

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

TOM WILKENING[†]

July 13, 2016

We study an experimental market in which some sellers are prone to moral hazard, and in which a private-order contract enforcement institution exists that can mediate trade and prevent sellers from renegeing on their contractual obligations. Using the institution to resolve the moral-hazard problem is costly. We demonstrate that in this market, the utilization of the private-order contract enforcement institution may make public and private market signals uninformative and inhibit learning. We study whether this potential information externality can limit adaptation away from the private-order institution when it is efficient to do so. Consistent with theory, we find inefficient persistence when the institution is used, but by contrast, efficient adaptation in other situations. Providing information to individuals who are using the private-order institution allows them to partially adapt.

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 also indebted to Georg Hartmann and Holger Herz for their help with translation and managing the laboratory in Zurich. I also thank Daron Acemoglu, Jim Berry, David Cooper, Nisvan Erkal, Greg Fischer, Simon Loertscher, Nikos Nikiforakis, Chris Smith, Rob Townsend, Patrick Warren, and four anonymous referees for valuable comments on the paper. 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 the agreed-upon service. This temporal delay – pervasive in financial transactions, labor relationships, agency relationships, and the exchange of experience goods – gives rise to what Greif (2000) terms the “fundamental problem of exchange,” whereby the second mover in a relationship has incentives to renege on their contractual obligations (Greif, 2005).

A myriad of contract-enforcement institutions (CEIs) 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, accreditation associations, banks, certifiers, and information intermediaries, provide supplemental enforcement capabilities, but at additional costs.

Despite their important role in facilitating trade, the forces that lead to the formation and persistence of private-order CEIs are not well understood. Evidence from the economic history literature (e.g., Wittfogel, 1957; Li Yang, 2002; Goetzmann and Köll, 2005) and the development literature (e.g., Fafchamps, 2004) suggests that private-order CEIs do not always 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).

The institutions-as-equilibrium literature seeks to understand how institutions might evolve by studying how the interactions of agents might lead to the adoption of new institutions. It further studies how previously adopted institutions might lead individuals to act in a manner that perpetuates the institution in the future.² Of particular interest to this literature is understanding how institutions might be self-enforcing through three interconnected channels: (1) the confirmation of beliefs about the types of others through observed outcomes; (2) the inter-temporal transmission of beliefs to newcomers; and (3) the reinforcement of actions through coordination (Greif, 2006).

This paper contributes to the institutions-as-equilibrium literature by exploring an information externality embedded in the services provided by the private-order CEIs that may limit the ability

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 in which identity is anonymous, and thus consider CEIs that would typically be classified as being designed.

²The broader institutions literature also studies how socio-political conflict and agency shape institutions. See, for example, the review articles by Acemoglu, Johnson, and Robinson (2005) and Ogilvie and Carus (2014).

for market participants to observe changes in their environment.³ 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 a 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 inhibit learning and can influence the responses of market participants to changes in their environment.

We consider an environment where sellers can produce high-quality and/or low-quality experience goods that the buyer cannot differentiate between at the point of sale. Sellers are heterogeneous in their production cost and, without use of a private-order CEI, face different incentives to renege on their contractual obligations. These incentives vary with 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 that can guarantee quality, but the use of this private-order CEI is costly.

Depending on the distribution of sellers' production costs, up to three different types of rational expectations equilibria may exist. We refer to these equilibria as unmediated, partially-mediated, and mediated, to reflect the use of the private-order CEI in mediating trades. We demonstrate that — relative to the other two types of equilibria — 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. In particular, while buyers can learn about the distribution of sellers' production costs from both their private experiences and by observing prices in the unmediated equilibrium, both channels are eliminated in the mediated equilibrium. Thus, the adoption of the private-order CEI may prevent traders from learning about their environment.

Given the theoretical differences in the information that exists when the private-order CEI is used and when 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 a setting where the distribution of sellers' production costs changes over time. Consistent with the model's predictions, subjects who establish an unmediated or a partially-mediated equilibrium in a "Safe" low-risk environment adapt to the mediated equilibrium when risk is increased. By contrast, subjects who establish the mediated equilibrium in a "Hazardous" high-risk environment do not adapt to a more efficient unmediated or partially-mediated equilibrium when risk is decreased.

Based on the underlying theoretical model, the private-order CEI not only creates the information externalities described above, but it is also predicted to generate a coordination problem for buyers and sellers who must simultaneously agree to trade without the private-order CEI at prices above the existing equilibrium price. To help disentangle the coordination channel from

³Warren and Wilkening (2012) study information externalities in a regulatory context in which adopting regulation limits what can be learned about the state of nature. Even with benevolent social planners, the information externality can lead to persistence in information-suppressing regulation. The current paper shows that these information externalities can arise endogenously from the utilization of private-order CEIs and provides experimental evidence regarding their influence.

the information channel, we conduct an additional set of experiments where we provide exogenous information about the distribution of sellers production costs. Consistent with the information channel, a small subset of buyers eventually trade without the private-order CEI when risk is decreased. Our results thus suggest that at least part of the persistence of the private-order CEI is due to the information externality that is inherent in its use.

Taken together, our results provide evidence that adopted institutions can have an 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 select efficient market institutions in the short run.

The model and experiments developed below may help us understand why utilization patterns of private-order CEIs do not respond strongly to changes in the underlying environment. The response of foreign firms that cross-listed their shares on the U.S. stock market to the 2002 Sarbanes-Oxley Act provides a useful example.

Equity markets across the world are supported by a combination of private-order exchanges and public-order regulators. Foreign firms that cross-list equity shares on U.S. stock exchanges must register with the U.S. Securities and Exchange Commission (SEC) and are generally subject to U.S. federal security laws.⁴ By cross-listing equity shares rather than using alternative securities that do not have registration requirements, foreign firms can voluntarily opt into a regulatory environment with strict disclosure requirements and a strong regulator.

The finance literature argues that cross-listing can be used as a signal of financial health (e.g., Baker, Nofsinger, and Weaver, 2002) and can reduce the potential for moral hazard by increasing transparency and introducing sanctions for fraudulent acts (e.g., Coffee, 1999, 2002; Stulz, 1999). However, meeting the reporting standards of the SEC has significant costs. Our model would predict that the extent to which firms choose to cross list depends on the level of moral hazard that exists in the home country, and the premium that exists for cross-listing. Consistent with the model, cross-listing is most common from companies where investor protection in their home country is weak and where the arbitrage opportunity of switching regulatory jurisdictions is large (Dojidge, Karolyi, and Stulz, 2009).

When the Sarbanes-Oxley Act was passed in 2002, very few exemptions were made for foreign cross-listing firms who faced new corporate governance requirements and significant additional compliance costs. As delisting was difficult, cross-listing firms objected strongly to the new legislation, and argued that they were trapped in a regulatory system from which they could not escape. In March 2007, the SEC adopted revised rules that relaxed barriers to deregistration and delisting of foreign companies.

At the time of the revision of delisting rules, there was significant concern that Sarbanes-Oxley reduced the competitiveness of the U.S. stock market and that there may be a flood of delistings from U.S. exchanges (Zingales, 2007; Small and Zhu, 2007). However, because cross-listing on the

⁴This is not the case for exchanges in most other countries. See Laby and Broussard (2009) and Bianconi, Chen, and Yoshino (2013) for a broader discussion.

U.S. exchange is used both to mitigate moral hazard in some firms and to signal quality in others, our model would predict that very few firms are likely to want to unilaterally exit from cross-listing if given the chance. It would further predict that if exit is observed, it should only occur for companies where investor protection in the home country is strong and where the proportion of firms cross-listing from a particular country is small. Consistent with these predictions, empirical evidence by Doidge, Karolyi, and Stulz (2009) suggests that the exit response of Sarbanes-Oxley was muted: only a very small number of firms chose to delist and the delisting firms were primarily from Europe, Canada, and Australia, where investor protections were strong and the initial proportion of cross-listing firms was small.

The paper also highlights potential issues related to the prudential regulation of banks. In fulfilling its role as a liquidity provider, banks issue liquid liabilities to small scale depositors but invest primarily in illiquid assets. This service improves the welfare of the depositors but, as shown in Diamond and Dybvig (1983), creates the potential for bank runs.⁵ Deposit insurance has the benefit of preventing bank runs, but these guarantees increase moral hazard since the owner of deposits lose their incentives to monitor banks (Allen, Carletti, Goldstein, and Leonello, 2015). Viewing deposit insurance as a mechanism to maintain a more liquid partially-mediated equilibrium, our paper suggests that deposit insurance is likely to only be effective in environments where other forms of regulation can properly mitigate moral hazard.⁶ It also highlights an information externality that may make the underlying level of moral hazard in the environment unobservable thereby making the outcome of regulatory changes uncertain.

The paper is organized as follows. After reviewing the literature in Section 2, Section 3 provides theoretical motivation for the experiment. Section 4 develops the experimental design. Section 5 reports the main experimental results and Section 6 concludes.

2 Related Literature

This paper relates to 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) that seeks to understand how institutions might be self-enforcing by (1) confirming beliefs about the types or actions of others through observed outcomes, (2) inter-temporally transmitting equilibrium beliefs to newcomers, and/or (3) reinforcing actions through coordination. We show that private-order CEIs may be self-enforcing in all three ways: the information externality embedded in private-order CEIs 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

⁵In previous experimental work, Madiès (2006) studies the formation of bank runs in a repeated coordination game and finds that once a bank run has occurred, the equilibrium with bank runs is persistent. The author finds that suspending trade can curb bank runs, but find little evidence that partial deposit insurance is effective. Schotter and Yorulmazer (2009) also study bank runs and finds that providing information about the solvency of banks and offering deposit insurance can mitigate bank runs.

⁶Consistent with this result, empirical work by Demirguc-Kunt and Kane (2002) finds strong cross-country variation in the impact of deposit guarantees.

nature. Further, the adoption of private-order CEIs makes trading without the institution more risky and leads to a coordination problem.

An alternative explanation of institutional persistence, studied in other institutional contexts by North (1981), Brainard and Verdier (1994), Coate and Morris (1999), and Acemoglu and Robinson (2006, 2008), 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), which suggests that private-order CEIs form from an initial informational advantage that enables them to gain from exchanging this information with others. Our paper makes no claim as to the relative importance of the two channels but instead uses experimental methods to control for market power.

The institutions-as-equilibrium literature builds on the coordination literature (e.g., Schelling, 1960; Van Huyck, Battalio, and Beil, 1990; Ochs, 1990) and 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 Cooper, Garvin, and Kagel (1997a,b) and Cooper and Kagel (2003, 2005, 2008) who study equilibrium selection, learning, and history dependence in the limit pricing game of Milgrom and Roberts (1982). The 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. Our paper contributes to these literatures by studying how information may play a role in the persistence of institutions.⁷

While this paper is the first to analyse information externalities created by 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 in place today limit the observation of potentially detrimental actions in the future. Warren and Wilkening (2012) study information externalities in a regulatory context. Fernandez and Rodrik (1991) and Friedrich (2013) combine information externalities with voting and rent seeking to study policy persistence.

Finally, the paper is related to the literature on asymmetric information in markets with multiple equilibria. 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.⁸

⁷The paper is also related to Deck and Nikiforakis (2012), who find that allowing for cheap talk via an unpaid continuous-time minimal effort game dramatically increases coordination to efficient equilibria when monitoring is perfect. In our paper, we allow individuals to trade in a continuous-time double auction environment where efficient equilibria often emerge. We show that even in an environment with a rich signalling space, inefficient equilibria may be persistent if information is garbled.

⁸See also Nöldeke and Samuelson (1997) for a dynamic model in which both pooling and separating equilibria might be stable.

3 Theoretical Motivations

In this section we provide the theoretical motivations for the experiment. We show that in a market with both adverse selection and moral hazard, multiple stable rational expectations equilibria may exist that vary in the use of a private-order CEI that can certify quality. We then study the informational properties of these equilibria to understand how public and private signals might be useful in the updating of beliefs about the underlying distribution of seller costs. We conclude by discussing how the lack of updating in the mediated equilibrium may lead to its overall persistence.

To emphasize intuition, we use a theory model with single unit demand and supply, and discuss the equilibria in relation to the parameters used in the actual experiment. The experiment allows for multiple units to be traded, but buyer demand functions and seller cost functions are constructed so that the rational expectations equilibria of the experiment coincide with the equilibria described here. A formal construction of the rational expectations equilibria is included in Appendix A.

3.1 Primitives

Consider an economy with experience goods of high (H) and low (L) quality which are referred to as “units”. There are N buyers indexed by $i \in \{1, \dots, N\}$. There are M sellers indexed by $j \in \{1, \dots, M\}$ divided into three types $s \in \{G, C, B\}$ (Good, Conditional, and Bad). The number of sellers who are of type s is M_s . The true proportions of type- G sellers and type- C sellers are \mathbf{g} and \mathbf{c} , respectively.

To restrict attention to the environment considered in the experiment, we make the following two assumptions about the distributions of buyers and sellers:

Assumption 1 *There is excess demand: $N > M$.*

Assumption 2 *There is at least one type- B seller: $M_b \geq 1$. M_b is common knowledge.*

As discussed below, Assumption 1 restricts attention to the part of the parameter space where prices are most informative. Assumption 2 ensures that buyer beliefs are always well defined and allows us to discuss beliefs about the seller’s type using a single distribution.

Each buyer can consume a single high- or low-quality unit. Likewise, each seller can produce a single high- or low-quality unit. We initially consider the case where buyers are homogeneous, have gross utilities for consuming the high and low quality good of $U^H > U^L$ 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 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 full support over $g \in \{0, \frac{1}{M}, \frac{2}{M}, \dots, \frac{M-M_b}{M}\}$ and expected value $\mathbb{E}(\hat{g})$. Assumption 2 implies that the distribution of seller beliefs can be fully expressed by $p(\hat{g})$.

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 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. The analysis is thus restricted to cases where all certified units are of high quality.

Sellers 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-neutral. Thus, the expected cost of producing a low-quality unit is $C^L := \alpha F$.

If a seller of type s produces and exchanges a high-quality unit, she pays a marginal cost C_s^H . We make two assumptions with regard to the marginal costs of the three types of sellers:

Assumption 3

$$C_B^H > C_C^H > \alpha F > C_G^H.$$

Assumption 4

$$C_B^H > \alpha F + U^H - U^L - T > C_C^H.$$

Assumption 3 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. Assumption 4 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.

In the experiment, we set $U^H = 200$, $U^L = 100$, $T = 60$, $C_B^H = 130$, $\alpha F = 50$, $C_C^H = 80$, and $C_G^H = 30$. Based on these parameters, $\alpha F < U^L$, which implies that trade is always welfare improving ex-ante. Further, $C_B^H - \alpha F < U^H - U^L$, which implies that the social optimum occurs when all three seller types produce high-quality units.

3.2 The Rational Expectations Equilibria

The experiment in the next section uses a continuous-time double auction to conduct trade. This mechanism is chosen as it provides subjects with a large strategy space that may allow individuals to signal to each other through bids and coordinate to a more efficient equilibrium. However, it is not a mechanism that is easily analyzed with standard game-theoretic tools.

In double auction experiments without uncertainty it has been found that the competitive equilibrium is a good predictor of equilibrium behavior. As the rational expectations equilibrium is a natural extension of the competitive equilibrium with uncertainty, we derive the rational expectations equilibria here and use them as a basis for the rest of the analysis.⁹

⁹Recent theoretical work also suggests that the rational expectations equilibrium is the limiting case for many

A rational expectations equilibrium 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 define the set of rational expectations equilibria, it is easiest to consider the trade of certified and uncertified units as two independent markets each with its own price. Let P^C be the price for high-quality certified units, P^{NC} be the price of uncertified units of unknown quality, and ΔP be the difference in these two prices.

We show in Appendix A that 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:** $P^C = U^H$, $P^{NC} = U^L$. 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:**¹⁰ $P^{NC} = U^H - (1 - \mathbb{E}(\hat{g}))(U^H - U^L)$, $P^C = U^H$. 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.

As can be seen by comparing the two equilibria, certification plays a different role for type- G and type- C sellers. Regardless of their decision to certify, type- G sellers produce high-quality units and have no incentive to renege on their contractual obligations even without the private-order CEI. For these sellers, certification acts as a signal of quality and can be used to resolve the adverse selection problem that exists in the uncertified market. By contrast, type- C sellers change the quality of their production when moving between the certified and uncertified markets. Certification thus resolves moral hazard for type- C sellers and prevents them from reneging on their contractual obligations.

As type- G sellers use certification to signal their quality, the unmediated equilibrium can only exist in cases where these sellers have no interest in separating from the other seller types. The following proposition provides conditions under which both potential equilibria exist:

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 $\mathbb{E}(\hat{g}) \geq 0.4$. Note that type- B sellers produce low-quality uncertified units in both equilibria, and thus both equilibria are inefficient relative to the first best.¹¹

game-theoretic models of exchange. For instance, Satterthwaite, Williams, and Zachariadis (2015) shows that under the buyers bid double auction, prices converge to the rational expectations equilibrium even for very small markets. Reny and Perry (2006) find a similar result in the case of a continuum of agents.

¹⁰Note 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.

¹¹The existence of both the mediated and unmediated equilibrium is due to Assumptions 3 and 4, which ensure that

3.3 Market Information

Having defined the mediated and unmediated equilibrium, we now return to the central question of information and the utilization of the private-order CEI. We begin with 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 3.3.1, we allow for buyers to have heterogeneous beliefs about the distribution of seller types 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.¹²

By contrast, in the unmediated equilibrium, the price of uncertified goods, $P^{NC} = U^L + (U^H - U^L)\mathbb{E}(\hat{g})$, and thus $\mathbb{E}(\hat{g}) = \frac{P^{NC} - U^L}{U^H - U^L}$. Hence, 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 the mediated equilibrium, no market signal generates information that can be used to update beliefs about the proportions of type- G and type- C sellers. In an unmediated equilibrium with a common prior, price is a sufficient statistic for $\mathbb{E}(\hat{g})$.*

The information properties of the mediated equilibrium is based on Assumptions 3 and 4, which imply that when certification occurs, type- G and type- C sellers take the same action in equilibrium and are thus indistinguishable. The result requires that there is both an adverse selection problem in addition to the moral hazard problem and that the utilization of the private-order CEI results in pooling. The result is robust to a modified environment where the certification is only partially effective in revealing quality as long as type- G and type- C sellers continue to take the same action in equilibrium.

The information properties of the unmediated equilibrium uses the fact that price is pinned down by the buyer's valuations and thus also uses Assumption 1. When Assumption 1 is relaxed, price is no longer informative in the unmediated equilibrium, but buyers may still learn about the

(i) type- G sellers always produce high-quality units and (ii) the utilization of certification in the mediated equilibrium resolves moral hazard for type- C sellers. Assumption 1 is not necessary for multiple equilibria to exist and can be replaced with the less strict assumption that $N > M_g$. However, with excess supply prices are less informative in the unmediated equilibrium.

¹²A new buyer could, however, ascertain the proportion of type- B sellers in the environment by observing the number of trades in the uncertified market. We focus on the distribution of type- G and type- C sellers as this is the most relevant information in evaluating what would happen in the counterfactual case of the unmediated equilibrium.

quality of goods based on their private purchase decisions and through ex-post information. We discuss learning through private experience in the next section.

3.3.1 Heterogeneous Beliefs and Learning

The discussion above highlights the relationship between the use of the private-order CEI 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 useful 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 buyer whose type pins down prices 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 mediated equilibrium no agent can distinguish between type- G and type- C sellers. 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 with support $g \in \{0, \frac{1}{M}, \frac{2}{M}, \dots, \frac{M-M_b}{M}\}$, and where the discrete distribution is single peaked. 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 most loss averse buyer who is willing to trade in the uncertified market, we require that each buyer must be willing to trade given the revelation that they are pivotal. In the unmediated equilibrium, this requires that for each buyer assigned an uncertified unit:

$$P^{\mathcal{N}C} \leq U^L + \mathbb{E}(\hat{g}|P^{\mathcal{N}C})(U^H - U^L). \quad (1)$$

Let P^* be the largest $P^{\mathcal{N}C}$ satisfying (1) for at least M buyers. Then, if $P^* \geq U^H - T$ an unmediated equilibrium exists where M buyers trade at the price P^* .¹³

The assumption of a single-peaked prior is not required for the convergence of beliefs, but ensures that the willingness of an individual to buy in the uncertified market is decreasing in the price of uncertified trades. Given two equilibrium where there are M buyers willing to trade at prices $P^{\mathcal{N}C}$ and $\tilde{P}^{\mathcal{N}C}$, a single-peaked prior ensures that $\mathbb{E}(\hat{g}|P^{\mathcal{N}C}) \geq \mathbb{E}(\hat{g}|\tilde{P}^{\mathcal{N}C})$ if $P^{\mathcal{N}C} > \tilde{P}^{\mathcal{N}C}$. This condition is enough to ensure that there exists a price that clears the uncertified market any time the unmediated equilibrium exists.¹⁴

¹³As 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.

¹⁴We can think of a single peaked prior 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.

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*

$$p_t^i(\hat{g}) \xrightarrow{a.s.} \mathbf{g}. \quad (2)$$

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, $\mathbb{E}(\hat{g}|P^{\mathcal{N}\mathcal{C}}) \rightarrow \mathbf{g}$ 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 mediated 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 ,*

$$p_0^i(\hat{g}) = \dots = p_{t+1}^i(\hat{g}) = \dots = p_\infty^i(\hat{g}). \quad (3)$$

As can be seen from Proposition 4, the mediated equilibrium eliminates all information that might be used to update beliefs when Assumptions 3 and 4 are satisfied. Thus, if a market reaches a mediated equilibrium and there is an exogenous shift in the proportion of type- G and type- C sellers, we would expect buyers' beliefs to remain unchanged.

3.3.2 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. These equilibria 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.¹⁵ Let $\mathcal{B} = \{\lambda_1, \lambda_2, \dots, \lambda_N\}$ be the set of buyer types, 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.

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 model without loss aversion, 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 equilibria may be stable since the willingness to pay for uncertified units is decreasing in loss aversion, leading to a downward sloping aggregate demand function. The partially-mediated equilibria may exist any time the unmediated equilibrium exists.

In the partially-mediated equilibrium $P^{\mathcal{N}C} = U^H - T$ and $P^C = U^H$. Thus 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. In particular, a buyer who observes S^C units traded in the certified market knows that there are at least S^C type- G sellers in the economy. Thus, starting from any prior $p(\hat{g})$, the posterior distribution $q(\hat{g}|S^C) = 0$ for all $\hat{g} \in \{0, \frac{1}{M}, \dots, \frac{S^C-1}{M}\}$.

Proposition 5 *In a partially-mediated equilibrium with a common prior, price conveys no information to market participants regarding the proportion of type- G sellers. However, market participants*

¹⁵The 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.

can determine the minimum number of type- G sellers in the economy by observing the number of trades in the certified market.

As with the unmediated equilibrium, buyers who are trading in the uncertified market each period can receive high- or low-quality units. Thus they can update their beliefs about the number of type- G sellers trading in this environment. Starting from a set of heterogenous priors, buyers who trade in the uncertified market each period can learn about the risk in this market and update their beliefs accordingly.

Proposition 6 *Consider a sequence of periods in which the partially-mediated equilibrium occurs each period with S^C certified trades. Then there exists at least $M - S^C$ buyers such that*

$$p_t^i(\hat{g}) \xrightarrow{a.s.} \mathbf{g}. \quad (4)$$

While Proposition 6 indicates that individuals in the partially-mediated equilibrium may learn the true distribution of types, this learning process may be slow and convergence speeds depend on the distribution of seller types. In the experiment described in the next section, we introduced a public information treatment in which we revealed the composition of trades in the uncertified market after each period. The following property makes clear that this public information is uninformative in the mediated equilibrium, but may be informative in both the unmediated and partially-mediated equilibrium if individuals have incorrect priors about the distribution of types.


Proposition 7 *In a partially-mediated equilibrium or an unmediated equilibrium, providing public information regarding the proportion of high-quality units traded in the uncertified market gives market participants perfect information regarding the proportion of type- G sellers. Public information is uninformative in the mediated equilibrium.*

4 The Experiment

Each experiment consisted of a fixed matching group of 5 buyers and 6 sellers who interacted for 24 periods. As summarized in Figure 1, subjects began each experiment by being assigned values and costs. They then played 12 periods each consisting of three stages: a trading stage, a bonus stage, and an information stage. Subjects were reminded of their cost and value schedules prior to each new trading stage. In period 13, four sellers were assigned new cost schedules and all other participants were reminded of their original cost and value schedules. Participants played an additional 12 periods using the new cost and value schedules.

The Cost and Valuation Assignment Stage Subjects in the experiment could trade experience goods of High (H) and Low (L) quality that the buyer could not differentiate between at the point of sale without the use of a private-order CEI. The cost of using the private-order CEI, called

Figure 1: Overview of the Experiment



a “certification cost,” was 60 points. This cost was known to both buyers and sellers, and paid for by the seller each time a certified unit was traded.

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 3, we designate as G , C , and B (Good, Conditional, and Bad).

Table 1: Seller Per-Unit Production Costs

	Low Quality	High Quality
Good	50	30
Conditional	50	80
Bad	50	130

Each of the five buyers could consume up to three units, creating an aggregate demand of 15 units. 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. Buyer types were fixed throughout the experiment.

Table 2: Buyer Valuations

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

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.

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 were predicted to range from 100 to 200 points, and the profits that buyers received from purchasing their inframarginal units. Since no buyer or seller had market power and there are no strategic incentives that arise from being able to trade multiple goods, there is no difference in the equilibria that exist in the experimental environment and the model environment.

The Trading Stage Trade was conducted through two computerized markets — a certified market and an uncertified market — 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.¹⁶ 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.¹⁷ If a seller transacted in the certified market, 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. Bids and offers could be changed or withdrawn at any time with no restriction on pricing.

In each period, a history of trades from the current period was available in graph form for all

¹⁶A 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. Departing slightly from the design developed by Smith (1964), subjects in this experiment were also free to enter the bid and ask queues at any price and accept any offer from the other side of the market. These changes gave sellers flexibility in their pricing strategies and allowed buyers a way to avoid offers that they believed to be of low quality.

¹⁷Buyers were free to request certified low-quality units. In practice, this rarely occurred.

subjects in the market. Certified trades were shown in the color of the actual unit traded while uncertified trades of all qualities 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.

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.¹⁸

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 2160 unique buyer-period observations.

The Bonus Game After each trading period, both buyers and sellers participated in a bonus game. In the bonus game, subjects were reminded that each of the six sellers might have a lower cost for producing a high-quality unit than a low-quality unit, and they were asked to guess 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 and was used as part of our information treatments described below.

The Information Screen Following the bonus game, subjects were given a summary sheet which varied by the information environment. In the *Private Information Environment*, individuals were only informed about the total number of units traded with and without certification. This information was a replication of the trade information that was observable in graph form in the Trading Stage. In the *Public Information Environment*, individuals were also informed about the actual number of high- and low-quality units traded in the uncertified market. Finally, in the *Full Information Environment*, subjects were reminded of their guess in the bonus game, told the correct number of type- G sellers in the environment, and told whether they had successfully guessed the correct number of type- G sellers. This information revealed the true number of type- G sellers to all market participants.

4.1 Experiments and Treatments

Subjects were assigned to one of five potential experiments each consisting of two 12-period blocks. Subjects in each experiment were initially assigned to one of two *moral hazard environments* — Safe (\mathcal{S}) and Hazardous (\mathcal{H}) — which varied in the number of sellers who were assigned to the three seller types. The distribution of seller types for both environments are shown in Table 3.

¹⁸One 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.

Table 3: Moral Hazard Environments

	Good	Conditional	Bad	g
Safe (\mathcal{S})	5	0	1	.833
Hazardous (\mathcal{H})	1	4	1	.167

In experiments that began in the 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 experiments beginning in the Hazardous environment. To distinguish between periods before and after the switch, *Pre* and *Post* superscripts are appended to the treatment identifiers.

As can be seen in Table 4, experiments differed both in the ordering of the Safe and the Hazardous environments, and in the amount of information that was revealed in the information screen. As there are no predicted theoretical differences between the Private and Public information environments, we pool data from these environments together for the purposes of analysis and compare only the full information environment to the other two information environments. Our final design has six treatments that are identified by the moral hazard environment, the set of periods considered, and whether the information environment is full or private/public.

Table 4: Experiments and Treatments

Experiment	Periods 1-12	Periods 13-24	Information	Treatments	Number of Groups
1	Safe	Hazardous	Private	$\mathcal{S}^{Pre}, \mathcal{H}^{Post}$	3
2	Safe	Hazardous	Public	$\mathcal{S}^{Pre}, \mathcal{H}^{Post}$	3
3	Hazardous	Safe	Private	$\mathcal{H}^{Pre}, \mathcal{S}^{Post}$	3
4	Hazardous	Safe	Public	$\mathcal{H}^{Pre}, \mathcal{S}^{Post}$	3
5	Hazardous	Safe	Full	$\mathcal{H}_{FI}^{Pre}, \mathcal{S}_{FI}^{Post}$	6

4.2 Protocol

Experiments 1-4 were run at The University of Zurich in 2007. Subjects were drawn from a centralized database comprised of undergraduate students from The University of Zurich and UTH-Zurich. Six sessions were conducted, where each session consisted of two independent groups of 11 subjects who remained in fixed groups and roles over all 24 periods. We ran two groups per session to increase anonymity of participants. Trades were conducted in points and converted to Swiss Francs at the end of the experiment, and subjects were paid for six randomly selected periods at a conversion rate of 30 points to 1 Swiss franc at the end of the session.¹⁹ We paid for only a subset of periods to discourage potential repeated game effects and to mitigate wealth effects.²⁰ We paid for more than one period to minimize the chance of a subject having negative earnings.

¹⁹Randomization of payments was done at the subject level.

²⁰In practice it would be very hard for subjects to develop repeated game strategies given the anonymous nature of the exchanges and the difficulty in detecting defection.

Experiment 5 was run at The University of Melbourne in 2014. Subjects were undergraduate students at the University of Melbourne and were randomly invited from a pool of more than 3,000 volunteers using ORSEE (Greiner, 2004). Three sessions were conducted, and each session consisted of two independent groups of 11 subjects who remained in fixed groups and roles over all 24 periods. Six additional sessions were run at Melbourne in 2015 where all sellers were switched to type- G in later periods; we discuss the result of these additional experiments in Appendix B.

In all sessions, subjects began by reading an extensive set of written instructions and taking a control quiz that they were required to correctly answer before proceeding. After the control quiz, a common verbal set of instructions were read aloud. These instructions summarized the trading environment and emphasized the rules of the bonus game. Finally all subjects went through a computerized set of instructions where they practiced placing trades, practiced accepting trades, and where all information concerning past trades was pointed out. The instruction period lasted between 40 and 50 minutes. After all 24 periods of the main experiment, risk aversion 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.

4.3 Predictions

The goal of our design is to first begin trade in one environment where the mediated equilibrium reliably forms, and one environment where the unmediated or partially-mediated equilibria reliably forms. We then perturb the underlying distribution of sellers in a way that should be undetectable in the mediated equilibrium, but which makes this equilibrium inefficient. A precondition for the rest of the study is that the mediated equilibrium reliably forms in experiments that start in the Hazardous environment, and that the unmediated or partially-mediated equilibria reliably form in experiments that start in the Safe environment.²¹

Prediction 1 *Trade converges to the mediated equilibrium in periods 7-12 of experiments beginning in the Hazardous environment. Trade converges to the unmediated or partially-mediated equilibrium in periods 7-12 of experiments beginning in the Safe environment.*

If Prediction 1 holds, Propositions 2, 3, and 5 predict that subjects in an unmediated or partially-mediated equilibrium learn from their private trades, prices, and public signals and thus can adapt to changes in the environment. By contrast, Propositions 2 and 4 predict that individuals who are in the mediated equilibrium cannot observe changes in the proportion of type- G and type- C

²¹Note that in the Hazardous environment, there is only one equilibrium while in the Safe environment there are three types of equilibria. For our prediction in the Safe environment to hold, behavior must not only converge to an equilibrium, but the mediated equilibrium must not be selected. We were guided here by results in Cooper, Garvin, and Kagel (1997a,b) that found that individuals naturally gravitated toward the pooling equilibria in the limit pricing game of Milgrom and Roberts (1982) when this equilibrium existed.

sellers in the environment. Thus there should be no way for individuals to adapt if the environment changes from Hazardous to Safe. We make the following prediction:

Prediction 2 *In the private or public information treatments, subjects that begin in the Safe environment will shift to the mediated equilibrium when the environment is changed to Hazardous.*

Prediction 3 *In the private or public information treatments, subjects that begin in the Hazardous environment will remain in the mediated equilibrium when the environment is changed to Safe.*

Due to the non-strategic nature of the rational expectations equilibria used as a solution concept, the mediated equilibrium is always an equilibrium in the Safe environment even with full information. As such, it is possible that persistence of the mediated equilibrium is a result of a coordination problem between buyers and sellers who must simultaneously agree to trade in the uncertified market at prices that are above the existing equilibrium prices.²² The Full Information treatment disentangles these two channels by eliminating the information channel while holding the environment constant. If the information channel is contributing to the persistence of the mediated equilibrium, we predict:

Prediction 4 *In the full information treatments, subjects that begin in the Hazardous environment will adapt to the unmediated or partially-mediated equilibrium when the environment is changed to Safe.*

Predictions 1-4 generate point predictions for each of our treatments as well as a set of within-experiment and between-experiments hypotheses tests. As seen in Panel (a) of Figure 2, we predict that the price of uncertified trades will converge to a value between [140, 183] in the \mathcal{S}^{Pre} and \mathcal{S}_{FI}^{Post} treatments, and that there will be more than two uncertified high-quality units traded. The model predicts a price of 100 in the \mathcal{S}^{Post} treatment and all Hazardous treatments. As seen in Panel (b) of Figure 2, we would predict to reject the null hypothesis generated from a Mann-Whitney-Wilcoxon test that the distribution of average uncertified prices in the \mathcal{S}^{Pre} and \mathcal{S}_{FI}^{Post} treatments are the same as in treatments \mathcal{H}^{Post} , \mathcal{S}^{Post} , \mathcal{H}^{Pre} , and \mathcal{H}_{FI}^{Pre} . We would not predict significant differences in any of the other pairwise tests. An identical rejection pattern is predicted for the number of uncertified high-quality units.

As our predictions are based on behavior that is predicted once an experiment has converged to an equilibrium, we restrict attention to periods 7-12 and 19-24 in all the analysis. The number of omitted periods was decided prior to running the experiment and is based on two initial pilots.

²²The double auction mechanism that we use in the experiment provides a large strategy space that may make it easier for buyers and sellers to coordinate away from inefficient equilibrium relative to more standard coordination games. For example, the mediated equilibrium could be eliminated by (1) a type- G seller (costlessly) posting an uncertified trade at $p \in (140, 150)$ and (2) a loss-neutral buyer accepting a trade at the price p conditional on observing two uncertified offers. Recall that a seller can have both a certified and an uncertified offer out at the same time. Thus, when the market price for certified units is 200, the offer to trade at price p is costless and can only increase the seller's profit if accepted. The buyer who knows the distribution of seller types knows that there is at most one type- B seller and thus at most one uncertified offer in the market. Conditional on seeing two uncertified

Figure 2: Predictions

(a) Point Predictions

Treatment	Uncertified Price	Certified Price	High-Quality Uncertified Quantity	Low-Quality Uncertified Quantity	Certified Quantity
\mathcal{S}^{Pre}	[140,183]	200	[2,10]	2	[0,10]
$\mathcal{H}^{\text{Post}}$	100	200	0	2	10
\mathcal{H}^{Pre}	100	200	0	2	10
$\mathcal{S}^{\text{Post}}$	100	200	0	2	10
$\mathcal{H}_{FI}^{\text{Pre}}$	100	200	0	2	10
$\mathcal{S}_{FI}^{\text{Post}}$	[140,183]	200	[2,10]	2	[0,10]

(b) Predicted Rejections from Pairwise Mann-Whitney-Wilcoxon Tests

	\mathcal{S}^{Pre}	$\mathcal{H}^{\text{Post}}$	\mathcal{H}^{Pre}	$\mathcal{S}^{\text{Post}}$	$\mathcal{H}_{FI}^{\text{Pre}}$	$\mathcal{S}_{FI}^{\text{Post}}$
\mathcal{S}^{Pre}	–	X	X	X	X	
$\mathcal{H}^{\text{Post}}$	X	–				X
\mathcal{H}^{Pre}	X		–			X
$\mathcal{S}^{\text{Post}}$	X			–		X
$\mathcal{H}_{FI}^{\text{Pre}}$	X				–	X
$\mathcal{S}_{FI}^{\text{Post}}$		X	X	X	X	–

5 Experimental Results

Empirical analysis is taken in three steps. We first test the point predictions of the model at the aggregate level and show that the pattern of pairwise hypotheses tests predicted by Predictions 1-4 is consistent with the data. We then delve deeper into the data to study behavior at the group level. Finally, we explore patterns of individual-level learning in Section 5.3. In Appendix B we provide details on a set of additional experiments where we eliminate all moral hazard in the environment. Additional robustness graphs and tables are included in Appendix C.

5.1 Aggregate Behavior

Predictions 1-4 generate a specific set of point predictions regarding the price of certified and uncertified trades, and the composition of trades in each of the six treatments. At an aggregate level, the data largely follows the predicted pattern:

Result 1 *Prices and the composition of trades are largely consistent with the point predictions*

offers, the buyer knows that one of the two offers is from a type- G seller. If he cannot distinguish between the two uncertified offers and the market price for certified units is 200, he would be willing to pay any amount below $100 + 0.5(U^H - U^L) = 150$ for an uncertified unit.

made by Propositions 1-4 in all six treatments. The pattern of rejections in the pairwise hypothesis tests matches the pattern in Predictions 1-4.

Panel (a) of Figure 3 reports average prices and the composition of trades for each of the six treatments using the last six periods of each treatment. The 95% confidence intervals were constructed by first averaging prices and trades at the period level, and then clustering observations by group.

As can be seen by looking at the results in the \mathcal{S}^{Pre} and \mathcal{S}_{FI}^{Post} treatments, prices in these treatments are close to 140, and on average there are more than 2 uncertified low-quality units traded. Both the average price and the composition of trades is consistent with the equilibrium predictions of one of the more inefficient partially-mediated equilibria. They are significantly different to the predictions made by the mediated equilibrium.

Looking at the other four treatments, the price of uncertified trades is close to 110, and slightly above the price of 100 that would be predicted in the mediated equilibrium. There are also slightly more uncertified low-quality units being traded than would be predicted in the mediated equilibrium. Despite the slightly higher than expected price of uncertified trades, the average difference between the price of uncertified and certified units is about 90 points and significantly larger than the certification cost of 60 points. Further, the number of high-quality uncertified goods is not significantly different to zero. These two factors suggest that the mediated equilibrium occurs in these treatments.

The left hand side of panel (b) of Figure 3 shows the p -value of pairwise Mann-Whitney-Wilcoxon tests on the average uncertified price of each treatment. An observation in these tests is the average price of uncertified trade for the last six periods of each treatment.²³ In support of Predictions 1-4, the distribution of uncertified prices in \mathcal{S}^{Pre} and \mathcal{S}_{FI}^{Post} are significantly different to those in the other four treatments. Consistent with Predictions 1-4, we also do not reject the null hypothesis that the distribution of uncertified prices in \mathcal{S}^{Pre} and \mathcal{S}_{FI}^{Post} are the same, nor do we reject the null hypothesis in all pairwise tests between \mathcal{H}^{Post} , \mathcal{S}^{Post} , \mathcal{H}^{Pre} , and \mathcal{H}_{FI}^{Pre} . The right hand side of panel (b) shows the same test for the average number of uncertified high-quality trades. As can be seen, the pattern matches the price data with the exception that there is a significant difference between the \mathcal{H}^{Post} and \mathcal{H}_{FI}^{Pre} treatments.

While we have shown all pairwise tests for completeness, the main predictions of the model are that the uncertified price and the number of uncertified high-quality units traded in the \mathcal{S}^{Pre} and \mathcal{S}_{FI}^{Post} treatments are significantly different to those in treatments \mathcal{H}^{Post} , \mathcal{S}^{Post} , \mathcal{H}^{Pre} , and \mathcal{H}_{FI}^{Pre} . In support of these predictions, the Kruskal-Wallis test of all six treatments rejects that uncertified prices and the number of high-quantity uncertified quantities are drawn from the same distributions at the .01 level (Uncertified Prices: $\chi^2(5) = 18.571$, p -value = .0023; Uncertified High-Quality Units: $\chi^2(5) = 21.18$, p -value = .0007).²⁴

²³Using a clustered version of the Mann-Whitney-Wilcoxon test developed by Datta and Satten (2005) also yields similar results. We report the averaged version as it is a more conservative test.

²⁴The Kruskal-Wallis test does not inform us of which treatments differ. To address this issue, we also considered two alternative specifications. First, we tested \mathcal{S}^{Pre} and \mathcal{S}_{FI}^{Post} separately against the other four treatments using a

Figure 3: Prices and the Composition of Trades in each Treatment

(a) Averages and 95% Confidence Intervals

Treatment	Uncertified Price	Certified Price	High-Quality Uncertified Quantity	Low-Quality Uncertified Quantity	Certified Quantity
\mathcal{S}^{Pre}	146.0 (118.4,173.7)	200.0 (193.8,206.2)	3.75 (1.4,6.1)	2.3 (1.8,2.8)	5.7 (3.2,8.2)
\mathcal{H}^{Post}	115.0 (104.9,125.1)	203.2 (198.5,208.0)	0.1 (0.0,0.1)	4.3 (2.6,6.1)	7.1 (5.2,8.9)
\mathcal{H}^{Pre}	110.4 (100.8,120.1)	202.0 (195.9,208.2)	0.2 (-0.1,0.4)	4.2 (2.7,5.6)	7.1 (5.5,8.6)
\mathcal{S}^{Post}	110.7 (95.7,125.7)	197.5 (191.8,203.2)	0.6 (-0.2,1.4)	2.5 (1.6,3.4)	8.9 (7.3,10.4)
\mathcal{H}_{FI}^{Pre}	119.7 (109.6,129.9)	204.8 (200.4,209.3)	0.5 (-0.1,1.2)	4.0 (3.2,4.1)	7.1 (6.1,8.1)
\mathcal{S}_{FI}^{Post}	138.3 (126.3,150.3)	198.9 (192.2,205.7)	2 (0.8,3.2)	2.1 (1.9,2.2)	7.8 (6.7,9.0)

95% Confidence intervals constructed by averaging prices and quantities to the period level and clustering data at the session level.

(b) p -values from Pairwise Mann-Whitney-Wilcoxon Tests

	Tests Based on Uncertified Prices						Tests Based on Number of Uncertified High-Quality Units						
	\mathcal{S}^{Pre}	\mathcal{H}^{Post}	\mathcal{H}^{Pre}	\mathcal{S}^{Post}	\mathcal{H}_{FI}^{Pre}	\mathcal{S}_{FI}^{Post}	\mathcal{S}^{Pre}	\mathcal{H}^{Post}	\mathcal{H}^{Pre}	\mathcal{S}^{Post}	\mathcal{H}_{FI}^{Pre}	\mathcal{S}_{FI}^{Post}	
\mathcal{S}^{Pre}	–	.04	.04	.02	.05	.52	\mathcal{S}^{Pre}	–	.01	.01	.02	.02	.20
\mathcal{H}^{Post}	.04	–	.26	.52	.63	.01	\mathcal{H}^{Post}	.01	–	.37	.12	.04	.01
\mathcal{H}^{Pre}	.04	.26	–	1.00	.11	.01	\mathcal{H}^{Pre}	.01	.37	–	.40	.25	.01
\mathcal{S}^{Post}	.02	.52	1.00	–	.26	.01	\mathcal{S}^{Post}	.02	.12	.40	–	.94	.02
\mathcal{H}_{FI}^{Pre}	.05	.63	.11	.26	–	.01	\mathcal{H}_{FI}^{Pre}	.02	.04	.25	.94	–	.02
\mathcal{S}_{FI}^{Post}	.52	.01	.01	.01	.01	–	\mathcal{S}_{FI}^{Post}	.20	.01	.01	.02	.02	–

5.2 Tests of Predictions 1-4 at the Group Level

Having established regularities in the data at the aggregate level, we now look at whether Predictions 1-4 hold at the more disaggregate group level. Analysis of this section is broken into three parts. First, we analyze whether Prediction 1 holds by looking at periods 7-12 of experiments 1-4 and determining: (i) whether an unmediated or partially-mediated equilibrium forms in the \mathcal{S}^{Pre} treatment; and (ii) whether the mediated equilibrium forms in the \mathcal{H}^{Pre} treatment. After establishing that Prediction 1 is consistent with the data, we next turn to Predictions 2 and 3 and ask how the utilization of the private-order CEI influences the response of buyers and sellers to ex-

Mann-Whitney-Wilcoxon test. All four of these tests reject the null at a p -value $< .01$. Using the Benjamini-Hochberg procedure to adjust for multiple hypothesis, all tests are significant using a false discovery rate of .05. Second we used Dunn's test for stochastic dominance for each set of 6 treatments and sequentially adjusted the p -values using the Benjamini-Hochberg adjustment to control for false discoveries. 14 of 16 pairwise tests between \mathcal{S}^{Pre} and \mathcal{S}_{FI}^{Post} and the other four treatments were significant and 0 of 14 of the other tests were significant.

ogenous changes in the number of type- G and type- C sellers in the environment. Finally, we study adaptation when the information externality is removed to determine the validity of Prediction 4.

Prediction 1: Initial Convergence A precondition for the rest of the study is that the mediated equilibrium forms in experiments that start in the Hazardous environment and that the unmediated or partially-mediated equilibria reliably form in experiments that start in the Safe environment. We find the following:

Result 2 *Prediction 1 is largely supported by the data. Five of the six groups that start in the Safe environment have trade prices and trade compositions that are consistent with the unmediated or partially-mediated equilibria. Likewise, five of the six groups that start in the Hazardous environment have trade prices and trade compositions that are consistent with the mediated equilibrium.*

Predictions 1-4 generate predictions regarding the relative price of certified and uncertified trades, and on the number of uncertified high-quality units traded. To show these predictions succinctly, it is useful to transform the price data by first calculating ΔP — the average difference between the price of uncertified and certified trades — and then plot each treatment in the space of ΔP and the average number of uncertified high-quality units traded.

In the space of ΔP and the number of uncertified high-quality units traded, groups that trade according to an unmediated or partially-mediated equilibrium in each period are predicted to have $\Delta P \in (13, 60)$ and to have between 2 and 10 uncertified low-quality trades. Groups that trade according to the mediated equilibrium are predicted to have $\Delta P = 100$ and no uncertified low-quality trades.


Figure 4 shows the estimated value of ΔP and the average number of uncertified trades for periods 7-12 of the \mathcal{S}^{Pre} and \mathcal{H}^{Pre} treatments. As can be seen, five of the six groups in the \mathcal{S}^{Pre} treatment have prices and uncertified trades consistent with an unmediated or partially-mediated equilibrium. Five of the six groups in the \mathcal{H}^{Pre} treatment have prices and uncertified low-quality trades consistent with the mediated equilibrium. The remaining group in \mathcal{S}^{Pre} converged to the mediated equilibrium while the last \mathcal{H}^{Pre} group does not fall into either set of equilibria.

The pattern seen in Figure 4 holds in a more formal regression analysis, where the prices of uncertified trades in the \mathcal{S}^{Pre} treatment are compared with the prices of uncertified trades in the \mathcal{H}^{Pre} treatment. Using group fixed effects, we estimate:

$$P_{i,g} = \alpha_0 + \sum \alpha_g + \beta_{Cert} I_{Cert} + \beta_{\mathcal{S}^{Pre}} I_{\mathcal{S}^{Pre}} + \epsilon_{i,g} \quad (5)$$

where $P_{i,s}$ is the price of an individual trade i in group g , α_g are individual group 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, group 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

Figure 4: Difference in Price of Uncertified and Certified Trades and the Number of Uncertified High-Quality Units Traded in the \mathcal{S}^{Pre} Treatment and the \mathcal{H}^{Pre} Treatment



partially-mediated equilibrium to form in the \mathcal{S}^{Pre} environment, we would predict $\alpha_0 = 100$, $\alpha_0 + \beta_{Cert} = 200$, $\alpha_0 + \beta_{\mathcal{S}^{Pre}} \in [140, 183]$.

Table 5 presents regression results from equation (5) 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 is predicted in a partially-mediated equilibrium.²⁵

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 in the \mathcal{S}^{Pre} 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 vary between 102 and 108 and are not statistically significant from the predicted price of 100.²⁶ 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.²⁷

Predictions 2 and 3: Persistence of the Mediated Equilibrium Having established that Prediction 1 holds for at least five of the six groups in each of the two initial environments, we

²⁵The 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.

²⁶Significance based on a Wald test of $\alpha_0 = 100$. p -value = 0.6148 for regression (1) and p -value = 0.1574 for regression (2).

²⁷Significance based on a Wald test of $\alpha_0 + \beta_{Cert} = 200$. p -value = 0.1103 for regression (1) and p -value = 0.8902 for regression (2).

Table 5: Estimation of Prices in Periods 7-12 of the \mathcal{S}^{Pre} Treatment and the \mathcal{H}^{Pre} Treatment

	(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*** (4.636)	107.973*** (5.255)
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 group-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 group level. Robust standard errors in parentheses clustered at the group level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

next look at how the equilibrium that formed in the initial 12 periods adapts to changes in the underlying environment. We find the following:

Result 3 *The prices of uncertified trades in the \mathcal{S}^{Post} treatment are significantly different to the prices observed in the \mathcal{S}^{Pre} treatment, and significantly different to the prices which are predicted in the unmediated and partially-mediated equilibria. Consistent with Predictions 2 and 3, the price of uncertified trades in the \mathcal{S}^{Post} treatment is not significantly different to those in the \mathcal{H}^{Pre} treatment.*

The persistence of the mediated equilibrium is most easily seen by comparing an individual group that began in the Safe environment to one that began in the Hazardous environment. Figure 5 makes this comparison, showing the complete trade history of group 6 and group 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 splits 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 at least two uncertified low-quality units are expected to be traded in all periods.

As can be seen in the top half of Figure 5, a group 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 groups that began in the Safe environment, the uncertified price converges from below to a partially-mediated equilibrium, with a subset of certified

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





Figure 5: Prediction 3 — Persistence of the Mediated Equilibrium

Group 12: Formation and Persistence of the Mediated Equilibrium



• Certified ◊ Uncertified Low ◊ Uncertified High

trades conducted in each period at a premium of 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.

As can be seen in the bottom half of Figure 5, a group that begins in the Hazardous environment converges to the mediated equilibrium 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.

The patterns of adaption and persistence evident in this example is typical of most of the groups.²⁸ Figure 6 shows average uncertified prices for the last six periods of the \mathcal{S}^{Pre} , \mathcal{H}^{Pre} , \mathcal{S}^{Post} , and \mathcal{H}^{Post} treatments. The uncertified price in the \mathcal{S}^{Post} treatment is not significantly different from the \mathcal{H}^{Pre} and \mathcal{H}^{Post} treatments, and is significantly different from the \mathcal{S}^{Pre} treatment, based on Mann-Whitney-Wilcoxon tests of the average price of uncertified trades in the last six periods of each treatment. Using the same test, the \mathcal{H}^{Post} treatment is significantly different to the \mathcal{S}^{Pre} treatment, but not to the \mathcal{H}^{Pre} or \mathcal{S}^{Post} treatments.²⁹

Figure 6: Average Uncertified Prices in the \mathcal{H}^{Pre} , \mathcal{S}^{Post} , \mathcal{S}^{Pre} , and \mathcal{H}^{Post} Treatments




Table 6 shows the results of an extended price regression where we include the last six periods of the \mathcal{S}^{Post} and \mathcal{H}^{Post} treatments. Consistent with Predictions 2 and 3, 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 $\alpha_0 + \beta_{\mathcal{S}^{Post}} = 140$ (p -value < 0.01 for regression (1) and p -value < 0.01 for regression (2)).³⁰

²⁸As noted in the previous section, one of the six groups that began in the Safe environment converged to the mediated equilibrium. 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.

²⁹ p -values for these Mann-Whitney-Wilcoxon tests are shown in panel (b) of Figure 3.

³⁰In Appendix C we show that there is also no difference in the number of high-quality uncertified units traded in the \mathcal{H}^{Pre} and \mathcal{S}^{Post} treatment.

Table 6: Estimation of Prices in the Last Six Periods of the \mathcal{H}^{Pre} , \mathcal{S}^{Post} , \mathcal{S}^{Pre} , and \mathcal{H}^{Post} Treatments

	(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 group level. Robust standard errors in parentheses clustered at the group level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.


Prediction 4: Information and Adaptation Having documented that the mediated equilibrium is persistent in the \mathcal{S}^{Post} treatment, we now study whether giving market participants the information that was suppressed in this equilibrium will lead market participants to adapt. We find the following:

Result 4 *Consistent with Prediction 4, subjects in the full information treatment initially converge to the mediated equilibrium but adapt to a partially-mediated equilibrium when the environment is changed to Safe. 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 groups in experiments 3-5. In the initial Hazardous environments, the distribution of uncertified prices in the \mathcal{H}^{Pre} treatment is not significantly different to uncertified prices in the \mathcal{H}_{FI}^{Pre} treatment based on a Mann-Whitney-Wilcoxon test using the average uncertified price in the last six periods of each treatment as an observation ($z = 1.60$, p -value = 0.11). In the subsequent Safe environment, however, uncertified prices are significantly higher in \mathcal{S}_{FI}^{Post} than in \mathcal{S}^{Post} using the same test ($z = 2.72$, p -value = 0.01). The uncertified price in \mathcal{S}_{FI}^{Post} in the last six periods is also not significantly different to 140 in all six groups using a standard t-test, suggesting that in all groups, a partially-mediated equilibrium was being played in these periods.³¹

³¹Group 13: $t = -0.49$; Group 14: $t = 1.53$; Group 15: $t = 0.84$; Group 16: $t = -1.44$; Group 17: $t = 1.02$; Group 18: $t = -1.86$.

Figure 7: Average Uncertified Prices in the \mathcal{H}^{Pre} , \mathcal{S}^{Post} , \mathcal{H}_{FI}^{Pre} , and \mathcal{S}_{FI}^{Post} Treatments



As before, the potential equilibria of the model make predictions not just about uncertified prices, but also about the relative price of certified and uncertified trades, and on the number of uncertified high-quality units traded. We jointly analyze these predictions in the space of ΔP and the number of uncertified high-quality units.


Figure 8 shows the estimated value of ΔP and the average number of uncertified trades for periods 19-24 of the \mathcal{S}^{Post} and \mathcal{S}_{FI}^{Post} treatments. Groups in the \mathcal{S}_{FI}^{Post} treatment appear to have prices consistent with the partially-mediated equilibrium, but had only 2.03 average uncertified high-quality units traded per period. This level of uncertified high-quality trades is low and suggests that there is at least some hysteresis that is being generated from the initial treatments.³² However, it is not significantly different to the average number of uncertified high-quality units traded in the last six periods of \mathcal{S}^{Pre} , where there was an average of 3.75 uncertified high-quality units traded per period (Mann-Whitney-Wilcoxon Test: $z = -1.29$, p -value = 0.20). These groups also do not differ in their average uncertified prices (Mann-Whitney-Wilcoxon test: $z = -0.64$, p -value = 0.52).

Taken together, the data in the full information treatment suggest 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 as a result of forming a mediated equilibrium in prior play.³³

³²As can be seen in Figure 7, prices in \mathcal{S}_{FI}^{Post} are not stable in the last six periods and the treatments do not appear to have fully converged to any equilibrium in periods 19 and 20. If data is restricted to only the last four periods, the average number of uncertified trades increases to 3.04, and all six groups are within 0.5 units of the partially-mediated equilibrium predictions. We stick to the pre-analysis plan here and show the results from periods 19-24.

³³One reason for history dependence is that in the partially-mediated equilibrium with only a small number of uncertified trades, the expected return that a buyer earns for trading in the uncertified market at prices that can support the partially-mediated equilibrium is small. If there are not multiple buyers who are willing to take reasonable amounts of risk, groups can get “stuck” in less efficient partially-mediated equilibria even if other partially-mediated equilibria could be supported by the risk preferences of the group. The experiment thus suggests that both information and coordination are important factors in the persistence of the mediated equilibrium. We discuss additional

Figure 8: Difference in Price of Uncertified and Certified Trades and the Number of Uncertified High-Quality Units Traded in the \mathcal{S}^{Post} Treatment and the \mathcal{S}_{FI}^{Post} Treatment



5.3 Are individuals learning from Public Signals or Private Experience?

Thus far we have looked at the experiment-level data and seen that the predictions made by the model closely matches the patterns of the observed market data. In this section, we take a more exploratory look at the actions of individual buyers, and document evidence of individual learning from both publicly observed market signals and private experience.³⁴

Result 5 *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

experiments in Appendix B designed to further distinguish between these channels.

³⁴We 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}_{FI}^{Post} treatment relative to \mathcal{H}^{Pre} .

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 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. Figure 9 shows these two Markov transition matrices for the public and private information environments combined, and for these information environments in isolation. Looking first at the pooled data, 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 making a Risky trade 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 < 0.01 ; Errors clustered at the individual level).

Likewise, individuals who last made a Risky trade have a 74.7% chance of continuing to make a Risky trade 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 < 0.01 ; Errors clustered at the individual level).

Proposition 7 of the theory model would predict that the Public Information environment accelerates learning in the unmediated or partially-mediated equilibria, but not in the mediated equilibria. This proposition is weakly supported in the data. Looking at the left hand column of Figure 9, individuals are 13.2 percentage points more likely to continue to make risky trades in the Public Information Environment than in the Private Information Environment when the difference in prices is less than the certification cost. This difference is weakly significant using the same probit regression as above (p -value = 0.10; Errors clustered at the individual level). By contrast, in the case where the difference in certification cost exceeds the certification cost, individuals in the Public Information Environment are 9.8 percentage points less likely to continue to make risky trades, a difference that is not significantly different (p -value = 0.45; Errors clustered at the individual level).

While the difference in the transition matrix may seem small, the difference in switch rates leads to very different proportions of uncertified trades in the aggregate, particularly for risk-neutral subjects. The left hand side of Figure 10 shows the proportion of risky trades in the Public and Private Information Environments of the \mathcal{S}^{Pre} treatment. Individuals who are willing to accept actuarially fair gambles dramatically increase the proportion of risky trades they are willing to take in the Public Information Environment, 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} treatment, the Public Information Environment appears to reduce individuals' propensities


Figure 9: Markov Transition Matrices for Treatments in Public and Private Environments

	Difference in Certified and Uncertified Prices is Less than Certification Cost			Difference in Certified and Uncertified Prices is Greater than Certification Cost		
Both Public and Private Information Environments		Safe	Risky		Safe	Risky
	Safe	0.755	0.235	Safe	0.923	0.077
	Risky	0.253	0.747	Risky	0.505	0.495
Private Information Environment Only		Safe	Risky		Safe	Risky
	Safe	0.743	0.257	Safe	0.920	0.080
	Risky	0.339	0.661	Risky	0.465	0.535
Public Information Environment Only		Safe	Risky		Safe	Risky
	Safe	0.781	0.219	Safe	0.926	0.074
	Risky	0.207	0.793	Risky	0.563	0.437

to experiment and decreases the number of risky trades occurring in the economy.

Finally, 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 percentage points more likely to trade again if they receive a high-quality uncertified unit instead of a low-quality uncertified unit. This difference is significant based on a probit regression, where the dependent variable is 1 if a risky trade is made and 0 otherwise, and the independent variables include the quality of the last risky trade and a dummy variable for the information treatment (p -value < 0.01).³⁵ Likewise, individuals who make a risky trade when $\Delta P > 60$ are 14.2 percentage points more likely to make another risky trade if they receive a high-quality unit (p -value = 0.088).

Figure 10: Proportion of Risky Trades in the Public Information Environment and Private Information Environment.



³⁵Only observations where (i) the last trade was risky and (ii) the difference in average price of other trades is lower than the certification cost are included. Errors are clustered at the individual level.

6 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. When information about the distribution of seller types 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 CEIs 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., S. JOHNSON, AND J. A. ROBINSON (2005): “Institutions as a Fundamental Cause of Long-Run Growth,” 1A, 385–472.
- 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.
- ALLEN, F., E. CARLETTI, I. GOLDSTEIN, AND A. LEONELLO (2015): “Moral Hazard and Government Guarantees in the Banking Industry,” *Journal of Financial Regulation*, 1(1), 30–50.
- AOKI, M. (2001): *Towards a Comparative Institutional Analysis*. MIT Press, Cambridge, MA.
- BAKER, H. K., J. R. NOFSINGER, AND D. G. WEAVER (2002): “International Cross-Listing and Visibility,” *Journal of Financial and Quantitative Analysis*, 37(3), 175–213.
- 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.
- BIANCONI, M., R. CHEN, AND J. A. YOSHINO (2013): “Firm Value, the Sarbanes-Oxley Act and Cross-Listing in the U.S., Germany and Hong Kong Destinations,” *North American Journal of Economics and Finance*, 24(1), 25–44.
- BLUME, A., AND A. ORTMANN (2007): “The Effects of Costless Pre-Play Communication: Ex-

- perimental 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.
- COFFEE, J. (1999): “The Future as History: The Prospects for Global Convergence in Corporate Governance and its Implications,” *Northwestern Law Review*, 93(7), 641–707.
- (2002): “Racing Towards the Top? The Impact of Cross-Listings and Stock Market Competition on International Corporate Governance,” *Columbia Law Review*, 102(7), 1757–1831.
- 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.
- DATTA, S., AND G. A. SATTEN (2005): “Rank-Sum Tests for Clustered Data,” *Journal of the American Statistical Association*, 100(471), 908–915.
- DECK, C., AND N. NIKIFORAKIS (2012): “Perfect and Imperfect Real-Time Monitoring in a Minimal-Effort Game,” *Experimental Economics*, 15(1), 71–87.
- DEMIRGUC-KUNT, A., AND E. J. KANE (2002): “Deposit Insurance Around the Globe: Where Does It Work?,” *Journal of Economic Perspectives*, 16(2), 175–196.
- DIAMOND, D. W., AND P. H. DYBVIK (1983): “Bank Runs, Deposit Insurance, and Liquidity,” *Journal of Political Economy*, 91(3), 401–419.
- DIXIT, A. K. (2004): *Lawlessness and Economics: Alternative Modes of Governance*. Princeton University Press, Princeton, NJ.
- DOIDGE, C., G. A. KAROLYI, AND R. M. STULZ (2009): “Why Do Foreign Firms Leave U.S. Equity Markets?,” Fisher College of Business Working Paper Series.

- 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.
- 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 Communal 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.
- GREINER, B. (2004): “The Online Recruitment System ORSEE 2.0 - A Guide for the Organization of Experiments in Economics,” Working Paper Series in Economics 10, University of Cologne, Department of Economics.
- 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.
- LABY, A. B., AND J. BROUSSARD (2009): "The Competition of Systems in the Market for Listings," in *Economic Law as an Economic Good: Its Rule Function and Its Tool Function in the Competition of Systems*, ed. by K. M. Meessen. Sellier.
- LI YANG, M. (2002): *Essays on Public Finance and Economic Development in a Historical Institutional Perspective*. Ph.D. Dissertation, Stanford University.
- MADIÈS, P. (2006): "An Experimental Exploration of Self-Fulfilling Banking Panics: Their Occurrence, Persistence, and Prevention," *Journal of Business*, 79(4), 1831–1866.
- 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.
- OCHS, J. (1990): "The Coordination Problem in Decentralized Markets: An Experiment," *The Quarterly Journal of Economics*, 105(2), 545–59.
- OGILVIE, S., AND A. CARUS (2014): "Institutions and Economic Growth in Historical Perspective," 2A, 403–513.
- RENY, P. J., AND M. PERRY (2006): "Toward a strategic foundation for rational expectations equilibrium," *Econometrica*, 74, 1231–1269.
- 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.
- SATTERTHWAITE, M., S. R. WILLIAMS, AND K. E. ZACHARIADIS (2015): "Price Discovery Using a Double Auction," SSRN working paper 2463425.
- SHELLING, T. C. (1960): *The Strategy of Conflict*. Harvard University Press, Massachusetts.
- SCHOTTER, A. (1981): *The Economic Theory of Social Institutions*. Cambridge University Press, Cambridge.

- SCHOTTER, A., AND T. YORULMAZER (2009): “On the Dynamics and Severity of Bank Runs: An Experimental Study,” *Journal of Financial Intermediation*, 18, 217–241.
- SMALL, K., AND H. ZHU (2007): “Has Sarbanes-Oxley Led to a Chilling Effect in the U.S. Cross-Listing Market?,” *The CPA Journal*, 77(3).
- SMITH, V. (1964): “Effects of Market Organization on Competitive Equilibrium,” *The Quarterly Journal of Economics*, 78(2), 181–202.
- STULZ, R. M. (1999): “Globalization of Equity Markets and the Cost of Capital,” *Journal of Applied Corporate Finance*, 12(7), 8–25.
- 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.
- VAN HUYCK, J., R. C. BATTALIO, AND R. O. BEIL (1990): “Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure,” *American Economic Review*, 80(1), 234–248.
- 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.
- YOUNG, P. (1993): “The Evolution of Conventions,” *Econometrica*, 61(1), 57–84.
- ZINGALES, L. (2007): “Is the U.S. Capital Market Losing its Competitive Edge?,” ECGI - Finance Working Paper No. 192/2007.

Appendix A: Theory Appendix

Formal Construction of the Rational Expectations Equilibrium

In this section, we formally define the rational expectations equilibria and develop the notation necessary for proving Propositions 1-7. 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 ,

$$u(m, P^m, \lambda_i, s) = \begin{cases} U^H - P^{\mathcal{C}} & \text{if } m \in \mathcal{C}, s \in \{G, C, B\} \\ U^H - P^{\mathcal{NC}} & \text{if } m \in \mathcal{NC}, s \in \{G\} \\ \lambda_i[U^L - P^{\mathcal{NC}}] & \text{if } m \in \mathcal{NC}, s \in \{C, B\}. \end{cases} \quad (6)$$

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:

$$v(m, P^m, b, s) = \begin{cases} P^{\mathcal{C}} - C_s^H - T & \text{if } m \in \mathcal{C}, s \in \{G, C, B\}, \\ P^{\mathcal{NC}} - C_s^H & \text{if } m \in \mathcal{NC}, s \in \{G\}, \\ P^{\mathcal{NC}} - C^L & \text{if } m \in \mathcal{NC}, s \in \{C, B\}. \end{cases} \quad (7)$$

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 (7) to be clear that both buyer and seller utility are defined over matches.

As in the main text, we define the number of sellers of type s as M_s . We further define the number of buyers of type λ_i as N_{λ_i} .

The description of the rational expectations equilibrium³⁶ 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 $D : \mathcal{M} \times \mathcal{B} \rightarrow \mathbb{I}_+$ 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

³⁶This 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.

market m is denoted by $S(m, s)$. An allocation of sellers is a function $S : M \times \{G, C, B\} \rightarrow \mathbb{I}_+$ 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

$$\sum_s u(m, P^m, b, s) \mu_b(m, s), \quad (8)$$

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 (8). Consequently, an equilibrium allocation must assign all buyers of type b to markets that are in the arg max of (8):

$$D(m^*, b) \neq 0 \Leftrightarrow m^* \in \arg \max_m \sum_s u(m, P^m, b, s) \mu_b(m, s) \quad \forall b. \quad (9)$$

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

$$\sum_b v(m, P^m, b, s) \mu_s(m, b), \quad (10)$$

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:

$$S(m^*, s) \neq 0 \Leftrightarrow m^* \in \arg \max_m \sum_b v(m, P^m, b, s) \mu_s(m, b) \quad \forall s. \quad (11)$$

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 is 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 $M_b \geq 1$ and where M_b is common knowledge, 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 $S^{\mathcal{C}}$ as the number of sellers trading in the certified market. Further define $p(\hat{g})$ and $q(\hat{g}|S^{\mathcal{C}}, S^{\mathcal{N}C})$ as the prior and posterior distribution regarding the proportion of good types in the economy, respectively, which both have support over $g \in \{0, \frac{1}{M}, \frac{2}{M}, \dots, \frac{M-M_b}{M}\}$. Finally, let $\mathbb{E}_q \hat{S}(m, s|S^{\mathcal{C}}, S^{\mathcal{N}C})$ be the expected number of sellers of type s in market m based on the posterior $q(\hat{g}|S^{\mathcal{C}}, S^{\mathcal{N}C})$ 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:

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

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

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

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

Analysis of the rational expectations equilibria is simplified by two characteristics of the benchmark environment. First, the sellers 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 (6), 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 the sellers' 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 $\Delta P \equiv P^C - P^{NC}$. Note that $\pi^H(\Delta P, \mathbb{E}(\hat{g})) = \mu_b(\mathcal{NC}, G)$, which is given by:

$$\mu_b(\mathcal{NC}, G) = \begin{cases} \mathbb{E}(\hat{g}) & \text{if } \Delta P < T \\ \frac{M\mathbb{E}(\hat{g}|S^C) - S^C}{M - S^C} & \text{if } \Delta P = T \\ 0 & \text{if } \Delta P > T \end{cases} \quad (12)$$

The conditioning of $\mathbb{E}(\hat{g}|S^C)$ by S^C in the partially-mediated equilibrium 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.

Proofs

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

- 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 $C_C^H \in (C^L, C^L + \overline{P}^C - \underline{P}^{NC} - T)$ 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 (7).

Lemma 2 *When all buyers are risk neutral, in equilibrium:*

- If $\Delta P > T$, then all buyers believe that all type- G sellers will certify their goods and therefore $\pi^H(\Delta P, \mathbb{E}(\hat{g})) = 0$. 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 $\Delta P \leq T$, then 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:

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

It follows:

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

$$\Sigma_s u(\mathcal{N}\mathcal{C}, P^{\mathcal{N}\mathcal{C}}, b_0, s) \mu_b(\mathcal{N}\mathcal{C}, s) = U^L - P^{\mathcal{N}\mathcal{C}}.$$

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 $P^{\mathcal{N}\mathcal{C}} = U^L, P^C = U^H$.

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

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

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

Proof of Proposition 1:

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

- (a) By Lemma 1, $S(\mathcal{N}\mathcal{C}, B) = M_b$, $S(\mathcal{C}, G) = M\mathbb{E}(\hat{g})$, and $S(\mathcal{C}, C) = M(1 - \mathbb{E}(\hat{g})) - M_b$.
- (b) By Lemma 2, if $P^{\mathcal{N}\mathcal{C}} = U^H, P^C = U^L$, $D(\mathcal{C}, \lambda_0) = [0, N_{\lambda_0}] \in \mathbb{I}_+$, $D(\mathcal{N}\mathcal{C}, \lambda_0) = [0, N_{\lambda_0}] \in \mathbb{I}_+$, $D(\emptyset, \lambda_0) = [0, N_{\lambda_0}] \in \mathbb{I}_+$ with $\Sigma_m D(m, \lambda_0) = N_{\lambda_0}$.

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

2. When $\Delta P > T$:

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

If $P^C - P^{\mathcal{N}\mathcal{C}} = (1 - \mathbb{E}(\hat{g}))(U^H - U^L) > T$, then $D(\mathcal{N}\mathcal{C}, \lambda_0) = M, D(\emptyset, \lambda_0) = N_{\lambda_0} - M$ 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, $P^{NC} = U^H - (1 - \mathbb{E}(\hat{g}))(U^H - U^L)$. Thus

$$\mathbb{E}(\hat{g}) = \frac{P^{NC} - U^L}{U^H - U^L} \quad (13)$$

and price is a sufficient statistic for $\mathbb{E}(\hat{g})$. Under the mediated equilibrium, both type- G and type- C individuals certify their product. As they are both in the same market, $P^{NC} = U^L$ and $P^C = U^H$, 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 $\hat{g} \in \{0, \frac{1}{M}, \dots, \frac{M-M_b}{M}\}$ 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

$$\sum_x q(x|\hat{g}_i) \log \left[\frac{q(x|\hat{g}_i)}{q(x|\hat{g}_j)} \right] > 0, \quad (14)$$

where $q(x|\hat{g}_i)$ is the posterior of receiving a good of quality x given the true parameter is \hat{g}_i .³⁷ Expanding condition (14) yields:

$$\hat{g}_i \log \left(\frac{\hat{g}_i}{\hat{g}_j} \right) + (1 - \hat{g}_i) \log \left(\frac{1 - \hat{g}_i}{1 - \hat{g}_j} \right). \quad (15)$$

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 (14) holds. Since $\mathbf{g} \in \{0, \frac{1}{M}, \dots, \frac{M-M_b}{M}\}$, 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 is known, prices and the allocation of sellers to markets do not lead to updating of priors by buyers. Further, buyers who purchase in the certified market receive 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.

³⁷The use of $q(x|\hat{g}_k)$ in this equation is to highlight that there are 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.

Proof of Proposition 5: When $\Delta P = T$, $S(\mathcal{NC}, B) = M_b$, $S(\mathcal{NC}, C) = M_c$, $S(C, B) = 0$, and $S(C, C) = 0$. Further $S(\mathcal{NC}, G) \in \{0, \dots, M_g\}$ and $S(\mathcal{NC}, G) \in \{0, \dots, M_g\}$. Thus in the partially-mediated equilibrium all the sellers in the certified market must be of type G . It follows that $M_g \geq S^C$, and thus $q(\frac{i}{M}|S^C, S^{\mathcal{NC}}) = 0$ for all $i < S^C$.

Proof of Proposition 6: $S^{\mathcal{NC}}$ buyers trade in the uncertified market each period. Thus, there must be at least $S^{\mathcal{NC}}$ buyers whose individual observations T go to infinity as the number of periods goes to infinity. Using the logic in the proof of Proposition 3 above, these individuals learn the probability of receiving a high-quality unit in the uncertified market almost surely. Let y denote this probability. Then, $M_g = yS^{\mathcal{NC}} + S^C$ by Lemma 1.

Proof of Proposition 7: In the partially-mediated equilibrium, only type- G sellers produce high-quality units in the uncertified market. Having observed a proportion y of high-quality uncertified units, a buyer knows that $M_g = yS^{\mathcal{NC}} + S^C$. In the mediated equilibrium, only type- B sellers trade uncertified units and thus $y = 0$.

Appendix B: Data Appendix

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

An interesting question that arises from our main treatments is whether the information externality documented in the main text might be powerful enough to sustain a private-order CEI even if there is no coordination problem and the private-order CEI 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 one type- G seller, three type- C sellers, and two type- B sellers. Unlike our original experiment where the type- B sellers always exist however, we slowly switch type- B sellers to type- C sellers in the first six periods. In the next six periods, we slowly replace type- C sellers with type- G sellers until all sellers are type- G . These transitions are shown in Table 7.

The goal of the No Moral Hazard Treatments is 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 has no moral hazard and where the private-order CEI serves no purpose. The type- B sellers were introduced in the early periods because 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 type- C to type- G as this limited information that sellers might have about each other at any point in time.

Table 7: 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

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 expectations equilibrium, it is strictly dominated by the unmediated equilibrium for all buyers and sellers and thus we predict that it will be eliminated in groups where exogenous information about the distribution of seller types is given.

We ran 12 additional groups of the follow-up No-Moral Hazard treatments across six sessions. Six of these groups were under our original Private Information Environment and are identified as the sessions with “Private Information and No Moral Hazard.” The remaining six groups used the Full Information Environment and are identified as “Full Information and No Moral Hazard.”

Result 6 *2 of 6 Private Information and No Moral Hazard groups 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 groups, 66.1% of trades in periods 19-24 use the private-order CEI. In the Full Information and No Moral Hazard treatment, all groups 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 groups, 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 groups 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 groups with full information and is predicted to be unchanged in the groups with no information.

Figure 11 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 group.

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 two treatments. There are an average of 8.08 certified high-quality units

in the Private Information and No Moral Hazard treatment and an average of 8.38 certified high-quality units in the Full Information and No Moral Hazard treatment, a difference that is not significant based on a Mann-Whitney-Wilcoxon test with period 7-12 trades averaged at the group level ($z = 0.16$, p -value = 0.87). Groups without information had an average of 3.81 uncertified trades (2.94 low-quality units and 0.86 high quality units) while groups 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: $z = -0.48$, p -value = 0.63; low-quality units: $z = -1.604$, p -value = 0.11; high-quality units: $z = 0.40$, p -value = 0.68).

Figure 11: Changes in the composition of trades in response to changes in information



In periods 13-24, by contrast, the distribution of trades between the Private Information and No Moral Hazard treatment and the Full Information and No Moral Hazard treatments differ quite strongly. Looking first at the Private Information and No Moral Hazard treatment, 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 = 0.81). Concentrating on periods 19-24 of the Private Information and No Moral Hazard treatment and looking across groups, there

is a large dispersion in the composition of trades across groups. In two groups, buyers and sellers use the private-order CEI for almost every transaction (69 out of 72 transactions in Group 19 and 66 out of 72 transactions in Group 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 and No Moral Hazard treatment, by contrast, there is a steady increase in the number of uncertified high-quality trades in each of the groups. The average number of certified trades in period 19-24 is only 2.47 and is falling over time in all groups. This average number of certified trades is significantly different to the number of trades in the same periods of the Private Information and No Moral Hazard treatment and is also significantly different to the number of certified trades in periods 7-12 of the same group (Full Information vs. Private Information: $z = -2.89$, p -value < 0.01 ; Periods 19-24 vs. Periods 7-12: $z = -2.89$, p -value < 0.01). In three of the groups, 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 groups 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 Private Information and No Moral Hazard treatment. Even when the environment has no moral hazard, individuals continue to use the private-order CEI for all transactions and buyers in these groups 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.

Additional Figures

Rationality Check: A general concern in all experiments is whether subjects understand their private incentives and are acting in a way that is consistent with these preferences. To check for this type of consistency, we look for cases where a buyer or a seller makes a trade that is unambiguously against their myopic pecuniary incentives. Figure 12 shows that while both buyers and sellers make trades that are in violation of their pecuniary incentives, such trades are rare and occur primarily in early periods. The most common violation in the data is the trade of a uncertified low-quality unit by a type- G seller, which occurs 2.0% of the time. These types of trades increase the risk of trading an uncertified units and reduce the predicted price of uncertified trades if an unmediated equilibrium forms, but do not occur with high enough frequency in any of the groups to change the set of possible equilibria. The majority of buyer violations also occur in early periods and 77% of these trades are within 10 points of the buyer's value suggesting that overall impact of these trades is likely to be small.

Quantity of Trades Time Series In the main text, we showed how prices evolved over time in the last six periods of each group. Here we present the results in terms of quantities.

Result 7 *Consistent with Predictions 2 and 3, there is little improvement in the quality of goods*

Figure 12: Trades Inconsistent with Myopic Pecuniary Incentives

	All Periods		Periods 7-12 and 19-24	
	Number of Trades	Percentage	Number of Events	Percentage
Sellers:				
Low-Quality uncertified unit traded by type-G seller	99	2.0%	40	1.6%
High-quality uncertified unit traded by type-C seller	33	0.7%	8	0.3%
High-quality uncertified unit traded by type-B seller	6	0.1%	1	0.0%
High-quality certified unit traded by type-B seller	15	0.3%	5	0.2%
Buyers:				
Unit traded to a buyer above his high-quality unit value	82	1.6%	33	1.3%
Total Number of Trades in Considered Periods	5031		2530	

traded in the uncertified market when groups that begin in the Hazardous environment are switched to the Safe environment.

Figure 13 shows 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 groups.³⁸

As shown in panel (b), the only significant change in the composition of trades for groups that began in the Hazardous environment is a significant shift away from uncertified low-quality units to certified high-quality units.³⁹ This is most likely a result of weaker incentives for type-*G* sellers to trade uncertified units relative to sellers of type-*C*.


Efficiency: Figure 8 reports the predicted efficiency of each of the potential mediated, partially-mediated, and unmediated equilibrium in the Safe environment. As can be seen, the equilibrium are Pareto ranked with the mediated equilibrium being least efficient and the unmediated equilibrium being most efficient. Based on Predictions 1-4, we would predict that efficiency in $\mathcal{S}^{Post} = 1460$ and that efficiency in this treatment is less than the efficiency in \mathcal{S}^{Pre} and \mathcal{S}_{FI}^{Post} . Figure 9 reports on a simple linear regression where data is restricted to these three treatments and where period earnings is regressed against treatment dummies for the \mathcal{S}^{Pre} and \mathcal{S}_{FI}^{Post} treatments. Errors in this regression are clustered at the session level. As can be seen, the predictions of the model hold,

³⁸There 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.


³⁹Significance is 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 < 0.01 with errors clustered at the group level. A similar regression with uncertified high-quality units as the dependent variable does not yield a significant treatment effect (p -value = 0.12).

Figure 13: Changes in the Composition of Trades in Response to Changes in the Environment

(a) Experiments 1-2 Beginning in the Safe Environment



(b) Experiments 3-4 Beginning in the Hazardous Environment



though efficiency levels are low when compared to the theoretical benchmarks of the unmediated equilibrium. This is in part due to the large number of certified trades that occur in most sessions. The \mathcal{S}^{Pre} and \mathcal{S}_{FI}^{Post} treatments are not significantly different ($F(1, 17) = 1.43$, p -value = .248).

Table 8: Predicted Per-Period Surplus For Each Potential Equilibrium in Safe Environment

Type of Equilibrium	Number of Certified Units	Predicted Surplus Each Period
Mediated	10	1460
Partially-Mediated	8	1580
Partially-Mediated	7	1640
Partially-Mediated	6	1700
Partially-Mediated	5	1760
Partially-Mediated	4	1820
Partially-Mediated	3	1880
Partially-Mediated	2	1940
Partially-Mediated	1	2000
Unmediated	0	2060

Table 9: Efficiency in the \mathcal{S}^{Post} , \mathcal{S}^{Pre} , and \mathcal{S}_{FI}^{Post} treatments

Treatment \mathcal{S}^{Pre}	175.6*** (51.1)
Treatment \mathcal{S}_{FI}^{Post}	106.7*** (33.9)
Constant	1450.8*** (15.0)
Adj. R^2	0.313
Observations	108

*Errors clustered at the session level. Significance levels: *** $p < 0.01$*