4 and 4.5.

¹³See Fischbacher (2007) for a description of Z-Tree.

¹⁴Counting the total number of safe gambles and setting a threshold for the number of safe choices yields a measure similar to the one used. Since some individuals had inconsistent choice patterns, this approach had a higher degree of subjectivity. Previous versions of this paper also used a loss aversion measure from the exit survey. This measure had greater variation across sessions and generated parameter estimates closer to theoretical predictions. Due to it being an *ex post* measure, the more conservative results are shown here.

4.1 Do markets converge to the efficient equilibrium?

4.1.1 Hypothesis and Empirical Strategy

The experimental treatments were designed so that absent an initial market equilibrium, the mediated equilibrium was expected in the Hazardous environment and the unmediated or partially-mediated equilibrium was expected in the Safe environment. As the convergence to these equilibria are important auxiliary assumptions to studying learning and adaptation, we begin by studying whether initial convergence takes place.

To test for initial convergence, we compare the prices of uncertified trades in the \mathcal{S}^{Pre} environment where the degree of moral hazard is low to the prices predicted in unmediated and partially-mediated equilibria. Similarly, we compare the uncertified price in the \mathcal{H}^{Pre} environment to the price predicted in the mediated equilibrium. To allow time for the market to converge, attention is restricted to periods 7-12.¹⁵ Using session fixed effects, we estimate:

$$P_{i,s} = \alpha_0 + \Sigma\alpha_s + \beta_{Cert}I_{Cert} + \beta_{\mathcal{S}^{Pre}}I_{\mathcal{S}^{Pre}} + \epsilon_{i,s} \quad (8)$$

where $P_{i,s}$ is the price of an individual trade i in session s , α_s are individual session fixed effects, I_{Cert} is an indicator for a certified trade, and $I_{\mathcal{S}^{Pre}}$ is an indicator variable for uncertified trades in the Safe environment. Note that since the estimation includes both certified and uncertified trades, session level fixed affects do not eliminate the variation in uncertified trades across treatments. Expecting the mediated equilibrium to form in the \mathcal{H}^{Pre} environment and the unmediated or partially-mediated equilibrium to form in the \mathcal{S}^{Pre} environment, we predict:

Hypothesis 1 $\alpha_0 = 100$, $\alpha_0 + \beta_{Cert} = 200$, $\alpha_0 + \beta_{\mathcal{S}^{Pre}} \in [140, 183]$.

4.1.2 Results

Result 1 *Hypothesis 1 cannot be rejected in the data. 5 of the 6 treatments that start in the Safe environment have certified and uncertified prices consistent with the unmediated or partially-mediated equilibria. All six treatments that start in the Hazardous environment have prices consistent with the mediated equilibrium.*

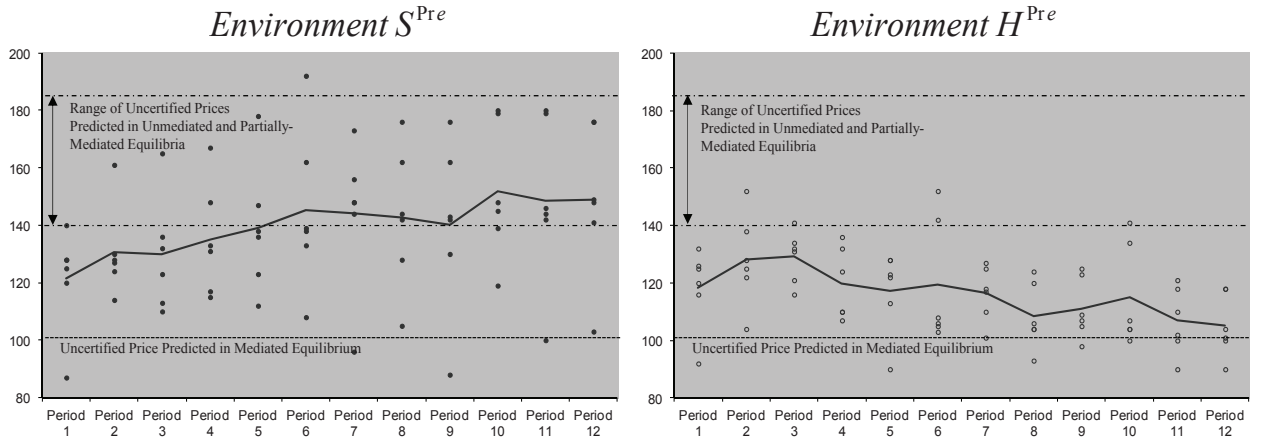
The predicted convergence of the Hazardous treatment to the mediated equilibrium and the Safe treatment to the partially-mediated or unmediated equilibrium is largely supported in the empirical data. Figure 2 shows the evolution of the average uncertified price of trades in the \mathcal{S}^{Pre} and \mathcal{H}^{Pre} environments over time. The black dots in each period are the average price of uncertified trades in each of the six sessions while the line shows the average of these session averages. As can be seen on the left hand side of Figure 2, the average uncertified price in the Safe environment

¹⁵The number of omitted periods was decided prior to running the experiment and based on two initial pilots. As can be seen in the individual experiment in section 4.2, the price of the uncertified market converges to the unmediated or partially-mediated equilibria from below. Thus, increasing the number of periods in the analysis decreases the estimated uncertified price for treatments that converge to the unmediated equilibrium.

increases over the 12 periods and falls within the region of prices predicted in the unmediated and partially-mediated equilibria in 5 out of 6 of the sessions. The variation in the uncertified price across sessions suggests that individual heterogeneity in risk aversion indeed may be influencing equilibrium selection, a hypothesis we discuss in detail below.¹⁶

As shown in the right hand side of the figure, the prices in the Hazardous environment fall over time, with average uncertified prices in four of the sessions falling to a price just above 100 and the remaining two sessions having average uncertified prices within 20 points of the benchmark prediction.¹⁷

Figure 2: Average Uncertified Prices in \mathcal{S}^{Pre} and \mathcal{H}^{Pre}



The consistency of the data with the predictions in hypothesis 1 can also be seen in the regression analysis. Table 6 presents regression results from equations 8 with varying degrees of control from the lottery treatment. As can be seen in column (1), the empirical uncertified price ($\alpha_0 + \beta_{\mathcal{S}^{Pre}} = 141.5$) is lower than the predicted unmediated equilibrium price of 183 but above the minimum price that could sustain a partially-mediated equilibrium.¹⁸

The likelihood that the partially-mediated equilibrium should form over the unmediated equilibrium is predicted to be related to the aversion to lotteries of the inframarginal buyer. We test for this in column (2), where we interact the (demeaned) number of buyers who are lottery averse with the safe treatment. Consistent with theory, the number of lottery averse individuals is negatively correlated with the price of uncertified trades.

Estimated prices for uncertified trades in the \mathcal{H}^{Pre} environment varies between 102 and 108

¹⁶As can be seen in the individual session data included in Appendix B, the intra-session variance in uncertified trades is declining over time suggesting at least partial convergence to one of the potential equilibria in all six sessions.

¹⁷At a session level, the intra-session variance of uncertified trades is very small in 4 of the 6 sessions that begin in the Hazardous environment. In these sessions, 95% of trades occur at prices between 90 and 110 in periods 7-12. In the remaining sessions, one session has a small number of trades above 150, but otherwise appears to be converging. The other session has at least one trade at a price above 150 in each period, suggesting that this session does not fully converge. Excluding this treatment from the analysis in the next section marginally increases the fit of the data to the model, but does not otherwise affect the analysis.

¹⁸The 95% confidence interval for $\alpha_0 + \beta_{\mathcal{S}^{Pre}}$ is [132.7, 150.26]. The null hypothesis is not rejected since 141.5 is within the predicted set of outcomes.

and is not statistically significant from the predicted price of 100.¹⁹ Likewise, the estimated trade price of certified trades varies between 194 and 198 in the two treatments and is not significantly different from the predicted value of 200 in either specification.²⁰

Table 6: Hypothesis 1: Convergence of *Pre* Treatments to the Unmediated or Mediated Equilibrium

	(1)	(2)
Certification (β_{Cert})	91.414*** (2.968)	91.414*** (2.970)
Treatment \mathcal{S}^{Pre} ($\beta_{\mathcal{S}^{Pre}}$)	39.100*** (8.105)	41.82*** (5.96)
Number of Lottery Averse Buyers in \mathcal{S}^{Pre} (β_{LA})		-24.887* (10.940) ^a
Constant (α_0)	102.401*** (3.500)	107.973*** (5.035)
Fixed Effects ^b	Yes	Yes
Adj. R^2	0.841	0.852
Observations (Trades in Period 7-12)	834	834

^aSince aversion to lotteries is an aggregate measure in specification (2) and there is serial correlation in prices, the standard error from the trade-level regression may be biased. As a better measure, randomization inference is used to construct a confidence interval. We begin by estimating the session-level regression $AvgP_s = \alpha_0 + \beta_{LA}(LA_s)$. We then take every permutation of possible assignments to construct placebo estimates of the lottery aversion parameter. This generates a distribution of possible parameters centered at zero. The empirically estimated value of β_{LA} lies outside the 90% confidence of this placebo distribution. See Bertrand, Duflo, and Mullainathan (2004)

^bFixed effects are at the session level. Robust standard errors in parenthesis clustered at the session level. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

4.2 Does utilization of the private-order CEI respond to changes in the environment?

4.2.1 Hypothesis and Empirical Strategy

Having established that the mediated equilibrium forms in all 6 markets that start in the Hazardous environment and a unmediated or partially-mediated equilibrium is selected in 5 out of 6 markets that start in the Safe environment, we next look at how the equilibrium that formed in the initial 12 periods adapts to changes in the underlying environment. In the theoretical model, we showed that when the mediated equilibrium is reached, there is no aggregate information observable when type-*C* sellers are replaced with type-*G* sellers. Thus the mediated equilibrium is predicted to persist even when it is no longer efficient. By contrast, when the unmediated equilibrium is reached, a replacement of type-*G* sellers with type-*C* sellers leads to a reduction in the uncertified price and an eventual change to the mediated equilibrium. This leads to:

¹⁹Significance based on a Wald test of $\alpha_0 = 100$. p -value = .6148 for regression (1) and p -value = .1574 for regression (2).

²⁰Significance based on a Wald test of $\alpha_0 + \beta_{Cert} = 200$. p -value = .1103 for regression (1) and p -value = .8902 for regression (2).

Hypothesis 2 *Any market equilibrium that reaches the mediated equilibrium will remain in this market equilibrium for any changes in the number of type-C and type-G sellers.*

This hypothesis is tested by comparing the price of uncertified trades that occur in the last six trading periods of each treatment. If there is no aggregate information observable when the environment changes from Hazardous to Safe, equilibrium prices in periods under the \mathcal{S}^{Post} treatment should be the same as those from \mathcal{H}^{Pre} and significantly differ from those in \mathcal{S}^{Pre} . We thus estimate:

$$P_{i,s} = \alpha_0 + \Sigma\alpha_s + \beta_{LA}(LA - \overline{LA}) * I_{\mathcal{S}^{Pre}} + \beta_{Cert}I_{Cert} + \beta_{\mathcal{S}^{Pre}}I_{\mathcal{S}^{Pre}} + \beta_{\mathcal{S}^{Post}}I_{\mathcal{S}^{Post}} + \beta_{\mathcal{H}^{Post}}I_{\mathcal{H}^{Post}} + \epsilon_{i,s}, \quad (9)$$

where $P_{i,s}$ is the price of an individual trade i in session s , α_s are individual session fixed effects, $(LA - \overline{LA}) * I_{\mathcal{S}^{Pre}}$ is the demeaned number of buyers who are lottery averse in the Safe treatment, I_{Cert} is an indicator for a certified trade, and $I_{\mathcal{S}^{Pre}}$, $I_{\mathcal{S}^{Post}}$, and $I_{\mathcal{H}^{Post}}$ are indicator variables for uncertified trades in their respective environment. We predict that $\alpha_0 + \beta_{\mathcal{S}^{Pre}} \in [140, 183]$, and $\beta_{\mathcal{S}^{Post}} = \beta_{\mathcal{H}^{Post}} = 0$.

4.2.2 Results

Result 2 *Consistent with Hypothesis 2, the price of uncertified trades in the \mathcal{S}^{Post} environment is not statistically significant to those in the \mathcal{H}^{Pre} environment and are consistent with prices predicted in the mediated equilibrium. The prices of uncertified trades in the \mathcal{S}^{Post} are also significantly below the prices observed in the \mathcal{S}^{Pre} environment and significantly below the prices which are predicted in the unmediated and partially-mediated equilibria.*

The persistence of the mediated equilibrium is most easily seen by comparing an individual session that began in the Safe environment to one that began in the Hazardous environment. Figure 1 makes this comparison, showing the complete trade history of session 6 and session 12. The horizontal dashed lines show the predicted price of the certified and uncertified market in the case of the unmediated equilibrium for the \mathcal{S}^{Pre} environment and the mediated equilibrium in the case of the other three environments. The vertical dashed lines split trades into six-period increments with the aggregate number of certified and uncertified trades reported at the bottom of each block. Note that in the Safe environment, there is always a single type- B seller. Thus the predicted composition of units without loss aversion is 60 uncertified high-quality units and 12 uncertified low-quality units in the mediated equilibrium. The unmediated and partially-mediated equilibria do not have a unique trade composition prediction but do require that some high-quality units be traded in the uncertified market.

As can be seen in the top half of Figure 3, a session that begins in the Safe environment converges to the partially-mediated equilibria in the first 12 periods and then adapts to the mediated equilibrium when the environment changes. Typical of all sessions that began in the Safe envi-

Session 6: Formation of the Unmediated Equilibrium and Adaptation to the Mediated Equilibrium





Figure 3: Hypothesis 2 — Persistence of the Mediated Equilibrium

Session 12: Formation and Persistence of the Mediated Equilibrium



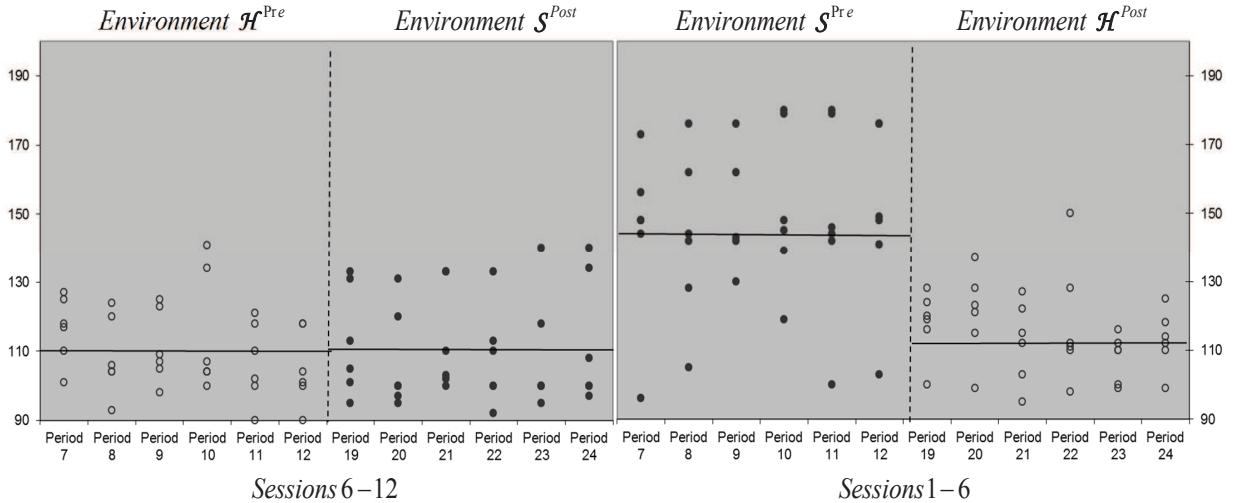
• Certified • Uncertified Low ◊ Uncertified High

ronment, the uncertified price converges from below to a partially-mediated equilibrium, with a subset of certified trades conducted in each period at a premium 60 points above the prevailing uncertified market price. When the environment changes, sellers who switched from type G to type C sell low-quality units leading to a decrease in price and the eventual establishment of a mediated equilibrium.

In the session that began in the Hazardous environment, the mediated equilibrium is established in the first 12 periods. When the environment switches to Safe at period 13, there is no noticeable change in the uncertified price nor in the composition of certified and uncertified trades. This is the case in the bottom half of Figure 3 where convergence to the mediated equilibrium is rapid and the convergence of the uncertified price is from above.

The patterns of adaption and persistence evident in this example is typical of most of the sessions.²¹ Figure 4 shows average uncertified prices for the last six periods of each environment. The uncertified price in the \mathcal{S}^{Post} environments is not significantly different from the \mathcal{H}^{Pre} and \mathcal{H}^{Post} treatments and significantly different from the \mathcal{S}^{Pre} treatment based on a Mann-Whitney-Wilcoxon test of the average price of uncertified trades in the last six periods of each treatment (\mathcal{S}^{Post} vs \mathcal{H}^{Pre} : $z = 0.00$, p -value = 1.00; \mathcal{S}^{Post} vs \mathcal{H}^{Post} : $z = -.64$, p -value = .52; \mathcal{S}^{Post} vs \mathcal{S}^{Pre} : $z = -2.40$, p -value = .016).

Figure 4: Average Uncertified Prices in Main Treatments



Turning to the price regression developed in equation 9, Table 7 extends the original regressions to include periods 18-24 of each session. In support of Hypothesis 2, there is no significant difference between the uncertified prices in the \mathcal{S}^{Post} and \mathcal{H}^{Post} environments relative to the baseline environment of \mathcal{H}^{Pre} . Further, the prices in \mathcal{S}^{Post} are significantly lower than those predicted in an unmediated or partially-mediated equilibrium based on a Wald test of $\alpha_0 + \beta_{\mathcal{S}^{Post}} = 140$

²¹As noted in the previous section, one of the six markets that began in the Safe environment had the certifying equilibrium form. One of the six markets that began in the Hazardous environment did not appear to converge in the first 12 periods and has a small number of high-quality uncertified trades in the second 12 periods.

(p -value < .01 for regression (1) and p -value < .01 for regression (2)).

Table 7: Hypothesis 2: Persistence of the Mediated Equilibrium

	(1)	(2)
Certification (β_{Cert})	89.229*** (2.566)	89.229*** (2.567)
Treatment \mathcal{S}^{Pre} ($\beta_{\mathcal{S}^{Pre}}$)	36.760*** (7.526)	37.024*** (6.397)
Treatment \mathcal{S}^{Post} ($\beta_{\mathcal{S}^{Post}}$)	2.323 (3.655)	2.323 (3.656)
Treatment \mathcal{H}^{Post} ($\beta_{\mathcal{H}^{Post}}$)	3.291 (4.199)	3.151 (4.107)
Number of Lottery Averse Buyers in \mathcal{S}^{Pre} (β_{LA})		-21.027* (10.654)
Constant (α_0)	107.109*** (3.715)	110.314*** (3.974)
Fixed Effects ^a	Yes	Yes
Adj. R^2	0.863	0.869
Observations	1675	1675

^aFixed effects are at the session level. Robust standard errors in parenthesis clustered at the session level. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

While we have thus far looked at the price data and shown that there is no observable difference in aggregate prices when the environment changes, a second prediction of the model is that individuals also cannot learn from their own experience since type- G sellers continue to trade certified goods and the uncertified market is full of only low-quality units. To see whether this prediction also holds, we next look at the composition of trades over time in each of the two treatment orderings. In the treatments that began in the Safe environment, the switch to the Hazardous environment should lead to an initial shift of units from uncertified high-quality units to uncertified low-quality units followed by a gradual transition to certified trades as the uncertified market price falls. In sessions that began in the Hazardous environment, theory would predict no change in the composition of goods when moral hazard is decreased.

Result 3 *Consistent with Hypothesis 2, there is little improvement in the quality of goods traded in the uncertified market when sessions that begin in the Hazardous environment are switched to the Safe environment.*


Figure 5 show the average number of certified and uncertified trades in treatments that start in the Safe environment and the Hazardous environment. Apparent in panel (a), the change in environment from Safe to Hazardous results in an immediate shift from uncertified high-quality units to uncertified low-quality units. Over time, uncertified low-quality units are replaced with

certified high-quality units leading to the mediated equilibrium in all sessions.²²


As shown in panel (b), the only significant change in the composition of trades for sessions that began in the Hazardous environment is a shift away from uncertified low-quality units to certified units.²³ This is most likely a result of weaker incentives for type- G sellers to trade uncertified units relative to sellers of type- C .

Figure 5: Changes in the composition of trades in response to changes in the environment

(a) Treatments Beginning in the Safe Environment



(b) Treatments Beginning in the Hazardous Environment



4.3 Are individuals learning from Public Signals or Private Experience?

Thus far we have looked at the aggregate data and seen that the intuition from the benchmark model closely matches the patterns of the observed market data. In this section, we take a more

²²There is also a small but consistent shift of transactions from certified high-quality units to uncertified low-quality units in the two periods following the change in treatment. Recall that in the partially-mediated equilibrium, it may be the case that the type- G sellers are indifferent between trading in the certified and uncertified markets while type- C sellers strictly prefer to sell uncertified units. Given a replacement of type- G sellers with type- C sellers, there is an increase in incentives to sell uncertified units. This effect may increase the speed of adaptation by increasing the number of uncertified low-quality units observed in the market.

²³Significance based on a probit regression, where the number of certified trades is the dependent variable and the treatment variable is the independent variable. p -value $< .01$ with errors clustered at session level. A similar regression with uncertified high-quality units as the dependent variable does not yield a significant treatment effect (p -value = .117).

exploratory look at the actions individual buyers and document evidence of individual learning from both publicly observed market signals and private experience.²⁴

Result 4 *There is evidence that buyers learn both from publicly observed market primitives and from their personal purchase experiences in markets where the unmediated or partially-mediated equilibrium has formed. There is little evidence of learning in environments where the mediated equilibrium has formed.*

In order to study the impact of market price on purchase decisions, we first generate a Markov transition matrix between (i) actions likely to be taken by individuals with optimistic beliefs about the trade environment and (ii) actions likely to be taken by individuals with pessimistic beliefs about the trade environment. We classify a trade as being made by a buyer with optimistic beliefs if the trade would produce a negative return in the event of a low-quality unit being supplied. These “Risky” trades are those made in the uncertified market where the price is greater than the buyer’s value. “Safe” trades are classified as those made in the certified market or trades made in the uncertified market where a profit is guaranteed. This is the case in the mediated equilibrium where the price of uncertified trades is equal to the marginal buyer’s valuation.

If price in the market is informative, the Markov transition matrix should have greater switching from Safe trades to Risky trades when the market price for uncertified trades is high and thus the gap between certified and uncertified prices, ΔP , is small. To study this conjecture, we generate two Markov transition matrices: one for trades where ΔP is less than the certification cost and one where the reverse is true. Table 8 shows these two Markov transition matrices over all treatments. As can be seen, when the difference in price is less than the certification cost, individuals who last made a Safe trade have a 23.5% chance of making a Risky trade. For individuals in an environment where this difference is greater than the certification cost, the likelihood of purchasing a Risky asset is only 7.7%. This difference is significant based on a probit regression which looks at the riskiness of the next trade of the same individual following a safe trade with an indicator variable for trades where the difference in average price of other uncertified and certified trades is less than the certification cost (p -value $< .01$; Errors clustered at the individual level).

Likewise, individuals who last purchased a Risky trade have a 74.7% chance of continuing to purchase a Risky asset in the next period when the price difference is small, while they have only a 49.5% chance of making another Risky trade when the price difference is large. This difference is also significant based on a probit regression which looks at the riskiness of the next trade of an individual following a risky trade with an indicator variable for trades where the difference in average price of other certified and uncertified trades is less than the certification cost (p -value $< .01$; Errors clustered at the individual level).

²⁴We initially planned to use the beliefs data here. However, in exit surveys, buyers reported that they were confused about the number of units sellers could trade and the relationship between the number of type- G sellers and overall risk. As confusion may be correlated with initial experiences that vary by treatment, the beliefs data has the potential for both classical and non-classical measurement error. Analysis of the beliefs data provides weak support for the theory model with no change in buyers’ beliefs in the \mathcal{S}^{Post} treatment relative to \mathcal{H}^{Pre} and a small but significant decrease in buyers’ beliefs in the \mathcal{S}^{Post} treatment relative to \mathcal{H}^{Pre} .

Table 8: Markov Transition Matrices Between Safe and Risky Trades as a Function of Prices

Markov Transition Matrix when Difference in Certified and Uncertified Prices is Less than Certification Cost			Markov Transition Matrix when Difference in Certified and Uncertified Prices is Greater than Certification Cost		
	Safe	Risky		Safe	Risky
Safe	.755	.235	Safe	.923	.077
Risky	.253	.747	Risky	.505	.495

In addition to the role of observable market prices, our data also suggests that an individual's trade experience also plays a role in his belief formation. In periods where ΔP is less than the certification cost, an individual who made a Risky trade in the previous period is 20.8% more likely to trade again if they receive a high-quality uncertified good instead of a low-quality uncertified unit. This difference is significant based on a probit regression where the left hand side is 1 if a risky trade is made and 0 otherwise, and the right hand side includes the quality of the last risky trade and a dummy variable for the information treatment (p -value $< .01$).²⁵ Likewise, individuals who make a risky trade when $\Delta P > 60$ are 14.2% more likely to make another risky trade if they receive a high-quality unit (p -value = .098.).

Finally, there is evidence that individuals learn from the composition of trades when the partially-mediated equilibria forms. The left hand side of Figure 6 shows the proportion of risky trades in the Public and Private Information treatments of the \mathcal{S}^{Pre} environment. As can be seen, individuals who are willing to accept actuarially fair gambles dramatically increase the proportion of risky trades they are willing to take, strongly suggesting that they are learning from the composition of trades. By contrast, when the mediated equilibrium forms, as is the case in the \mathcal{S}^{Post} , the information treatment appears to reduce individuals' propensity to experiment and decrease the number of risky trades which occur in the economy.

4.4 Is the Persistence of the Mediated Equilibrium Due to Equilibrium Selection or Information?

While the Safe Environment was designed so that only a single pair of buyers and sellers needed to coordinate away from the mediated equilibrium, it is possible that our results from the initial experiments are due to the fact that the mediated equilibrium is an equilibrium even in the Safe environment and that once selected, individuals will continue to play this equilibrium when the environment changes. In order to differentiate between this equilibrium selection channel and the information channel, we ran an additional 6 sessions that began in the Hazardous environment and switched to the Safe environment, but where individuals were informed about the distribution of seller types at the end of each period.

²⁵Only observations where the last trade was risky and where the difference in average price of other trades is lower than the certification cost are included. Errors clustered at the individual level

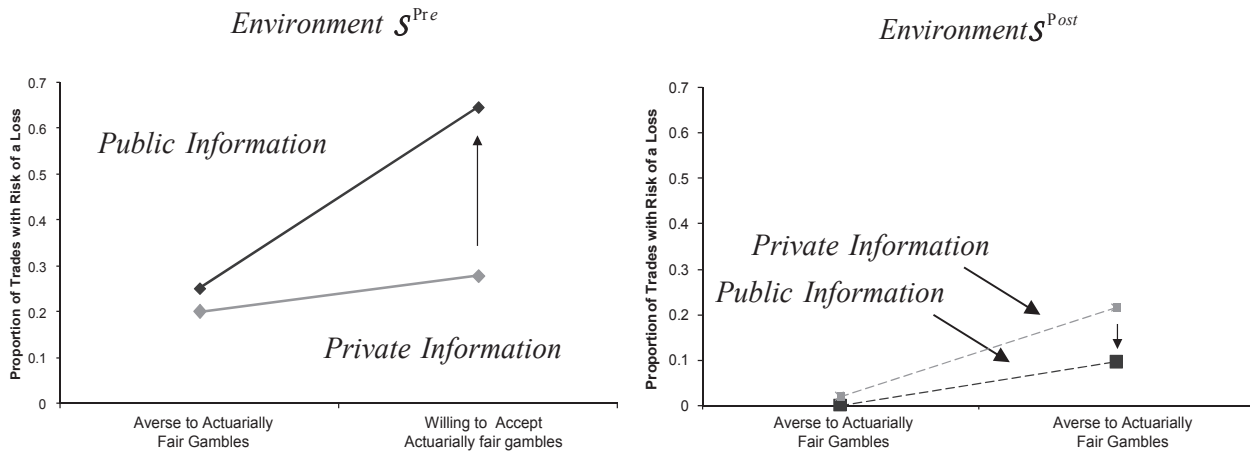


Figure 6: Proportion of risky trades in the Public information and No Information treatments.

The Full Information Treatment was identical to Treatment 3 of our original experiments, with the exception of additional information displayed on a results screen that occurred at the end of each period after individuals finished the bonus game. In the original experiments, subjects were given only their profit from trade in the previous period on this results screen. In the new full information treatment, the results screen also displayed an individual's guessed number of type- G sellers, the true number of type- G sellers, and the amount the individual earned in the bonus game.

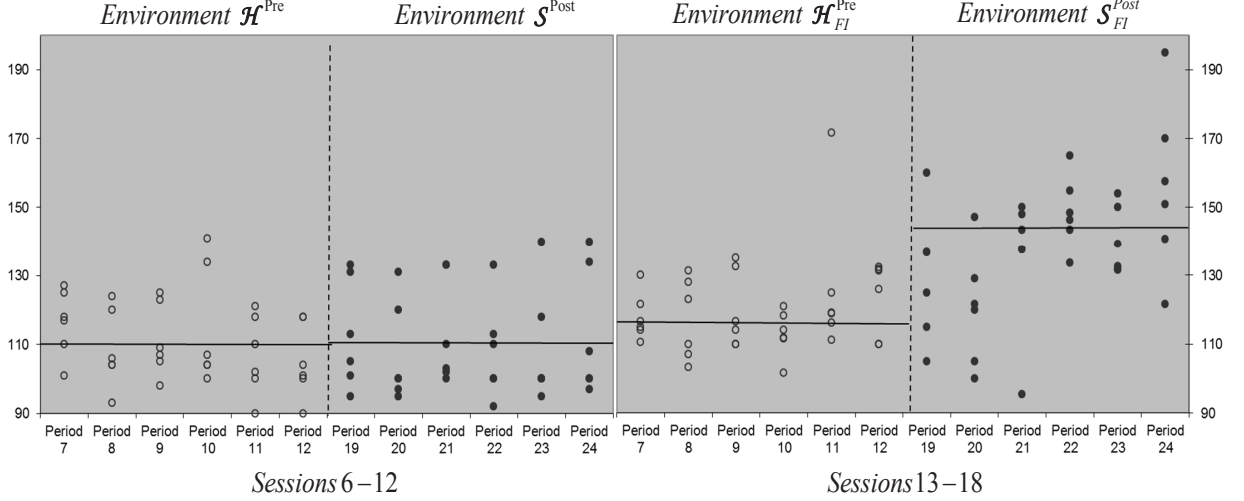
If equilibrium selection is the main driver of our initial results, buyers and sellers in this treatment should stay in the mediated equilibrium regardless of the additional information. However, if information also has an effect, we predict a shift away from the mediated equilibrium and toward the partially-mediated and unmediated equilibria.

Result 5 *In the Full Information treatments, all sessions initially converged to the mediated equilibrium in the Hazardous environment and converge to a partially-mediated equilibrium in the subsequent Safe environment. The total number of buyers and sellers trading certified units remains high, however, suggesting that only a subset of buyers are willing to adapt away from their initial actions.*

Figure 7 shows the average uncertified price of the last six periods of each treatment for sessions that began in the Hazardous environment in the original sessions (left) and the follow-up Full Information Treatments (right). In the initial Hazardous environments, uncertified prices are not significantly different in these two treatments based on a Mann-Whitney-Wilcoxon test using the average uncertified price in the last 6 periods of each treatment as an observation ($z = 1.60$, p -value = .11). In the subsequent Safe environment, however, uncertified prices are significantly higher in \mathcal{S}_{FI}^{Post} than \mathcal{S}^{Post} using the same test ($z = 2.72$, p -value = .01). The uncertified price in \mathcal{S}_{FI}^{Post} in the last 6 periods is also not significantly different to 140 in all six sessions using a standard t-test, suggesting that in all sessions a partially mediated equilibrium was being played in these periods.²⁶

²⁶Session 13: $t = -0.49$; Session 14: $t = 1.52$; Session 15: $t = .84$; Session 16: $t = -1.44$; Session 17: $t = 1.02$; Session 18: $t = .157$.

Figure 7: Average Uncertified Prices in Main Treatments



Looking again at the last six periods of each treatment, \mathcal{S}_{FI}^{Post} had an average of 4.11 uncertified trades per period. This level of uncertified trade is low and suggests that there is at least some hysteresis that is being generated from experience. However, it is not significantly different to the average number of uncertified trades in the last six periods of \mathcal{S}^{Pre} where there was an average of 6.06 uncertified trades per period (Mann-Whitney-Wilcoxon Test: $z = -1.537$, p -value = .12). These sessions also do not differ in their average uncertified prices based on a Mann-Whitney-Wilcoxon test ($z = -.641$, p -value = .52).

Taken together, the data in the full information treatment suggests that information is an important driver of equilibrium selection in our environment. However, the small amount of uncertified trades suggests that there may be some history dependence that is a result of forming a mediated equilibrium in prior play.

4.5 Can Information lead to the utilization of Private-Order CEIs that are useless?

An interesting question that arises from our results thus far is whether the information externality documented above might be powerful enough to sustain a private-order CEI even if that institution serves no purpose. We explore this idea in a series of follow-up experiments where we gradually eliminate the moral hazard in the environment and study how behavior adapts.

In our follow up No Moral Hazard Treatments, subjects initially start in an environment similar to our Hazardous environment, with a single type- G seller, 3 type- C sellers, and 2 type- B sellers. Unlike our original experiment where the type- B sellers always existed, however, we slowly switch type- B sellers to type- C sellers. We then slowly replace type- C sellers with type- G sellers until all sellers are type- G and where no seller has an incentive to produce low-quality units. These transitions are shown in table 9.

Table 9: Number of Good, Conditional, and Bad Sellers in No Moral Hazard Treatments

	Good	Conditional	Bad
Period 1-2	1	3	2
Period 3-4	1	4	1
Period 5-6	1	5	0
Period 7-8	2	4	0
Period 9-10	3	3	0
Period 11-12	4	2	0
Period 13-24	6	0	0

The goal of the No Moral Hazard Treatments was to induce a mediated equilibrium in periods 7-12 and then to study adaptation in periods 13-24 after this equilibrium formed but where the environment had no moral hazard and where the private-order CEI serves no purpose. The type-*B* sellers were introduced in the early periods as sellers of this type tended to trade uncertified low-quality units at low prices and thus were expected to have the biggest impact on changing buyer’s beliefs and pushing the market toward the mediated equilibrium. We staggered the switch points of the sellers from condition to good as this limited information that sellers might have about each other at any point in time.

Recall that in the mediated equilibrium, both type-*G* and type-*C* sellers will be selling certified high-quality units. This implies that if the mediated equilibrium occurs, all trades will be certified and there will be no way to distinguish between type-*G* and type-*C* sellers. We hypothesize that the information externality can sustain the mediated equilibrium in periods 13-24 when no exogenous information is given since the shift from type-*C* sellers to type-*G* sellers is not observable from market primitives. By contrast, under full information, while the mediated equilibrium technically exists as a rational expectation equilibrium, it is strictly dominated by the unmediated equilibrium for all buyers and sellers and thus we predict that it will be eliminated in sessions where exogenous information about the distribution of seller types is given.

We ran 12 additional sessions of the follow-up No-Moral Hazard treatments. Six of these sessions were under our original no information treatment and are identified as the sessions with “No Information and No Moral Hazard” in Periods 13-24 . The remaining six sessions provided information about the distribution of seller types through the results page of the bonus game and are identified as “Full Information and No Moral Hazard.”

Result 6 *2 of 6 No Information and No Moral Hazard Sessions converge to the mediated equilibrium in periods 7-12 and remain in this mediated equilibrium in periods 19-24. Averaging across all 6 of these sessions, 66.1% of trades in periods 19-24 use the private-order CEI. In the Full Information and No Moral Hazard treatment, all sessions converge to the mediated equilibrium in periods 7-12 and converged to the unmediated or partially-mediated equilibrium in periods 19-24. Averaging across all 6 of these sessions, only 20.6% of trades in periods 19-24 use the private-order CEI.*

As we have no type- B seller in this environment, not all sessions had uncertified trade in each period and it is difficult to perform analysis based on uncertified prices as was done in previous sections. We instead study the composition of trades over time, which is predicted to change in the sessions with full information and is predicted to be unchanged in the sessions with no information.

Figure 8 shows the number of certified high-quality units, uncertified high-quality units, and uncertified low-quality units traded in periods 7 - 24 of the information and no information treatments. The grey bar graph shows the average number of trades while the black dashes show the number traded in each individual session.

As can be seen looking at the left hand side of each figure, the distribution of trades in periods 7-12 are similar in the noise and no-noise treatments. There are an average of 8.08 certified high-quality units in the No Information treatment and an average of 8.38 certified high-quality units in the Full Information treatment, a difference that is not significant based on a Mann-Whitney-Wilcoxon test with period 7-12 trades averaged at the session level ($z = .161$, p -value = .87). Sessions without information had an average of 3.805 uncertified trades (2.94 low-quality units and 0.86 high quality units) while sessions with information had an average of 3.05 uncertified trades (1.97 low-quality units and 1.08 high-quality units). The difference in the number of uncertified trades is also not significantly different across treatments (All uncertified: $z = -.481$, p -value = .63; low-quality units: $z = -1.604$, p -value = 0.11; high-quality units: $z = .404$, p -value = .68).


In periods 13-24, by contrast, the distribution of trades between the No Information and Information treatments differ quite strongly. Looking first at the no information sessions, the average number of certified trades in periods 19-24 is large (7.86) and not significantly different to the number of certified trades in periods 7-12 ($z = -.241$, p -value = .81). Concentrating on periods 19-24 of the No Information treatment and looking across sessions, there is a large dispersion in the composition of trades across sessions. In two sessions, buyers and sellers use the private-order CEI for almost every transaction (69 out of 72 transactions in session 19 and 66 out of 72 transactions in session 22), while in another treatment, the private-order CEI in only 33.8% of transactions. Overall 66.1% of trades occur with the private-order CEI, suggesting that this institution was well utilized on average.

In the Full Information treatment, by contrast, there is a steady increase in the number of uncertified high-quality trades in each of the sessions. The average number of certified trades in period 19-24 is only 2.47 and is falling over time in all sessions. This average number of certified trades is significantly different to the number of trades in the same periods of the No Information sessions and is also significantly different to the number of certified trades in periods 7-12 of the same session (Full Information vs. No Information: $z = -2.89$, p -value < .01; Periods 19-24 vs. Periods 7-12: $z = -2.89$, p -value < .01). In three of the sessions, all trades occur without the certification technology in periods 24; the private-order CEI is used in only 20.6% of period 19-24 transactions across all sessions suggesting that the institution is being abandoned.


Taken together, there is strong evidence that information is inhibiting the elimination of the private-order CEI in the No Information sessions. Even when the environment has no moral hazard,

Figure 8: Changes in the composition of trades in response to changes in information

(a) Sessions with No Information and No Moral Hazard in Periods 13-24



(b) Treatments with Full Information and No Moral Hazard in Periods 13-24



■ Average • Individual Session

individuals continue to use the private-order CEI in a subset of the sessions and cannot observe the fact that the institution is now useless. By contrast, when exogenous information is introduced, all markets adapt away from the private-order CEI and efficiency is improved.

5 Conclusion

This paper represents a first step in studying the relationship between the utilization of private-order CEIs and information. We demonstrated that, in a market where a private-order CEI becomes utilized, observable information about changes in the underlying environment could be lost. This lost information could lead to the persistence of an equilibrium where all participants in the environment are weakly worse off relative to a world without the private-order CEI. In laboratory experiments, the inefficient persistence of the mediated equilibrium was striking. Without exception, markets that utilized the private-order CEI and formed a mediated equilibrium failed to respond to a change in the underlying distribution of seller types. This failure to adjust led to efficiency losses when compared to sessions where participants were unhindered by the early adoption of the private-order CEI. In follow up sessions where information was given exogenously, adoption to the more efficient partially-mediated equilibria occurred even after the formation of the mediated equilibrium, suggesting that the information channel was important to the adaptation process.

The information externality highlighted in this paper suggests a general phenomenon that may extend beyond the certification private-order CEI considered here. Common private-order CEIs designed to mitigate moral hazard such as regulation, certification, monitoring, process management, and credit scoring all share the common characteristic that they group heterogeneous agents into the same action. Given the ubiquity of these institutions in everyday markets and organizations, developing an understanding of how information externalities dynamically alter the institutional landscape is of great importance.

References

- ACEMOGLU, D., AND J. A. ROBINSON (2006): “De Facto Political Power and Institutional Persistence,” *American Economic Review*, 96(2), 326–330.
- (2008): “Persistence of Power, Elites and Institutions,” *American Economic Review*, 98(1), 267–293.
- AOKI, M. (2001): *Towards a Comparative Institutional Analysis*. MIT Press, Cambridge, MA.
- BERTRAND, M., E. DUFLO, AND S. MULLAINATHAN (2004): “How Much Should We Trust Differences-in-Differences Estimates?,” *The Quarterly Journal of Economics*, 119(1), 249–275.
- BLUME, A., AND A. ORTMANN (2007): “The effects of costless pre-play communication: Experimental evidence from games with Pareto-ranked equilibria,” *Journal of Economic Theory*, 131(1), 274–290.

- BRAINARD, L. S., AND T. VERDIER (1994): "Lobbying and Adjustment in Declining Industries," *European Economic Review*, 38(3), 568–595.
- BRANDTS, AND HOLT (1992): "An Experimental Test of Equilibrium Dominance in Signaling Games," *American Economic Review*, 82(5), 1350–1365.
- CACHON, G., AND C. CAMERER (1996): "Sunk Costs and Forward Induction in Experimental Coordination Games," *The Quarterly Journal of Economics*, 111(1), 165–194.
- CALVERT, R. (1995): "Rational Actors, Equilibrium, and Social Institutions," in *Explaining Social Institutions*, ed. by J. Knight, and I. Sened. University of Michigan Press.
- COATE, S., AND S. MORRIS (1999): "Policy Persistence," *American Economic Review*, 89(5), 1327–1336.
- COOPER, D., AND J. KAGEL (2003): "The Impact of Meaningful Context on Strategic Play in Signalling Game," *Journal of Economic Behavior and Organization*, 50, 311–337.
- (2005): "Are Two Heads Better Than One? Team Versus Individual Play in Signaling Games," *American Economic Review*, 95, 477–509.
- (2008): "Learning and Transfer in Signaling Games," *Economic Theory*, 34, 415–439.
- COOPER, D. J., S. GARVIN, AND J. H. KAGEL (1997a): "Adaptive Learning vs. Equilibrium-Refinements in an Entry Limit Pricing Game," *The Economic Journal*, 107, 553–575.
- (1997b): "Signaling and Adaptive Learning in an Entry Limit Pricing Game," *RAND Journal of Economics*, 28, 662–683.
- COOPER, R., D. DEJONG, R. FORSYTHE, AND T. ROSS (1990): "Selection Criteria in Coordination Games: Some Experimental Results," *American Economic Review*, 80.
- DIXIT, A. K. (2004): *Lawlessness and Economics: Alternative Modes of Governance*. Princeton University Press, Princeton, NJ.
- FAFCHAMPS, M. (2004): *Market Institutions in Sub-Saharan Africa: Theory and Evidence*. MIT Press, Cambridge.
- FERNANDEZ, R., AND D. RODRIK (1991): "Resistance to Reform: Status Quo Bias in the Persistence of Individual-Specific Uncertainty," *American Economic Review*, 81(5), 1146–1155.
- FISCHBACHER, U. (2007): "z-Tree: Zurich Toolbox for Ready-made Economic Experiments," *Experimental Economics*, 10(2), 171–178.
- FOSTER, D., AND P. YOUNG (1990): "Stochastic Evolutionary Game Dynamics," *Theoretical Population Dynamics*, 38(1), 219–232.
- FRIEDRICH, S. (2013): "Policy Persistence and Rent Extraction," *Public Choice*, 157(1–2), 287–304.
- GALE, D. (1992): "A Walrasian Theory of Markets with Adverse Selection," *Review of Economic Studies*, 59(2), 229–255.
- GERSCHENKRON, A. (1962): *Economic Backwardness in Historical Perspective: A Book of Essays*. The Belknap Press of Harvard University Press, Cambridge.
- GOETZMANN, W. N., AND E. KÖLL (2005): "The History of Corporate Ownership in China: State Patronage, Company Legislation, and the Issue of Control," in *A History of Corporate*

- Governance around the World: Family Business Groups to Professional Managers*, ed. by R. K. Morck, vol. 0-226-53680-7 of *NBER*, pp. 149–183. University of Chicago Press.
- GREIF, A. (1994): “Cultural Beliefs and the Organization of Society: A Historical and Theoretical Reflection on Collectivist and Individualist Societies,” *Journal of Political Economy*, 102(5), 912–950.
- (1998): “Historical and Comparative Institutional Analysis,” *American Economic Review*, 88(2), 80–84.
- (2000): “The Fundamental Problem of Exchange,” *Review of European Economic History*, 4(3), 251–284.
- (2002): “Institutions and Impersonal Exchange: From Comunal to Individual Responsibility,” *The Journal of Institutional and Theoretical Economics*, 158(1), 168–204.
- (2005): “Commitment, Coercion and Markets: The Nature and Dynamics of Institutions Supporting Exchange,” in *Handbook of New Institutional Economics*, ed. by C. Ménard, and M. Shirley, vol. 53, pp. 727–88. Springer.
- (2006): *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*. Cambridge University Press, Cambridge.
- GREIF, A., AND C. KINGSTON (2011): “Institutions: Rules or Equilibria,” in *Political Economy of Institutions, Democracy and Voting*, ed. by N. Schofield, and G. Caballero, pp. 13–43. Springer-Verlag.
- HELLWIG, M. (1987): “Some Recent Developments in the Theory of Competition in Markets with Adverse Selection,” *European Economic Review*, 1(2), 319–325.
- HOFFMAN, P. T., G. POSTEL-VINAY, AND J.-L. ROSENTHAL (2000): *Priceless Markets: The Political Economy of Credit in Paris, 1660-1870*. University of Chicago Press, Chicago.
- HOLT, C., L. LANGAN, AND A. P. VILLAMIL (1986): “Market Power in Oral Double Auctions,” *Economic Inquiry*, 24, 107.
- HOLT, C., AND S. LAURY (2002): “Risk Aversion and Incentive Effects,” *American Economic Review*, 5, 1644–1655.
- JEHIEL, P., AND A. NEWMAN (2014): “Loopholes and the Evolution of Contract Form,” mimeo.
- KANDORI, M. (1992): “Social Norms and Community Enforcement,” *Review of Economic Studies*, 59(1), 61–80.
- KINGSTON, C., AND G. CABALLERO (2011): “Comparing Theories of Institutional Change,” *Journal of Institutional Economics*, 5(1), 151–180.
- MENGER, C. (1963 [1883]): *Problems of Economics and Sociology*. Translated by F. J. Nock. University of Illinois Press, Urbana, Illinois.
- MILGROM, P., AND J. ROBERTS (1982): “Limit Pricing and Entry Under Incomplete Information: An Equilibrium Analysis,” *Econometrica*, 50, 443–459.
- NÖLDEKE, G., AND L. SAMUELSON (1997): “A Dynamic Model of Equilibrium Selection in Signaling Markets,” *Journal of Economic Theory*, 73, 118–156.
- NORTH, D. C. (1981): *Structure and Change in Economic History*. Norton, New York.

- RILEY, J. (1979): “Informational Equilibrium,” *Econometrica*, 47, 331–359.
- ROTHSCHILD, M., AND J. STIGLITZ (1976): “Equilibrium in Competitive Insurance Markets: an Essay on the Economics of Imperfect Information,” *The Quarterly Journal of Economics*, 90(2), 629–650.
- SCHOTTER, A. (1981): *The Economic Theory of Social Institutions*. Cambridge University Press, Cambridge.
- SMITH, V. (1964): “Effects of Market Organization on Competitive Equilibrium,” *The Quarterly Journal of Economics*, 78(2), 181–202.
- VAN HUYCK, J., R. BATTALIO, AND R. O. BEIL (1993): “Asset markets as an equilibrium selection mechanism: Coordination failure, game form auctions, and tacit communication,” *Games and Economic Behavior*, 5(3), 485–450.
- WARREN, P., AND T. WILKENING (2012): “Regulatory Fog: The Role of Information in Regulatory Persistence,” *Journal of Economic Behavior and Organization*, 84(3), 840–856.
- WITTFOGEL, K. (1957): *Oriental Despotism: A Comparative Study of Total Power*. Yale University Press, New Haven.
- YANG LI, M. (2002): *Essays on Public Finance and Economic Development in a Historical Institutional Perspective*. Ph.D. Dissertation, Stanford University.
- YOUNG, P. (1993): “The Evolution of Conventions,” *Econometrica*, 61(1), 57–84.

6 Appendix

6.1 Formal Construction of the Rational Expectations Equilibrium

In this section, we formally define the rational expectation equilibrium and develop the notation necessary for proving Propositions 1-4. Following Gale (1992), it is convenient to define an *interim* utility where an individual’s utility is a function of a match and market environment. A buyer of type $b \in \mathcal{B}$ who matches with a seller of type $s \in \{G, C, B\}$ in market $m \in \{\mathcal{C}, \mathcal{NC}, \emptyset\}$ at price P^m receives utility $u(m, P^m, b, s)$. The market affects this utility by restricting the set of actions that a seller can take. For instance, if a buyer matches with a type- C seller in market \mathcal{NC} , the conditional seller is free to exchange a unit of either high or low quality and optimally supplies a low-quality unit. If the buyer had matched with the same seller in market \mathcal{C} , the conditional seller is constrained and would supply a high-quality unit.

Buyers in our model are either risk and loss neutral, in which we denote their type as λ_0 , or loss averse with type corresponding to their loss aversion parameter λ_i . For a given type λ_i ,

$$u(m, P^m, \lambda_i, s) = \begin{cases} U^H - P^{\mathcal{C}} & \text{if } m \in \mathcal{C}, s \in \{G, C, B\} \\ U^H - P^{\mathcal{NC}} & \text{if } m \in \mathcal{NC}, s \in \{G\} \\ \lambda_i[U^L - P^{\mathcal{NC}}] & \text{if } m \in \mathcal{NC}, s \in \{C, B\}. \end{cases} \quad (10)$$

Similarly, a seller of type s who matches with a buyer of type b in market m at price P^m receives

utility $v(m, P^m, b, s)$. A seller maximizes expected value and thus, given optimal action in both markets, has a utility function of:

$$v(m, P^m, b, s) = \begin{cases} P^{\mathcal{C}} - C_s^H - T & \text{if } m \in \mathcal{C}, s \in \{G, C, B\}, \\ P^{\mathcal{NC}} - C_s^H & \text{if } m \in \mathcal{NC}, s \in \{G\}, \\ P^{\mathcal{NC}} - C^L & \text{if } m \in \mathcal{NC}, s \in \{C, B\}. \end{cases} \quad (11)$$

Note that the seller's value is independent of the buyer type in which she is matched. We leave the parameter b in the left hand side of equation 11 to be clear that both buyer and seller utility are defined over matches.

The description of the rational expectations equilibrium²⁷ is comprised of three parts: an attainable allocation (D, S) , a belief system μ , and a price system P .

Attainable Allocations: The number of buyers of type b who demand from market m is denoted by $D(m, b)$. An allocation of buyers is a function $D : \mathcal{M} \times B \rightarrow \mathbb{I}_+$ such that $\sum_{m \in \mathcal{M}} D(m, b) = N_b$. Likewise, the number of sellers of type $s \in \{G, C, B\}$ who supply in market m is denoted by $S(m, s)$. An allocation of sellers is a function $S : \mathcal{M} \times \{G, C, B\} \rightarrow \mathbb{I}_+$ such that $\sum_{m \in \mathcal{M}} S(m, s) = M_s$. An allocation (D, S) is *attainable* iff $\sum_{s \in \{G, C, B\}} S(m, s) = \sum_{b \in B} D(m, b)$ for $m \in \{\mathcal{C}, \mathcal{NC}\}$. Note that this market clearing condition is not binding in the \emptyset market.

Belief System: Buyers and sellers form beliefs about the types of agents exchanging within a market. Let $\mu_b(m, s)$ denote the subjective probability that a unit purchased in market m by a buyer is in fact supplied by a seller of type s . Let $\mu_s(m, b)$ denote the subjective probability that a unit sold in market m by a seller is in fact bought by a buyer of type b . A belief system is a pair of beliefs $\mu = (\mu_b, \mu_s)$ such that $\mu_b(m, s) : \mathcal{M} \times \{G, C, B\} \rightarrow \mathbb{R}_+$ satisfies $\sum_s \mu_b(m, s) = 1$ for every m and $\mu_s(m, b) : \mathcal{M} \times B \rightarrow \mathbb{R}_+$ satisfies $\sum_b \mu_s(m, b) = 1$ for every m .

Price System: A price system is a function $P : \mathcal{M} \rightarrow \mathbb{R}_+$. For convenience, we define $P^{\mathcal{C}}, P^{\mathcal{NC}}, P^{\emptyset}$ as the prices in each market.

Suppose that a buyer of type b purchases a unit in market m at price P^m . If the buyer's beliefs are given by $\mu_b(m, s)$, his expected utility is given by

$$\sum_s u(m, P^m, b, s) \mu_b(m, s), \quad (12)$$

where $u(m, P^m, b, s)$ is the utility received when a seller sells her market constrained optimal unit to the buyer. A buyer will choose a market that maximizes (12). Consequently, an equilibrium

²⁷This formulation is also defined as a price equilibrium, competitive equilibrium or information equilibrium depending on author. As it is most often discussed in relation to macroeconomic rational expectations models, the most common term is used here.

allocation must assign all buyers of type b to markets that are in the arg max of (12):

$$D(m^*, b) \neq 0 \Leftrightarrow m^* \in \arg \max_m \sum_s u(m, P^m, b, s) \mu_b(m, s) \quad \forall b. \quad (13)$$

Likewise, suppose that a seller sells a unit in market m at price P^m . If the seller's beliefs are given by $\mu_s(m, b)$ her expected utility is given by

$$\sum_b v(m, P^m, b, s) \mu_s(m, b), \quad (14)$$

where $v(m, P^m, b, s)$ is the value the seller receives from selling her optimal unit to a buyer of type b subject to the constraints of the market she has entered. Like the buyer, any rational expectations equilibrium requires:

$$S(m^*, s) \neq 0 \Leftrightarrow m^* \in \arg \max_m \sum_b v(m, P^m, b, s) \mu_s(m, b) \quad \forall s. \quad (15)$$

Finally, the rational expectations equilibrium requires that beliefs perfectly forecast the rational actions of others and are updated according to Bayes rule. For the sellers, where the distribution of buyer types is known, this simply requires that the belief that a unit in a market is bought by a buyer of type b is equal to the actual proportion of type- b buyers in the market.

For the buyers, who do not know the distribution of seller types, we require that the buyer forms expectation of matching with each seller type based on his (correct) beliefs about the actions of each type of sellers and his (potentially incorrect) posterior of the number of sellers of each type. This is done in three steps. For any market in which there are a positive number of sellers, a buyer evaluates the likelihood of each seller type being in each market given the prices. Given this evaluation and the number of sellers allocated to each market, the buyer next updates his prior about the distribution of seller types, ruling out seller distributions where the rational allocation of sellers could not generate the observed allocation. This will only occur in the partially-mediated equilibria where all trades in the certified market are made by type- G buyers. Finally, the buyer forms an expectation of matching with each seller type based on his (correct) beliefs about the actions of the sellers and his (potentially incorrect) posterior of the seller distribution. If a market has no trades in equilibrium, then these proportions are not well-defined and beliefs may be arbitrary.

As in the main text, we restrict attention to the case where there is exactly one type- B seller so that buyers' beliefs about the uncertified market are always well defined and the distribution of seller types can be expressed by the number of type- G sellers in the market. Define $S^{\mathcal{N}C}$ as the number of sellers trading in the uncertified market and S^C as the number of sellers trading in the certified market. Further define $p(\hat{g})$ and $q(\hat{g}|S^C, S^{\mathcal{N}C})$ as the prior and posterior distribution regarding the proportion of good types in the economy, which has support over $g \in \{0, \frac{1}{M}, \frac{2}{M}, \dots, \frac{M-1}{M}, \}$. Finally, let $\mathbb{E}_q \hat{S}(m, s|S^C, S^{\mathcal{N}C})$ be the expected number of sellers of type s in market m based on the posterior $q(\hat{g}|S^C, S^{\mathcal{N}C})$ and the assumption that all sellers behave rationally.

Definition 1 *Rational Expectations Equilibrium:* A Rational Expectations Equilibrium is a

triple $\langle (D \times S), \mu, P \rangle$ consisting of an attainable allocation $(D \times S)$, beliefs μ , and a price system P that satisfy:

$$E.1 : \quad S(m^*, s) \neq 0 \Leftrightarrow m^* \in \arg \max_m \Sigma_b v(m, P^m, b, s) \mu_s(m, b) \quad \forall s,$$

$$E.2 : \quad D(m^*, b) \neq 0 \Leftrightarrow m^* \in \arg \max_m \Sigma_s u(m, P^m, b, s) \mu_b(m, s) \quad \forall b,$$

$$E.3a : \quad \mu_b(m, s) = \frac{\mathbb{E}_q \hat{S}(m, s | S^C, S^{\mathcal{N}C})}{\Sigma_s \mathbb{E}_q \hat{S}(m, s | S^C, S^{\mathcal{N}C})} \quad \text{if } \mathbb{E}_q \hat{S}(m, s | S^C, S^{\mathcal{N}C}) > 0,$$

$$E.3b : \quad \mu_s(m, b) = \frac{D(m, b)}{\Sigma_b D(m, b)} \quad \text{if } \Sigma_b D(m, b) > 0.$$

Analysis of the rational expectation equilibria is simplified by two characteristics of the benchmark environment. First, the seller's valuation $v(m, P^m, b, s)$ is independent of the buyer that she is matched with and thus $\mu_s(m, b)$ does not affect the seller's decision. It follows that condition (E.1) can be reduced to

$$E.1b : \quad S(m^*, s) \neq 0 \Leftrightarrow m^* \in \arg \max_m \Sigma_b v(m, P^m, b, s) \quad \forall s,$$

which is the requirement that all sellers enter the market where the difference between price and the cost of their constrained optimal production choice is largest. Second, since all buyers share the same utility function given in equation (10), only beliefs about $\mu_b(\mathcal{N}C, G)$, the probability of matching with a type- G seller in the uncertified market, affect utility. Since seller's actions only depend on prices, we define a function $\pi^H(\Delta P, \mathbb{E}(\hat{g}))$ where $\pi^H : P \rightarrow [0, 1]$ is a buyer's belief about the proportion of high-quality units in the uncertified market for a difference in prices of $\Delta P \equiv P^C - P^{\mathcal{N}C}$. Note that $\pi^H(\Delta P, \mathbb{E}(\hat{g})) = \mu_b(\mathcal{N}C, G)$, which is given by:

$$\mu_b(\mathcal{N}C, G) = \begin{cases} \mathbb{E}(\hat{g}) & \text{if } \Delta P < T \\ \frac{M\mathbb{E}(\hat{g}|S^C) - S^C}{M - S^C} & \text{if } \Delta P = T \\ 0 & \text{if } \Delta P > T \end{cases} \quad (16)$$

The conditioning of $\mathbb{E}(\hat{g}|S^C)$ by S^C in the partially-mediated market is due to the fact that only type- G sellers are willing to certify their goods when $\Delta P = T$. Thus, observing S^C rules out some initial seller distributions that have less than S^C type- G sellers.

6.2 Proofs

Lemma 1 For a set of prices where $U^L \leq \underline{P}^{\mathcal{N}C} \leq P^{\mathcal{N}C} \leq P^C \leq \overline{P}^C \leq U^H$:

- A seller of type G has $C_G^H \leq C^L$ and will always produce high-quality units. A type- G seller will trade in the uncertified market if $\Delta P \leq T$.
- A seller of type C has $C_C^H \in (C^L, C^L + \overline{P}^C - \underline{P}^{\mathcal{N}C} - T)$ and will produce either low-quality units to the uncertified market or high-quality units to the certified market. A type- C seller will trade to the uncertified market if $\Delta P \leq T + (C_C^H - C^L)$.

- A seller of type B has $C_B^H \geq C_L + \overline{P}^C - \underline{P}^{NC} - T$. Given the bounds on possible prices, type-B sellers never sell high-quality units and will always produce low-quality units in the uncertified market.

Proof of Lemma 1: This lemma follows directly from the optimality condition in E1.b and the seller's utility function given in equation 11.

Lemma 2 *In Equilibrium:*

- If $\Delta P > T$ all buyers believe that all type-G sellers will certify their goods and thus that $\pi^H(\Delta P, \mathbb{E}(\hat{g})) = 0$. In this case, a buyer prefers to purchase the certified unit as long as $\Delta P < U^H - U^L \equiv \overline{P}^C - \underline{P}^{NC}$ and is indifferent between buying a non-certified unit and not purchasing if $P^{NC} = U^L$.
- If $\Delta P \leq T$ the buyers believe that all sellers trade in the uncertified market. In this case $\pi^H(\Delta P, \mathbb{E}(\hat{g})) = \mathbb{E}(\hat{g})$ and a risk neutral buyer prefers to purchase the uncertified unit as long as $\Delta P \geq (1 - \mathbb{E}(\hat{g}))(U^H - U^L)$.

Proof of Lemma 2: In the baseline model, there is only one type of buyer which we denoted as λ_0 whose utility is given as:

$$u(m, P^m, \lambda_0, s) = \begin{cases} U^H - P^C & \text{if } m \in \mathcal{C}, s \in \{G, C, B\} \\ U^H - P^{NC} & \text{if } m \in \mathcal{NC}, s \in \{G\} \\ U^L - P^{NC} & \text{if } m \in \mathcal{NC}, s \in \{C, B\}. \end{cases}$$

It follows:

1. When $\Delta P > T$, $v(\mathcal{C}, P^C, b, G) > v(\mathcal{NC}, P^{NC}, b, G)$ and thus $\mathbb{E}_q \hat{S}(\mathcal{NC}, G | S^C, S^{NC}) = 0$. By the definition of the rational expectations equilibrium, $\mu_b(\mathcal{NC}, G) = 0$ and thus

$$\sum_s u(\mathcal{NC}, P^{NC}, b_0, s) \mu_b(\mathcal{NC}, s) = U^L - P^{NC}.$$

Since $\forall s$, $u(\mathcal{C}, P^C, b_0, s) = U^H - P^C$ and $u(\emptyset, P^\emptyset, \lambda_0, s) = 0$, it follows that an agent is indifferent between all three markets when $P^{NC} = U^L$, $P^C = U^H$.

2. When $\Delta P \leq T$, $\forall s$, $v(\mathcal{C}, P^C, b, s) < v(\mathcal{NC}, P^{NC}, b, s)$ and thus $\mathbb{E}_q \hat{S}(\mathcal{NC}, G | S^C, S^{NC}) = M\mathbb{E}(\hat{g})$. By the definition of the rational expectations equilibrium, $\mu_b(\mathcal{NC}, G) = \mathbb{E}(\hat{g})$. It follows that

$$\sum_s u(\mathcal{NC}, P^{NC}, b_0, s) \mu_b(\mathcal{NC}, G) = \mathbb{E}(\hat{g})U^H + (1 - \mathbb{E}(\hat{g}))U^L - P^{NC}.$$

A buyer is indifferent across all three markets if $P^{NC} = U^H - (1 - \mathbb{E}(\hat{g}))(U^H - U^L)$ and $P^C = U^H$.

Proof of Proposition 1:

1. When $\Delta P = U^H - U^L$:

(a) By Lemma 1, $S(\mathcal{NC}, B) = 1$, $S(\mathcal{C}, G) = M\mathbb{E}(\hat{g})$, and $S(\mathcal{C}, C) = M(1 - \mathbb{E}(\hat{g})) - 1$.

(b) By Lemma 2, if $P^{\mathcal{NC}} = U^H$, $P^{\mathcal{C}} = U^L$, $D(\mathcal{C}, \lambda_0) = [0, N_{\lambda_0}] \in \mathbb{I}_+$, $D(\mathcal{NC}, \lambda_0) = [0, N_{\lambda_0}] \in \mathbb{I}_+$, $D(\emptyset, \lambda_0) = [0, N_{\lambda_0}] \in \mathbb{I}_+$ with $\sum_m D(m, \lambda_0) = N_{\lambda_0}$.

Thus the attainable allocation where $P^{\mathcal{NC}} = U^H$, $P^{\mathcal{C}} = U^L$, $D(\mathcal{C}, \lambda_0) = M - 1$, $D(\mathcal{NC}, \lambda_0) = 1$, and $D(\emptyset, \lambda_0) = N_{\lambda_0} - M$ always exists.

2. When $\Delta P > T$:

(a) By Lemma 1, $S(\mathcal{NC}, B) = 1$, $S(\mathcal{NC}, G) = M\mathbb{E}(\hat{g})$, and $S(\mathcal{NC}, C) = M(1 - \mathbb{E}(\hat{g})) - 1$.

(b) By Lemma 2, a buyer is indifferent between all three markets if $P^{\mathcal{NC}} = U^H - (1 - \mathbb{E}(\hat{g}))(U^H - U^L)$ and $P^{\mathcal{C}} = U^H$.

If $P^{\mathcal{C}} - P^{\mathcal{NC}} = (1 - \mathbb{E}(\hat{g}))(U^H - U^L) > T$, then $D(\mathcal{NC}, \lambda_0) = M$, $D(\emptyset, \lambda_0) = N_{\lambda_0} - M$ is an equilibrium. Otherwise, there does not exist a set of prices such that $\Delta P > T$ and a buyer is indifferent between the certified and uncertified markets.

Proof of Proposition 2: When a unmediated equilibrium exists, $P^{\mathcal{NC}} = U^H - (1 - \mathbb{E}(\hat{g}))(U^H - U^L)$. Thus

$$\mathbb{E}(\hat{g}) = \frac{P^{\mathcal{NC}} - U^L}{U^H - U^L} \quad (17)$$

and price is a sufficient statistic for $\mathbb{E}(\hat{g})$. Under the certifying equilibrium, both type- G and type- C individuals certify their product. As they are both in the same market, $P^{\mathcal{NC}} = U^L$ and $P^{\mathcal{C}} = U^H$, there is no new information regarding the relative proportions of type- G and type- C sellers. If the number of type- B sellers is unknown, they can be distinguished in the separating equilibrium as they are the only ones left in the uncertified market.

Proof of Proposition 3: Let $\mathbf{x} = (x_1, \dots, x_T)$ be observations of a single buyer trading in the uncertified market T times, where $x_i \in \{H, L\}$. As before, let $\hat{g} \in \{0, \frac{1}{M}, \dots, \frac{M-1}{M}\}$ be the possible number of type- G sellers in the market. Given an initial prior $p_0^i(\hat{g}) = \{p_0^i(\hat{g}_0), p_0^i(\hat{g}_1), \dots, p_0^i(\hat{g}_{M-1})\}$ where $p_0^i(\hat{g}_k) > 0$ and $\sum_k p_0^i(\hat{g}_k) = 1$, the posterior $p_t(\hat{g}|\mathbf{x})$ converges almost surely to the true proportion as $T \rightarrow \infty$ as long as $\mathbf{g} \in \hat{g}$ and

$$\sum_x q(x|\hat{g}_i) \log \left[\frac{q(x|\hat{g}_i)}{q(x|\hat{g}_j)} \right] > 0, \quad (18)$$

where $q(x|\hat{g}_i)$ is the posterior of receiving a good of quality x given the true parameter is \hat{g}_i .²⁸

²⁸The use of $q(x|\hat{g}_k)$ in this equation is to highlight that there is actually two steps taking place in updating the posterior over types. The first is an empirical update on the likelihood of getting a high-quality unit in the uncertified market. The second is mapping this empirical data back into implications about the proportion of type- G sellers in the environment under the assumption that sellers do not play dominated strategies.

Expanding condition (18) yields:

$$\hat{g}_i \log\left(\frac{\hat{g}_i}{\hat{g}_j}\right) + (1 - \hat{g}_i) \log\left(\frac{1 - \hat{g}_i}{1 - \hat{g}_j}\right). \quad (19)$$

Rewriting $\hat{g}_j = \hat{g}_i + z$ and taking the derivative with respect to z , the first derivative is zero at $z = 0$ and the second derivative is strictly positive for all z . Thus condition (18) holds. Since $\mathfrak{g} \in \{0, \frac{1}{M}, \dots, \frac{M-1}{M}\}$, convergence is guaranteed as $t \rightarrow \infty$.

Returning to the original problem, M buyers purchase each period. Thus, there must be at least M individuals whose individual observations T go to infinity as the number of periods goes to infinity.

Proof of Proposition 4: Since $M_B = 1$ is known, prices and the allocation of sellers to markets does not lead to updating by buyers. Further, buyers who purchase in the certified market get a high-quality unit by either a type- G or type- C seller while those in the uncertified market receive a low-quality unit by a type- B seller. Thus, individual experiences again yield no new information about the distribution of seller types.

6.3 Supplementary Material: Experiment Instructions (Seller Version)

Sellers Instructions

Before the experiment, subjects were randomly split into two groups: buyers and sellers. These are a translated version of the instructions given to the sellers. Instructions for the buyers as well as the computerized instructions are available upon request.

Today you will take part in a market experiment. Please read through the following instructions carefully. All the information you need to successfully participate in this experiment is written here. If you have questions regarding the experiment or the instructions, please raise your hand. An instructor will come to your desk and will answer your question.

By participating in this experiment, you automatically receive a show-up fee of **10 Francs**. In the course of the experiment you can earn additional money by earning points through trading. The amount of points you will earn depends on your decisions and the decisions of other participants during the experiment.

The experiment is split up into **24 separate periods**. In each period you will interact with other participants in the experiment using the computer in front of you. The points that you earn during this experiments are converted into francs at the end of the experiment. The conversion rate is:

30 Points = 1 Swiss Franc

At the end of the experiment, six periods are randomly chosen and you will receive the amount of money you earned in these periods plus the 10 francs show-up fee in cash.

Please be aware that communication is strictly forbidden during the time you are in the laboratory. Also note that the use of the computer is restricted to the experimental program only. Communication or manipulating of the computer will result in exclusion from the experiment. If you have any questions please raise your hand and an instructor will answer them.

Overview of the course of the experiment

In this study you are a **seller** in a market with RED and BLUE products. The market consists of 5 buyers and 6 sellers. As a seller, you may sell **up to two** products. You will earn a number of points on a transaction equal to the price that you sell a unit minus the cost for producing the unit and any certification costs that you incur.

Your Earnings = Price – Production Cost – Certification Fee

In the market, you may sell two types of products: RED and BLUE. These products are of different quality and may have different valuations to the buyers in the market. **A buyer earns money if he pays less than his valuation for a product.** A buyer's valuation for a product depends on the quality of the product that he receives and the total number of units that he has already bought in the period.

Initially, the buyers and other sellers can not observe the quality of the unit that you are selling. You may choose to offer certified units instead of normal units which guarantee a specific color to the buyer. If you sell a certified unit, you will be charged 60 points in certification fees at the time of transaction.

In total the experiment consists of 24 Periods. The course of each period is as follows:

1. The Trading Phase: In the trading phase, you will trade with buyers in the market. The trading phase in the first 3 periods will be 4 minutes. The trading phase for the remaining periods will be 2 minutes. During the trading phase, you may complete trades either by posting offers that a buyer accepts or by accepting bids from the buyers.

Your offer to sell:

- Your offer to sell consists of the following specifications:
 - 1) the price that buyers have to pay for a unit of the product
 - 2) the quality of the product
 - 3) whether there is a certificate for the product
- The other participants can only see the actual quality of a product if the product is certified. If the product is not certified, the product quality will be labeled "UNKNOWN".

The offers from buyers:

- A buyer's bid to buy consists of the following specifications:
 - 1) the price he is willing to pay for a unit of the product
 - 2) the desired quality of the product
 - 3) whether the buyer requires a certificate or not
 - If a buyer requests a certificate you **must** sell the buyer his desired quality. If the buyer doesn't request a certificate you can sell either quality.
- 2. The Bonus Phase:** The next phase is the bonus phase. In this phase you have to guess how many of the sellers had lower cost producing the RED quality than producing the BLUE quality during the respective period. If your guess is correct you will earn 20 points.
- 3. The Earnings Screen:** At the end of each period you will see the earnings screen. Each participant is informed how much he has earned during the last trading period.

6 out of the 24 Periods are randomly chosen and the earnings of these periods and the show-up fee will be paid out in cash at the end of the experiment.

Detailed course of the experiment

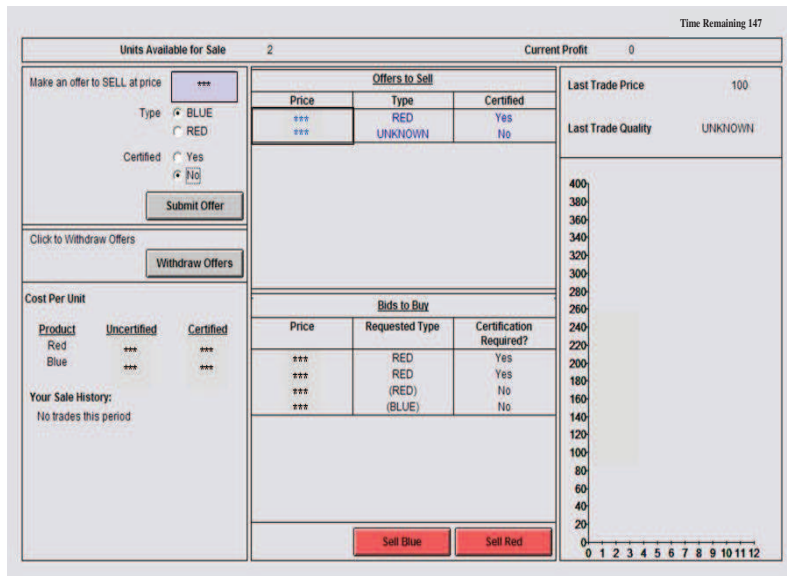
During the experiment you will enter your decisions using the computer. In the following instructions, all the functions will be explained in detail.

1. The Trading Phase

At the beginning of the trading phase, you will be informed of the production costs for the following period. When all players have reviewed their cost and value information, the trading phase will begin.

During the first three periods the trading phase will last for **4 minutes**. In the remaining periods, the trading phase will last **2 minutes**. The clock in the upper right hand corner of the screen will show the remaining time in a period in seconds. When this clock reaches zero the game will immediately end and you will not be able to make any more trades.

During each trading phase you will see the following screen:



Product Quality

There are two possible product qualities: RED and BLUE. Your production costs as well as the valuations of the buyers differ with the quality. In each period either the RED or the BLUE quality can be cheaper for you to produce.

Sellers Production Costs

The production costs of a product depend on two things. First the quality (RED or BLUE) of the product influence the costs and second certification increases the production costs. In every period you will see your costs on the lower left side of the trading screen.

Your costs can change from period to period, so please pay close attention to your production costs.

The following cost structures can occur during the experiment. In each period one of the three following cost structures will be applicable. Please note that different sellers may have different costs during each period.

Case 1, RED Quality is cheaper to produce:

Quality	Costs without certification	Costs with certification
RED	30	30+60 = 90
BLUE	50	50 + 60 = 110

Case 2a, BLUE Quality is cheaper to produce:

Quality	Costs without certification	Costs with certification
RED	80	80 + 60 = 140
BLUE	50	50 + 60 = 110

Case 2b, BLUE Quality is cheaper to produce:

Quality	Costs without certification	Costs with certification
RED	130	130 + 60 = 190
BLUE	50	50 + 60 = 110

Certification

The other participants, buyers and sellers, can only see the quality of a product if the product is certified. A buyer can see the quality of products without a certificate only after the purchase. In this case the quality of the product will be labeled "UNKNOWN".

To reveal the quality of a product to the buyers, you can elect to certify your product. **As you can see in the table above, certification increases the production cost by 60 Points.** The certification costs only occur when a product is sold. So you don't have to pay certification costs for an unsold unit.

Your offers to buyers

You and all the other sellers can post offers to buyers during the whole period. If you want to post an offer you have to specify the following:

- You have to specify a price, which the buyer has to pay for the product. The price has to lie between 0 and 400:

$$0 \leq \text{Price} \leq 400$$

- You have to specify the quality:

$$\text{Quality} = \text{RED or BLUE}$$

- You have to decide whether you will issue a certificate:

$$\text{Certificate} = \text{Yes or No}$$

$$\text{Costs of certification} = 60$$

As soon as you have made all the required specifications you can validate your offer by clicking on the "post offer"-button.

This information will appear on the screen in the field offers to sell and all the other participants, buyers and sellers can see it. Your own offers will appear in blue, the offers of all the other sellers appear in black. **The offers to sell appear in descending order of the price on the screen.**

As soon as a buyer accepts an offer, the respective offer disappears from the screen. If you want to post the same offer again, you have to reenter all the specifications.

As long as you can sell at least one unit you can have two standing offers, one that is certified and one that is not certified. After your second sale all of your standing offers will be deleted.

If you have a standing offer, and you enter a new offer, the new offer replaces the old one, if both offers have the same certification status.

Example:

You have the following standing offers:

Quality	Price	Certified
RED	400	Yes
BLUE	50	No

Now you enter an offer for a RED quality product at the price of 350 and you offer a certificate. Your standing offers will change to:

Quality	Price	Certified
RED	350	Yes
BLUE	50	No

Now you enter an offer for a RED quality product at the price of 250 and you do not offer a certificate. Your standing offers will change to:

Quality	Price	Certified
RED	350	Yes
RED	250	No

To withdraw offer you can click the “withdraw offers”-button and **all your offers are withdrawn**.

Accepting offers from buyers

The offers to buy are sorted in descending order of the price.

To accept an offer from a buyer, you select the line of the respective offer and click the “sell RED”-button, if you want to sell the RED Quality or click the “sell BLUE”-button if you want to sell the blue quality.

- If the buyer doesn't request certification, you can sell either quality.
- If the buyer request certification, you have to sell the desired quality AND you have to pay the certification cost.

History

On the bottom left side of the screen, you will see your personal history. There you will see detailed information about the products you have sold so far during the respective period. For every product purchased you will see:

- the quality
- whether the product was certified
- the price you got
- the resulting earnings

On the right side of the screen you will see the market history. On the top you will find the information of the last traded good. Below you find a chart with all the trades of the period.

On the axis to the right you will find the amount of products traded. On the other axis you will find the price that has been paid for the product. Depending on the quality and certification of the product, the entry is of a different color:

- RED certified products appear in **red**
- BLUE certified products appear in **blue**
- Uncertified products appear in **black**

2. The bonus phase

Following the trading phase is the bonus phase. In this phase you have to guess how many of the sellers had lower cost producing the RED Quality than producing the BLUE quality during the respective period. If your guess is correct you will get 20 points.

3. The earning screen

At the end of each period you will see the earnings screen. There you will find your market earnings of the period.

Six out of the 24 Periods are randomly chosen and the earnings of these periods and the showupfee will be paid out in cash at the end of the experiment.

Omitted: Examples of How Earnings Is Calculated, Example of Randomized Payment

Exercises

The experiment starts only after all participants are fully accustomed with the experiment. To ensure this, we ask you to solve the exercises on this page.

Please also write down intermediary steps.

After these exercises you will have the possibility to get to know the trading screen before the first period starts. The options you have will be presented again in detail and you can do some trial trades.

For these exercises please use the following cost structure:

	Cost without certification	Cost with certification
ROT	80	140
BLAU	60	120

Exercise 1: A buyer bids 180 for a product and doesn't request a certificate, how much do you earn with this sale?

Earnings if you sell a BLUE quality product =
Earnings if you sell a RED quality product =

Exercise 2: You sell a RED Quality good for which a buyer paid 150. How high are your earnings if the buyer requests a certificate and what do you earn if he doesn't request a certificate?

Earnings with certificate =
Earnings without certificate =

Exercise 3: There are the following two standing offers of buyers:

Offer number	Price	Quality	Certificate requested
1	220	BLUE	Yes
2	180	RED	No

Through which sale can you make the higher earnings?

Possible earnings through offer number 1 =
Possible earnings through offer number 2 =