Information and the Organization of Markets: Theory and Experiments

Tom Wilkening*,†

June 20, 2012

This paper studies the link between market structure and the informativeness of publicly and privately observed market signals in an environment where sellers are heterogeneous in production costs and prone to moral hazard. Sellers in the market have the option to credibly certify their goods at a cost, giving rise to both a pooling market equilibrium where no sellers use certification and a separating market equilibrium where some sellers certify their goods. The separating equilibrium eliminates important information regarding the underlying distribution of seller types and makes evaluating counterfactual market structures difficult. In a dynamic experimental setting, we study whether the information externality in the separating equilibrium can limit adaptation to the more efficient pooling equilibrium when market conditions improve. Consistent with theory, we find inefficient persistence when the separating equilibrium has initially formed, but efficient adaptation in other situations. These results suggest that information externalities inherent in markets *themselves* can have a profound impact on their adaptability and long-run efficiency.

Keywords: Information, Externalities, Experimental Economics.
JEL Classification Codes: C92, D02, D40, D82

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

†Department of Economics, The University of Melbourne, 619, Economics and Commerce Building, Victoria, 3010, Australia. E-mail: Tom.Wilkening@unimelb.edu.au

1 Introduction

Imperfectly informed individuals use a variety of information sources to make decisions regarding their pricing, purchase and investment decisions. Publicly observed prices are one such source of information that emerges organically from previous transactions in the market itself. By observing the prices at which potentially better-informed individuals are willing to trade, endogenous signals are passed on to less well-informed individuals. These signals provide information that is typically superior to that based on private transactions and exogenous information alone.

This paper studies the link between the organization of markets and the informativeness of these publicly observed market signals in an environment where sellers are heterogeneous and prone to moral hazard. To this end, we consider an environment with experience goods of high and low quality where a costly certification technology is required to guarantee quality. Heterogeneity in production costs divide sellers into three categories: good, bad and conditional. Good sellers have incentives to always produce high-quality units, while bad sellers have incentives to produce low-quality units. Conditional sellers are prone to moral hazard and, depending on the organization of the market, produce either high-quality certified units or low-quality uncertified units.

We first show theoretically that in a rational expectations equilibrium, changes in the adopted market equilibria can significantly alter the information being generated from market primitives and private consumption regarding the underlying distribution of seller types. For some initial distributions of seller types, two rational expectations equilibria exist — pooling and separating — which vary in the adoption of the certification mechanism. These equilibria differ in terms of efficiency and in the informativeness of public market signals and private experience which might be used to update beliefs about the underlying environment.

In the pooling equilibrium, the cost of certifying a unit exceeds the difference in price between certified and uncertified goods. As a result, no seller chooses to certify their product. High- and low-quality products are traded within a single market and conditional sellers produce low-quality units. As there is uncertainty regarding the quality of the product and only good sellers produce high-quality units, the competitive market price contains information about the expected proportion of good sellers. An exogenous decrease in the number of good sellers therefore leads to an *observable* decrease in the competitive price. This decline in price can lead to an arbitrage opportunity for good and conditional sellers by adopting certification and provides a natural channel by which a market may endogenously adopt certification.

In the separating equilibrium the certification technology is adopted by both good and conditional sellers. Consequently, their actions no longer reveal their types and market prices provide no new information. Hence, there is no direct way for individuals to share information necessary to adapt away from a market structure once certification has been adopted.

The model shows that while the separating equilibrium is more informative about the quality of the good at the point of sale, the information available to update beliefs important to selecting the optimal market structure is actually diminished relative to the pooling equilibrium. Thus the organization of the market can have a profound impact on the aggregation of information in environments where there is uncertainty and variation in the market environment over time.

Given the strong theoretical difference in the informativeness of market equilibria, a natural conjecture is that market forces which lead to efficient prices *within* a market will not always select *between* market equilibria efficiently. To explore this idea, laboratory experiments are next used to study equilibrium selection and the persistence of market equilibria in an environment where the underlying population of sellers changes over time. Subjects initially trade in one of two environments — Safe and Hazardous — which vary in the composition of sellers in the market. In the Safe environment, the proportion of good sellers in the market is large, thus favoring the formation of a pooling equilibrium. In the Hazardous environment, good sellers are replaced with conditional sellers, leading to substantial amounts of moral hazard and a predisposition toward a separating equilibrium. Subjects who begin in the Safe environment are switched to the Hazardous environment midway through the experiment. Likewise subjects who begin in the Hazardous environment are switched to the Safe environment.

Consistent with the theoretical model, individuals who begin in the Safe environment establish a pooling equilibrium and then adapt to the separating equilibrium in response to a change in the underlying environment. Subjects who begin in the Hazardous environment form the separating equilibrium and remain in this equilibrium when the environment is changed to Safe. Looking at individual decision making, we find evidence of learning in the pooling equilibrium both through an individual's personal purchase experiences and through his or her observation of other buyers' trades. By contrast, there is little evidence of learning in markets where the separating equilibrium has formed. Taken together, these results provide strong evidence that the market structure can have a large effect on the ability of individuals to learn thereby opening a channel by which long-term inefficient market equilibria can arise even under conditions where market forces efficiently select the optimal market equilibrium in the short run.

The combination of heterogeneous seller costs, moral hazard, and costly certification is a common combination in agriculture markets. In these markets, certification is typically done through large intermediaries who establish a reputation for reselling high-quality products and earn a proportion of the trade surplus. In India's grain market, for instance, it is difficult to differentiate between different types of seed at the time of purchase, making it difficult to assess its quality prior to milling. Some farmers have high quality initial stocks of seed while other farmers have low quality stocks leading to different marginal costs for providing high quality grain.

In 1999, the Agri Business Division of one of India's largest conglomerates, ITC, introduced the e-choupal program which embedded web kiosks and intermediaries into thousands of Indian villages. Along with a number of other services, the kiosk and intermediary provide a single point at which farmers can purchase high-quality farm inputs and have their commodities purchased from their farm door. These services allow farmers to certify their product and reduce the uncertainty of their supply chain, but also embed ITC into each transaction potentially leading to large intermediation costs. While it is likely the case that the introduction of certification in this market improves the welfare of farmers in the short run, the theoretical and empirical results of this paper suggests that ITC's position as intermediary may be permanent despite the overall input mixes of farmers improving through repeated purchase of higher quality seed. This persistence may occur even if ITC is passive and does not strategically exploit its market position.

At a more general level, this paper suggests that care must be taken in developing mechanisms which mitigate moral hazard. As these mechanisms typically take discretion away from individuals and agglomerate the actions of heterogeneous types, they can have negative long-run consequences in environments where the underlying population is stochastic and therefore the optimal institution varies over time. For example, in markets, the persistence of certification institutions may lead to needless verification costs and intermediation.[1] In government, the persistence of regulation can lead to regulatory burden and red tape.[2] In organizations, the persistence of monitoring can lead to a decrease in intrinsic motivation and experimentation.[3]

[1] The Agriculture Marketing Service, for instance, offers voluntary certification programs for a variety of US agriculture goods. Similar decentralized certification institutions exist for management standards, business school accreditation, health and safety management, and some environmental laws. See King, Lenox & Terlaak (2005) for more examples.

[2] The Sarbanes-Oxley Act, for instance, requires that all publicly traded companies implement standardized auditing and risk management as part of an effort to constrain publicly traded firms from taking undisclosed risks. These programs have high fixed costs, however, which potentially limits access to equity markets for small firms. See Section III of the SEC's *Final Report of the Advisory Committee on Smaller Public Companies* (2006).

[3] See Benner & Tushman (2003) for an empirical study of the effect of process management on firm innovation.

While policies, institutions and market structures which eliminate information could be studied in market and non-market settings, analyzing a market setting allows for new insights into the informativeness of prices. As pointed out by Hayek (1945), market prices are important mechanisms for communicating information; the adjustment of prices in response to the decisions of buyers and sellers provides new information which is not available to an observer or central planner. The findings in this paper suggest that in the presence of moral hazard, information from market primitives alone may be insufficient in efficiently organizing markets. This result sheds new light on why the "market for markets" may require auxiliary institutions to operate effectively.

The model presented here relates to the literature on history dependence and herding. History dependent models establish links between actions today and global actions in the future. Multiple equilibria exist due to market frictions (Diamond (1982)), non-convexity in investment costs (Arthur (1994)), imperfect competition (Cooper & John (1988)), expectations (Krugman (1991)), or imperfect reputation (Tirole (1996)). In our setup, history dependence arises due to informational differences between market structures; global coordination to the pooling equilibrium can be welfare improving for every individual in the economy but requires information that is obscured by the market itself.

In the herding literature, pioneered by Banerjee (1992), Bikhchandani, Hirshleifer & Welch (1992), and Welch (1992), the ability to observe the actions of past actors may lead individual agents to follow past play rather than their own signal.[4] This can lead to an information cascade where individuals discard their private signals and all agents continue to make the wrong, inefficient choice. Whereas the herding literature finds that private signals are suppressed by public signals, the model here studies an environment where both private and public signals may be suppressed simultaneously by the market mechanism endogenously adopted in the past.

The current paper is also related to the literature on incomplete learning. Incomplete learning models study environments where the dynamic aggregation of information does not lead to optimal policy and correct beliefs. This can occur in single decision theory problems (McLennan (1984); Berentsen, Bruegger & Loertscher (2008)) or in group decision making (Piketty (1995)) when experimentation is costly and individuals are impatient. This paper proposes a new channel through which incomplete learning can occur; the institutions adopted today carry the seeds of their own persistence by altering the informativeness of market prices and private consumption in the future.[5]

[4]For more general theoretical treatments of herding, see Chamley (1999) and Smith & Sorensen (2000).

[5]In concurrent work, Warren & Wilkening (2011) study a similar channel but concentrate on optimal regulation in a single decision maker framework. In that paper, it is shown that policies that eliminate information are more likely to be persistent since sampling alternative policies is potentially costly. The current

There is a long tradition in experimental economics of studying equilibrium selection and learning.[6] In the current setting, experiments allow for the study of equilibrium selection in a replicable environment where there is exogenous control of supply, demand, information, and the number of equilibria. This allows for an experimental study of market dynamics with minimal assumptions about the strategies of agents. Experiments also allow for the elimination of other channels of persistence, such as the strategic actions of the intermediary, sunk investment, and reputation.

The empirical results of this paper are closest to Brandts & Holt (1992), who find that learning from a sequence of historical interactions has a large effect on equilibrium selection. The theme of history dependence is echoed in the coordination literature (e.g., Cooper, DeJong, Forsythe & Ross (1990), Cachon & Camerer (1996)), where pre-play actions and communication can lead to coordination on Pareto efficient equilibria. It can also be seen clearly in the information cascades literature (e.g. Anderson & Holt (1997); Goeree, Palfrey, Rogers & McKelvey (2007)) which observes at least partial herding in the lab.[7]

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

2 The Model

In this section we determine the rational expectations equilibria for a market with heterogeneity in seller costs and costly certification. We begin by showing that multiple stable equilibria exist, which vary in the use of the certification technology. We then study the informational properties of these equilibria to understand how public and private signals might be used to update beliefs about the underlying distribution of seller costs. We con-

paper studies information externalities in decentralized markets where market structure is endogenous.

[6] See, for instance, Ochs (1990).

[7] History can also matter through learning channels where individuals can apply lessons learned from one game into the next. For instance, in Cooper & Kagel (2008), individuals play two games in which either a unique pooling or a unique separating equilibrium exists for different parameter configurations. Individuals who participated in experiments where the pooling equilibrium existed had faster adaptation speeds to the separating equilibrium in follow up experiments where the parameters were altered. While not the main focus of our experiment, similar increases in adaptation speed (relative to myopic learning) are observed when moving from the pooling to separating equilibrium in treatments which began in the Safe environment. Given this result, it is all the more interesting that there is no adaptation away from the separating equilibrium in treatments that start in the Hazardous environment.

clude by discussing how the lack of updating in the separating equilibria may lead to its overall persistence.

Rational expectations models typically adopt the assumptions of common knowledge about a correct common prior. These assumptions allow for individuals to consistently assess the behaviors of others and ensures that all participants have at their disposal all information needed to optimally react to the play of others. In a dynamic game where individuals are trading over multiple rounds and observe both public prices and private signals from past trade, a fully rational model typically requires extending the model to one with heterogeneous priors and requires infinite regress of assessments and beliefs.[8] The main complexity of these dynamic models is that private signals feed into public signals which in turn change the beliefs and signals of all buyers and sellers.

The approach used in this paper is to look at public and private signals in isolation and show that both channels are uninformative in the separating equilibrium. We start with a static rational expectations model where individuals have a (potentially incorrect) common prior about the distribution of seller types and analyze the set of rational expectation equilibria that exist and might be selected. For each of these equilibria, we then determine the maximal amount of information which could be extracted by an outsider who observes the publicly observable price data, knows the potential types of buyers and sellers, but does not have full information about the common prior. The change in the outsider's posterior represents the maximal informativeness of the market signals. We then study the polar case in which individuals are updating their priors myopically over a sequence of periods but do not use any public information to update their beliefs. This model shuts down the public information channels but allows for beliefs to update over time.

To demonstrate that the separating equilibrium can suppress all information necessary for adaptation, we first show that for pessimistic beliefs about the seller types, only the separating equilibrium exists. Further, when trade occurs in this separating equilibrium, the posterior of all market participants and outsiders is the same as their priors and thus both the public and private signals in the market are uninformative. It follows that if market participants are initially pessimistic, their beliefs do not adjust and they can become stuck in an inefficient equilibrium.[9] By contrast, we show that in the pooling equilibrium, buyers in the market are privy to private information about the value of goods they received and thus the prior and posterior of some traders differ. Even in the case where buyers do not

[8] For a current development of such models see Čopič & Galeotti (2012).

[9] Studying initial beliefs where only the separating equilibrium exists allows us to make explicit assumptions on how individuals select equilibria from one period to the next. Under an additional assumption that individuals select the same equilibrium given the same common beliefs, the market can get stuck in the separating equilibrium for any initial prior.

use public information to improve the precision of their signals, a subset of buyers will have beliefs which converge to the true state and the market price will become fully informative.

2.1 Primitives

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

Each buyer can consume a single high- or low-quality unit. Likewise, each seller can produce a single high- or low-quality unit. Initially we consider the case where there is only one type of buyer denoted by λ_0. Buyers of type λ_0 have gross utilities for consuming the high and low quality good of U^H and 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 of type λ_0 also have a common (though potentially incorrect) prior about the proportion of type-G sellers in the environment. Let $p(\hat{g})$ be the prior distribution regarding the proportion of good types in the economy, which has support over $g \in \{0, \frac{1}{M}, \frac{2}{M}, \dots, \frac{M-1}{M}, \}$ and expected value $\mathbb{E}(\hat{g})$.

The quality of units being traded is initially unknown to buyers. However, sellers have available a costly technology that certifies quality. Certification costs $T \in (0, U^H - U^L)$ and eliminates all uncertainty over the quality of the unit to the buyer. This certification cost is common knowledge and is paid by the seller when a trade occurs. Since $U^H > U^L$, certifying the low-quality unit can not increase its value and thus a certified low-quality unit will never be offered by a profit maximizing firm. Analysis is thus restricted to cases where all certified units are of high quality.

If a seller produces and exchanges a low-quality unit, she pays a cost of C^L which is the same for all sellers. If a seller of type s produces and exchanges a high-quality unit, she pays a cost C^H_s. Types are defined such that

$$C^H_B > C^H_C > C^L > C^H_G \tag{1}$$

and

$$C^H_B > C^L + U^H - U^L - T > C^H_C. \tag{2}$$

Condition 1 distinguishes type-G sellers from the other types because they have incentives to produce high-quality units if they trade in the uncertified market.[10] Condition (2) distinguishes type-B sellers from the other types because they never have an incentive to produce certified goods for any potential set of equilibrium prices. It also ensures that type-C sellers will find it worthwhile to certify their goods along with type-G sellers in the separating equilibrium.

To focus on the most interesting case of the model, two additional assumptions are made on the relative value and cost of units. Let $C^L < U^L$ so that trade is always welfare improving and assume $C_B^H - C^L < U^H - U^L$ so that the social optimum occurs when all three seller types produce high-quality units. Note that because type-B sellers always produce low-quality units, all equilibria are inefficient.

2.2 The Rational Expectation Equilibria

While a formal construction of the rational expectations equilibria is provided in the appendix, an informal construction is included here. It is easiest to think of the certification process as splitting certified and uncertified units into independent markets. Given the choice over certification, buyers and sellers may exchange in three markets $m \in \mathcal{M} = \{\mathcal{C}, \mathcal{NC}, \varnothing\}$, where $\mathcal{C}$ is a market for high-quality certified units, $\mathcal{NC}$ is a market of uncertified units, and $\varnothing$ is a "market" without trades. In the certified market, all three types of sellers produce the high-quality unit and all trade occurs at the price $P^{\mathcal{C}}$. in the uncertified market, a seller is free to exchange a unit of either quality and all trade occurs at the price $P^{\mathcal{NC}}$.

For a given set of prices, sellers optimally select the market which is best for them. As buyers will receive a high-quality unit in the uncertified market only from a type-G seller, buyers form beliefs over this probability which vary in the difference in prices $\Delta P = P^{\mathcal{C}} - P^{\mathcal{NC}}$. Denote this belief as $\pi^H(\Delta P, \mathbb{E}(\hat{g}))$ and note that this belief is a function both of price and the prior. Based on this belief, buyers choose the market which is best for them. A rational expectation equilibrium is one in which supply equals demand in every market and where the buyers have correct beliefs about the type of good supplied by each seller type.[11]

[10] One might question whether type-G sellers are likely to exist in reality. As mentioned in the introduction, many agriculture industries have this characteristics. A wine maker, for instance, with very high-quality grapes and a standardized production line must actively change his practices to produce lower quality wine. Vice versa, firms without initial stocks of high-quality grapes must invest in such inputs to improve the quality of their goods. Another way to view the environment is as a simplification of one where there are two possible certification regimes of different cost and likelihoods of verification. In the low verification environment, type-G sellers strictly prefer to produce high-quality units.

[11] In order to avoid equilibria which are based on out-of-equilibrium belief, the case where $M_b > 0$ is studied so that there always exists at least one seller in the uncertified market. This ensures that $\pi^H(\Delta P, \mathbb{E}(\hat{g}))$ is always well defined. Typically in rational expectations problems it is assumed that all parties have a correct prior about the seller types and that all parties have correct beliefs about the quality of objects they will

Starting with the seller's market selection problem, a seller will prefer to exchanges in the certified market over the uncertified market if:

$$\Delta P \geq T + max(0, C_s^H - C_L). \tag{3}$$

Define $\overline{P}^{\mathcal{C}}$ as the maximum willingness to pay for a certified unit across all buyers. Similarly, define $\underline{P}^{\mathcal{NC}}$ as the minimum willingness to pay across all buyers for an uncertified unit. In the baseline model $\overline{P}^{\mathcal{C}} = U^H$ and $\underline{P}^{\mathcal{NC}} = U^L$. In equilibrium it will be the case that $\underline{P}^{\mathcal{NC}} \leq P^{\mathcal{NC}} \leq P^{\mathcal{C}} \leq \overline{P}^{\mathcal{C}}$ so that i) ΔP is always either zero or positive and ii) both buyers and sellers have incentives to trade in either the certified or uncertified market for prices within these bounds. Given the definition of Good, Conditional, and Bad seller types:

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

- *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 $C_C^H \in (C^L, C^L + \overline{P}^{\mathcal{C}} - \underline{P}^{\mathcal{NC}} - T)$ 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 $C_B^H \geq C_L + \overline{P}^{\mathcal{C}} - \underline{P}^{\mathcal{NC}} - T$. 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.*

Repeating the exercise for the buyers, a buyer's decision to purchase in the certified or uncertified market is based on the probability of receiving a high-quality unit in the uncertified market, $\pi^H(\Delta P, \mathbb{E}(\hat{g}))$. As buyers are on the long end of the market, the utility gained from buying a certified and uncertified unit must be exactly equal to zero and thus equal to each other. This requires that:

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

As buyers rationally predict the actions of the sellers for each set of prices:

Lemma 2 *In Equilibrium:*

purchase. As we are interested in learning about the underlying distribution of seller types, this assumption is relaxed. In our model we require that buyers correctly predict the quality they get from each type of seller they may match with, such that when $\mathbb{E}(\hat{g}) = \mathfrak{g}$, individuals correctly forecast their risk.

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

- *If $\Delta P \leq T$ the buyers believe that all sellers trade in the uncertified market. In this case $\pi^H(\Delta P, \mathbb{E}(\hat{g})) = \mathbb{E}(\hat{g})$ and a risk neutral buyer prefers to purchase the uncertified unit as long as $\Delta P \geq (1 - \mathbb{E}(\hat{g}))(U^H - U^L)$.*

Given lemma 1, lemma 2 and the indifference condition in equation 4, the market has two rational expectations equilibria:[12]

- **Separating Equilibrium:** $P^{\mathcal{C}} = U^H, P^{\mathcal{NC}} = U^L$. 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.

- **Pooling Equilibrium:**[13] $P^{\mathcal{NC}} = U^H - (1 - \mathbb{E}(\hat{g}))(U^H - U^L), P^{\mathcal{C}} = U^H$. 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.

The stability and existence of the two potential equilibria can be seen by plotting the supply of high-quality uncertified units and the demand for uncertified goods as a function of the market prices ΔP. Given the incentives of each seller type, the probability of receiving a high-quality unit in the uncertified market is based on the difference in market prices ΔP. This is shown by the black line in figure 1. When the difference in price of certified units is less than T, type-G sellers trade in the uncertified market and thus the expected proportion of high-quality units in the uncertified market is equal to the expected proportion of type-G sellers, $\mathbb{E}(\hat{g})$. In contrast, when the difference in price between the certified and uncertified market is greater than T, type-G sellers certify their goods and no high-quality units are traded in the uncertified market. The corresponding probability of receiving a high-quality unit in the uncertified market is thus zero.

[12]In general, a partial-pooling equilibrium will also exist where $\Delta P = T$ and type-G sellers are indifferent to trading in the certified and uncertified market. In the baseline model, since (1) all buyers have the same beliefs and utility functions and (2) seller types are discrete, the partial-pooling equilibrium exists only in very special cases. See section 2.3.2 for an extension of the model where partial-pooling equilibria are more robust.

[13]Note that in the Pooling 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.

Likewise, the incentives of buyers to purchase in the certified and uncertified market can also be plotted in the space of ΔP and $\pi^H(\Delta P, \mathbb{E}(\hat{g}))$ as shown by the downward sloping blue line. When the difference in price between certified and uncertified goods is small, buyers always want to purchase certified units and avoid the chance of matching with a bad or a conditional seller. As the difference in price increases, however, buyers are willing to accept some uncertainty to pay a lower price.

Rational expectations equilibria exist in each location where the supply function of high-quality goods intersects the buyers' indifference condition. When the proportion of type-G individuals is small or T is small, the two lines intersect only once and only the separating equilibrium exists. When the proportion of type-G sellers increases, however, a second equilibrium emerges in which no seller certify and high and low-quality units are traded within a single market. It follows:

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

2.3 Market Information

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

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

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

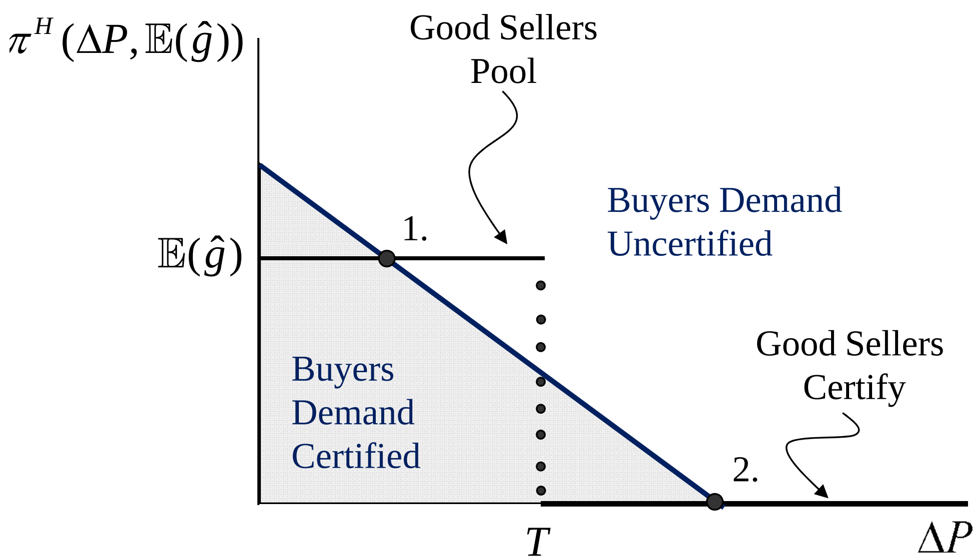

Figure 1: The two potential rational expectations equilibrium as a function of the difference in price for a certified and uncertified good (ΔP) and the corresponding expected probability of receiving a high-quality unit in the uncertified market ($\pi^H(\Delta P, \mathbb{E}(\hat{g}))$). The black line represents the expected proportion of high quality units in the uncertified market based on the optimal actions of the sellers and beliefs of the buyer. When the difference in price of certified units is less than the certification cost T, good sellers trade in the uncertified market leading to an expected proportion $\mathbb{E}(\hat{g})$ of high-quality units in the uncertified market. When the difference in price is greater than T, good sellers certify their goods, eliminating high-quality units from the uncertified market. As buyers are on the long end of the market, the utility gained from buying a certified and uncertified unit must be exactly equal to zero and thus equal to each other. This is represented by the blue line. The pooling and separating equilibrium occur at points (1.) and (2.) respectively. The partial-pooling equilibrium typically does not exist in the case of risk-neutral buyers due to units being traded in integer values.

By contrast, in the pooling equilibrium, the price of uncertified goods, $P^{NC} = U^H - (1 - \mathbb{E}(\hat{g}))(U^H - U^L)$, carries information about the proportion of good sellers. Given only the pooling price and knowledge about U^H and U^L, a new buyer can determine $\mathbb{E}(\hat{g})$.

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

2.3.1 Heterogeneous Beliefs and Learning

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

To begin, let $p^i_t(\hat{g})$ be the prior distribution of buyer i at time t regarding the proportion of good types in the economy which has support over $g \in \{\frac{2}{M}, \frac{1}{M}, ...\frac{M-1}{M}, 1\}$ and where the discrete distribution is single peaked.[14] Further, define the type of an individual by his prior.

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

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

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

Let P^* be the largest $P^{\mathcal{NC}}$ satisfying equation (5) for at least M buyers. Then, if $P^* \geq T + U^L$ a rational expectations equilibrium exists where M buyers trade at the price P^*.[15]

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 of the good 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 pooling equilibrium occurs each period and individuals update their beliefs only from their private purchases. Then there exists at least M buyers such that*

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

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 $\mathfrak{g}$ over time. As P^* is pinned down by the value of the Mth buyer, and his beliefs are accurate, $\mathbb{E}(\hat{g}|P^{\mathcal{NC}}) \longrightarrow \mathfrak{g}$ and thus 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 separating equilibrium, individuals in the market for certified and uncertified goods learn no new information from their purchases since the qualities are guaranteed. Further, the market price carries no information about the priors of the buyers in each period of time. It follows that beliefs regarding the proportion of type-G sellers in the certifying equilibrium may be arbitrary and that there is no reason to expect convergence to true beliefs over time.

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

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

As can be seen from the example and proposition 4, the certifying equilibrium eliminates all information that might be used to update beliefs. Thus, if a market reaches a certifying

[15] As 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.

equilibrium, it is likely to get stuck in this market organization. Further, if there is an exogenous shift in the proportion of type-G and type-C sellers, buyer beliefs will remain unchanged.

2.3.2 Heterogeneous Preferences, Partial Certification, and Public Information

One interesting corollary from the previous discussion is that if a market has converged to a separating equilibrium, *ex post* revelation of uncertified trades does not generate new information about the distribution of seller types. In the case of the pooling equilibrium, this information may increase the rate of convergence but does not provide new information once beliefs of the individuals inside the market have converged.

In an experimental setting, agents typically exhibit some aversion toward accepting actuarially fair gambles. This heterogeneity can lead to a partial-pooling equilibrium where *ex-post* disclosure of trade quality can generate new information. Due to its tractable nature and players' responses to survey questions at the end of the experiment, we model the aversion toward gambles using loss aversion with a reference point of zero.[16] All the results of this section carry over to alternative models using risk or regret aversion.

Suppose that some buyers are loss averse and put a greater weight on aggregate losses than gains. Let $\mathcal{B} = \{\lambda_1, \lambda_2, \ldots, \lambda_N\}$ where λ_i is the idiosyncratic loss aversion parameter for buyer i with $\lambda_i \geq 1$ for $i \in \{1, 2, ..., N\}$ 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 $\lambda_1 \leq \lambda_2 \leq ... \leq \lambda_N$ and again normalize the utility obtained from not trading to zero.

In the pooling equilibrium, the market price $P^{\mathcal{NC}} > U^L$ and there is a potential for losses. When a buyer receives a low quality unit in the pooling equilibrium, his net utility is $-\lambda_i(P^{\mathcal{NC}} - U^L)$ 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, partial separating equilibria may form. Let $S^{\mathcal{C}}$ be the number of certified units in an equilibrium. Then for each $S^{\mathcal{C}} < M_G$, a partial-pooling equilibrium may exist with the following properties:

- **Partial-Pooling Equilibrium:** $P^{\mathcal{NC}} = U^H - T$, $P^{\mathcal{C}} = U^H$. Type-C and type-B

[16] 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. We thus view loss aversion with the status quo of zero profit to be a reasonable assumption of preferences.

sellers produce uncertified low-quality units. $S^{\mathcal{C}}$ type-G sellers produce certified high quality goods. $M_G - S^{\mathcal{C}}$ type-G sellers produce uncertified high quality goods. Buyers $i \in \{1, \ldots, M - S^{\mathcal{C}}\}$ buy uncertified units. $S^{\mathcal{C}}$ other buyers buy certified units.

In the benchmark model, the partial-pooling equilibrium was degenerate 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, partial-pooling equilibrium may be stable since the willingness to pay for uncertified units is decreasing in loss aversion leading to a downward sloping aggregate demand function.

In the partial separating equilibrium, since $P^{\mathcal{NC}} = U^H - T$ and $P^{\mathcal{C}} = U^H$, price alone does not convey information about the proportion of type-G sellers. While a lower bound on the number of type-G sellers can be constructed using the number of sellers in the certified market (where all sellers are of type-G) and on the decision of the M^{th} buyer to trade in the uncertified market, public information about the proportion of high-quality units in the uncertified market can generate new information unavailable from market signals. Information about the proportion of high-quality units traded in the uncertified market in conjunction with the size of the certified market once again allows an outside observer to determine the proportion of type-G sellers in the environment.

3 The Experiment

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

In the remaining sections of the paper, we study an experimental market in which the differences in the informativeness of public and private signals between the two equilibria are predicted to have consequences with regard to the adaptability and efficiency of market organizations. The goal of our design is to begin trade in environments in which the separating and pooling equilibria reliably form and then perturb the underlying distribution of sellers in a way that should be undetectable in the separating equilibrium, but which makes this equilibrium highly inefficient. We study both the way in which markets respond to these perturbations as well as studying how individuals learn in each environment and equilibrium.

3.1 Valuations and Costs

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

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

Table 1: Seller Production Costs

	Uncertified Low Units	Uncertified High Units	Certified High Units
Good	50	30	90
Conditional	50	80	140
Bad	50	130	190

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

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

Buyers and sellers were allowed to trade multiple units in order to increase the thickness of the market and to avoid using passive buyers who might cause noise in the experiment by trying to participate. The supply and demand curves were constructed so that no seller or buyer could change the equilibrium price by more than 10 points by withholding their entire supply or demand from the market. This was small relative to the market prices which ranged from 100 to 200 points. Since no buyer or seller had market power, the separating and pooling equilibrium for the experimental environment are the same as the simplified model of section 2.1.[17]

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 & Villamil (1986) to improve the speed of convergence in markets. Conditional on buying a unit, the valuation of both the high- and low-quality units declined for each unit purchased. Thus, if buyer 1 had purchased a low-quality unit and then purchased a high-quality unit, his valuation for the two units would have been 140 and 220 respectively. The demand functions of buyers four and five were staggered slightly to smooth the aggregate demand function.

Table 2: Buyer Valuations

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

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

3.2 The Trading Mechanism

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

Each exchange was conducted as a double auction.[18] Departing slightly from the design

[17] The set of potential partial-pooling equilibria is slightly smaller in the experimental environment since the loss aversion coefficient for multiple units is from the same individual buyer. However, the price and informational properties of these partial-pooling equilibria remain the same.

[18] A 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 buyers 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

developed by Smith (1964), subjects in this experiment were free to enter the bid and ask queues at any price. Subjects were also free to accept any offer from the opposite side of the market and were not bound to accept the lowest possible price. These changes allowed sellers some flexibility in their pricing strategies and allowed buyers a way to avoid offers that they believed to be of low quality.

In the uncertified market, a seller who posted an offer publicly submitted an asking price and secretly selected the quality of the offered unit. A buyer who bid in the uncertified market publicly submitted a bid price and a quality request. Quality requests in the uncertified market were not binding and a seller who filled a request had the option of supplying either quality good. 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.[19] If a seller transacted in the certified market, either by having an offer accepted or fulfilling a buyer's trade request, she was charged the certification fee of 60 points.

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

In the first three periods of the experiment, each trading period lasted four minutes to allow for subjects to become accustomed to the interface. In the remaining periods, the trading period lasted two minutes.[20]

3.3 Information

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

side of the market.

[19] Buyers were free to request certified low-quality units. In practice, this never occurred.

[20] One 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.

than low-quality units.

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

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

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

3.4 Treatments

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

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

Table 3: Moral Hazard Environments

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

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

Table 4: Treatments

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

As the Hazardous and Safe environments are our main treatment variable, it is useful to discuss their design. The goal of our design was to study the link between market structure and information. Thus, we wanted to begin with two environments where the separating equilibrium was likely to form in one environment and the pooling or partial-pooling equilibria would form in the other. To this end, the Hazardous environment was designed so that, under full information about the distribution of types, only the separating equilibrium existed. Our prediction here was that individuals who started in this environment and traded for uncertified market early on would update their beliefs downward and drive the risk premium past the certification cost. This would lead to the formation of the certifying equilibrium in markets that started in the Hazardous environment. Note that the consistent formation of the certifying equilibrium hinges on the ability of buyers to update their beliefs (either through trade or market signals) from early periods and the assumption that the price of uncertified units would adjust downward as a function of these beliefs. In the Hazardous environment, the predicted separating equilibrium had the following properties:

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

The Safe environment was designed so that under the full information about the distribution of types, the separating equilibrium was extremely unlikely to form or persist. Due to the non-strategic nature of the rational expectation equilibria used as a solution concept, the separating equilibrium is always an admissible outcome as a full-information equilibrium outcome. Nonetheless, the Safe environment was designed so that under full information, if a single type-G sellers switched to the uncertified market, a loss neutral buyer who knew the proportion of agents in each market would be willing to pay $.5U^H + .5U^L$ for an uncertified good and U^H for an uncertified good. Since $U^H - U^L$ was 100 points across all units, the difference in willingness to pay for a certified and an uncertified unit was $.5(U^H - U^L) = 50$. This difference was less than the certification cost of 60 points. Thus under full information, a paired deviation from the separating equilibrium by a seller and risk neutral buyer could eliminate the separating equilibrium. The consistent formation of the pooling and partial-pooling equilibria in the Safe environment hinges on individuals having a high enough initial belief that buyers are willing to trade uncertified units at high prices in early periods and that the distribution of loss aversion was such that at least some buyers were willing to trade uncertified units even if the underlying distribution was known. If all buyers were loss neutral, the pooling and separating equilibrium under the safe environment were as follows:

- **Pooling Equilibria for Safe Environment :** $P^{NC} = 183$. Type-G sellers produce uncertified high-quality units for a surplus of 153 points per unit. Type-B sellers produce uncertified low-quality units for a surplus of 133 per unit. All trades occur in the uncertified market.

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

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

As the separating equilibrium is always an equilibrium regardless of environment, a final concern in the design of treatments is that sellers who have their type changed must wish to reveal this information to the broader environment and shift the market to a different market structure. Comparing the two equilibria, type-G sellers receives a surplus of 153 points in the pooling equilibrium versus 110 points in the separating equilibrium. The type-B seller receives a surplus of 133 points in the pooling equilibrium versus 50 points in the separating equilibrium. Thus, all sellers were better off in the pooling equilibrium and had group incentives to coordinate to this equilibrium.[21] Equilibria were efficiency ranked in the Safe environment with the pooling equilibrium being most efficient and the separating equilibrium being the least efficient. As noted in Table 5, all possible equilibria were inefficient relative to the first best due to inefficient production by the type-B seller.

Table 5: Efficiency

	Perfect Information	Pooling Equilibrium	Separating Equilibrium
Safe	2100*	2060	1460
Hazardous	1700*	1100*	1060

*not supportable as an equilibrium

3.5 Protocol

Subjects in this experiment were drawn from a centralized database comprised of undergraduate students from The University of Zurich and UTH-Zurich. 12 sessions were run each composed of 11 subjects who remained in fixed groups and fixed roles over all 24 periods. Trades were conducted in points and converted to Swiss Francs at the end of the experiment at a conversion rate of 30 points to 1 Swiss franc. A session lasted on average 140 minutes and paid an average of 45 Swiss Francs ($38 at the time of the experiment). The first 40 minutes of each session was devoted to an extensive set of written, oral, and computerized instructions which included a control quiz. All programs for this experiment were written in Z-Tree.[22]

After all 24 periods of the main experiment, aversion to gambles was measured via a series of lottery choices similar to those used in Holt & Laury (2002). Subjects made a series of decisions between a guaranteed return of 90 points and a 50-50 gamble between earning 0 and x, where x varied between 60 and 360 in increments of 30. Individuals were considered averse to gambles if they rejected the 50/50 gamble with high payment of 210. Interpreted

[21]While no conditional sellers existed in the Safe environment, this type of seller also would have preferred the pooling equilibrium.

[22]See Fischbacher (2007) for a description of Z-Tree.

as risk aversion with initial wealth of zero, this corresponds to a $\sigma = .19$ in a CRRA utility function of the form $u(x) = \frac{x^{1-\sigma}}{1-\sigma}$. Interpreted as loss aversion with the earnings from the safe gamble used as the reference point, this corresponds to a loss aversion $\lambda = 1.333$.[23]

4 Experimental Results

The theoretical model predicts that when a market reaches the separating equilibrium, no new information is generated when the number of good and conditional sellers changes in the underlying population. It is our hypothesis that this lack of information may lead to persistence in the separating equilibrium since the market has no way to observe and respond to changes in the underlying distribution of seller types and select a more efficient market structure.

In order to empirically evaluate this conjecture, empirical analysis is taken in two steps. We first establish that absent a pre-existing market structure, the efficient pooling or partial-pooling equilibria forms in the Safe environment, while the separating equilibrium forms under the Hazardous environment. This would suggest that absent an established market structure, buyers and sellers select the most efficient equilibria starting from an uninformed prior. We then turn to our main question of how markets that have established a pooling or separating equilibria adapt to exogenous changes in the number of type-G and type-C sellers in the environment.

For convenience, average price information for the last six periods of the pre and post treatments are included in table 6. The $\mathcal{S}$ and $\mathcal{H}$ letters correspond to the Safe and Hazardous environments while the *Pre* and *Post* superscript correspond to the first and second half of the experiment. The average number of buyers averse to gambles in a session was 2.33. Individual session level data is located in the data appendix.

Table 6: Summary Statistics Across Treatments

	$\mathcal{S}^{Pre}$	$\mathcal{H}^{Post}$	$\mathcal{H}^{Pre}$	$\mathcal{S}^{Post}$
Average Uncertified Price	151.6	116.2	113.2	116.2
Quantity (Observations)	217	158	156	110
Average Certified Price	198.2	203.3	201.4	197.1
Quantity (Observations)	205	254	255	319

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

4.1 Do markets converge to the efficient equilibrium?

4.1.1 Hypothesis and Empirical Strategy

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

To test for initial convergence, we begin by comparing the prices of uncertified trades in the $\mathcal{S}^{Pre}$ environment where the degree of moral hazard is low to the prices predicted in pooling and partial-pooling equilibria. Similarly, we comparing the uncertified price in the $\mathcal{H}^{Pre}$ environment to the price predicted in the separating equilibrium. To allow time for the market to converge, attention is restricted to periods 7-12.[24] Using session fixed effects, we estimate:

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

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

In markets where the separating equilibrium forms, the predicted equilibrium prices for certified and uncertified units are 200 and 100. In markets where the partial-pooling or pooling equilibrium forms, the predicted equilibrium price for uncertified units is between 140 and 183. The predicted price for certified units remains 200. Expecting the the separating equilibrium to form in the $\mathcal{H}^{Pre}$ environment and the pooling or partial-pooling equilibrium to form in the $\mathcal{S}^{Pre}$ environment, the empirical predictions are as follows:

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

The likelihood that the partial-pooling equilibrium should form over the pooling equilibrium is directly tied to the proportion of the buyer population that are unwilling to accept actuarially fair lotteries. As a simple control for aversion toward lotteries, the total number of buyers categorized as lottery averse in the lottery treatment is used. Interacting this

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

number with the safe treatment we further estimate

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

where LA is the total number of buyers in a session who were averse to the lottery and $\overline{LA}$ is the average number of buyers who are averse to lotteries over all sessions. β_{LA} is expected to be negative since treatments with more lottery averse individuals is expected to have lower uncertified prices. As the lottery measurement has been demeaned, it is not expected to have an impact on $\alpha_0 + \beta_{\mathcal{S}^{Pre}}$.

4.1.2 Results

The predicted convergence of the Hazardous treatment to the separating equilibrium and the Safe treatment to the partial-pooling or pooling equilibrium is largely supported in the empirical data. Figure 2 shows the evolution of the average uncertified price of trades in the $\mathcal{S}^{Pre}$ and $\mathcal{H}^{Pre}$ environments over time. The black dots in each period are the average price of uncertified trades in each of the six sessions while the line shows the average of these session averages. As can be seen on the left hand side of Figure 2, the average uncertified price in the Safe environment increases over the 12 periods and falls within the region of prices predicted in the pooling and partial-pooling equilibria in 5 out of 6 of the sessions. The variation in the uncertified price across sessions suggests that individual heterogeneity in risk aversion indeed may be influencing equilibrium selection, a hypothesis we discuss in detail below.[25]

As shown in the right hand side of the figure, the prices in the Hazardous environment fall over time, with average uncertified prices in four of the sessions falling to a price just above 100 and the remaining two sessions having average uncertified prices within 20 points of the benchmark prediction.[26]

The consistency of the data with the predictions in hypothesis 1 can also be seen in the regression analysis. Table 7 presents regression results from equations 8 and 9 with varying degrees of control from the lottery treatment. As can be seen in column (1), the empirical

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

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

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

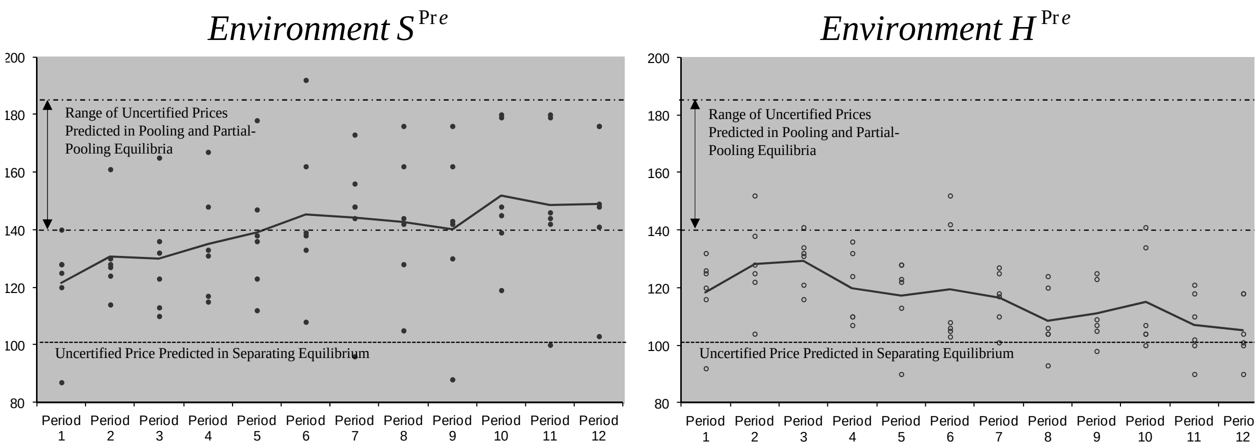

uncertified price ($\alpha_0 + \beta_{\mathcal{S}^{Pre}} = 141.5$) is lower than the predicted pooling equilibrium price of 183 but above the minimum price that could sustain a partial-pooling equilibrium.[27] As can be seen in column (2), the number of lottery averse individuals is negatively correlated with the price of uncertified trades. This is consistent with the theory, which predicts that an aversion to lotteries by the inframarginal buyer will lead to lower overall prices.

Estimated prices for uncertified trades in the $\mathcal{H}^{Pre}$ environment varies between 102 and 108 and is not statistically significant from the predicted price of 100.[28] 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.[29] We summarize the results of the initial 12 periods as follows:

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

[27]The 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.

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

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

Table 7: Hypothesis 1: Convergence of *Pre* Treatments to the Pooling or Separating Equilibrium

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

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

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

4.2 Do market structures adapt to changes in the environment?

4.2.1 Hypothesis and Empirical Strategy

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

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

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

$$\begin{aligned} P_{i,s} &= \alpha_0 + \Sigma\alpha_s + \beta_{LA}(LA - \overline{LA}) * I_{\mathcal{S}^{Pre}} + \beta_{Cert} I_{Cert} \\ &\quad + \beta_{\mathcal{S}^{Pre}} I_{\mathcal{S}^{Pre}} + \beta_{\mathcal{S}^{Post}} I_{\mathcal{S}^{Post}} + \beta_{\mathcal{H}^{Post}} I_{\mathcal{H}^{Post}} + \epsilon_{i,s}, \end{aligned} \tag{10}$$

where $P_{i,s}$ is the price of an individual trade i in session s, α_s are individual session fixed effects, I_{Cert} is an indicator for a certified trade, and $I_{\mathcal{S}^{Pre}}$, $I_{\mathcal{S}^{Post}}$, and $I_{\mathcal{H}^{Post}}$ are indicator variables for uncertified trades in their respective environment. We predict that $\alpha_0 + \beta_{\mathcal{S}^{Pre}} \in [140, 183]$, and $\beta_{\mathcal{S}^{Post}} = \beta_{\mathcal{H}^{Post}} = 0$.

4.2.2 Results

The persistence of the separating equilibrium is most easily seen by comparing an individual session that began in the Safe environment to one that began in the Hazardous environment. Figure 1 makes this comparison, showing the complete trade history of session 6 and session 12. The horizontal dashed lines show the predicted price of the certified and uncertified market in the case of the pooling equilibrium for the $\mathcal{S}^{Pre}$ environment and the separating equilibrium in the case of the other three environments. The vertical dashed lines split trades into six-period increments with the aggregate number of certified and uncertified trades reported at the bottom of each block. Note that in the Safe environment, there is always a single type-B seller. Thus the predicted composition of units without loss aversion is 60 uncertified high-quality units and 12 uncertified low-quality units in the separating equilibrium. The pooling and partial-pooling equilibria do not have a unique trade composition prediction but do require that at least 40% of trades in the uncertified market be high quality under the auxiliary assumption that individuals are not willing to take actuarially unfair gambles.

As can be seen in the top half of Figure 3, a session that begins in the Safe environment converges to the partial-pooling equilibrium in the first 12 periods and then adapts to the separating equilibrium when the environment changes. Typical of all sessions that began in the Safe environment, the uncertified price converges from below to a partial-pooling equilibrium, with a subset of certified trades conducted in each period at a premium 60 points above the prevailing uncertified market price. When the environment changes, sellers who switched from type G to type C sell low-quality units leading to a decrease in price and

Figure 3: Hypothesis 2 — Persistence of the Separating Equilibrium

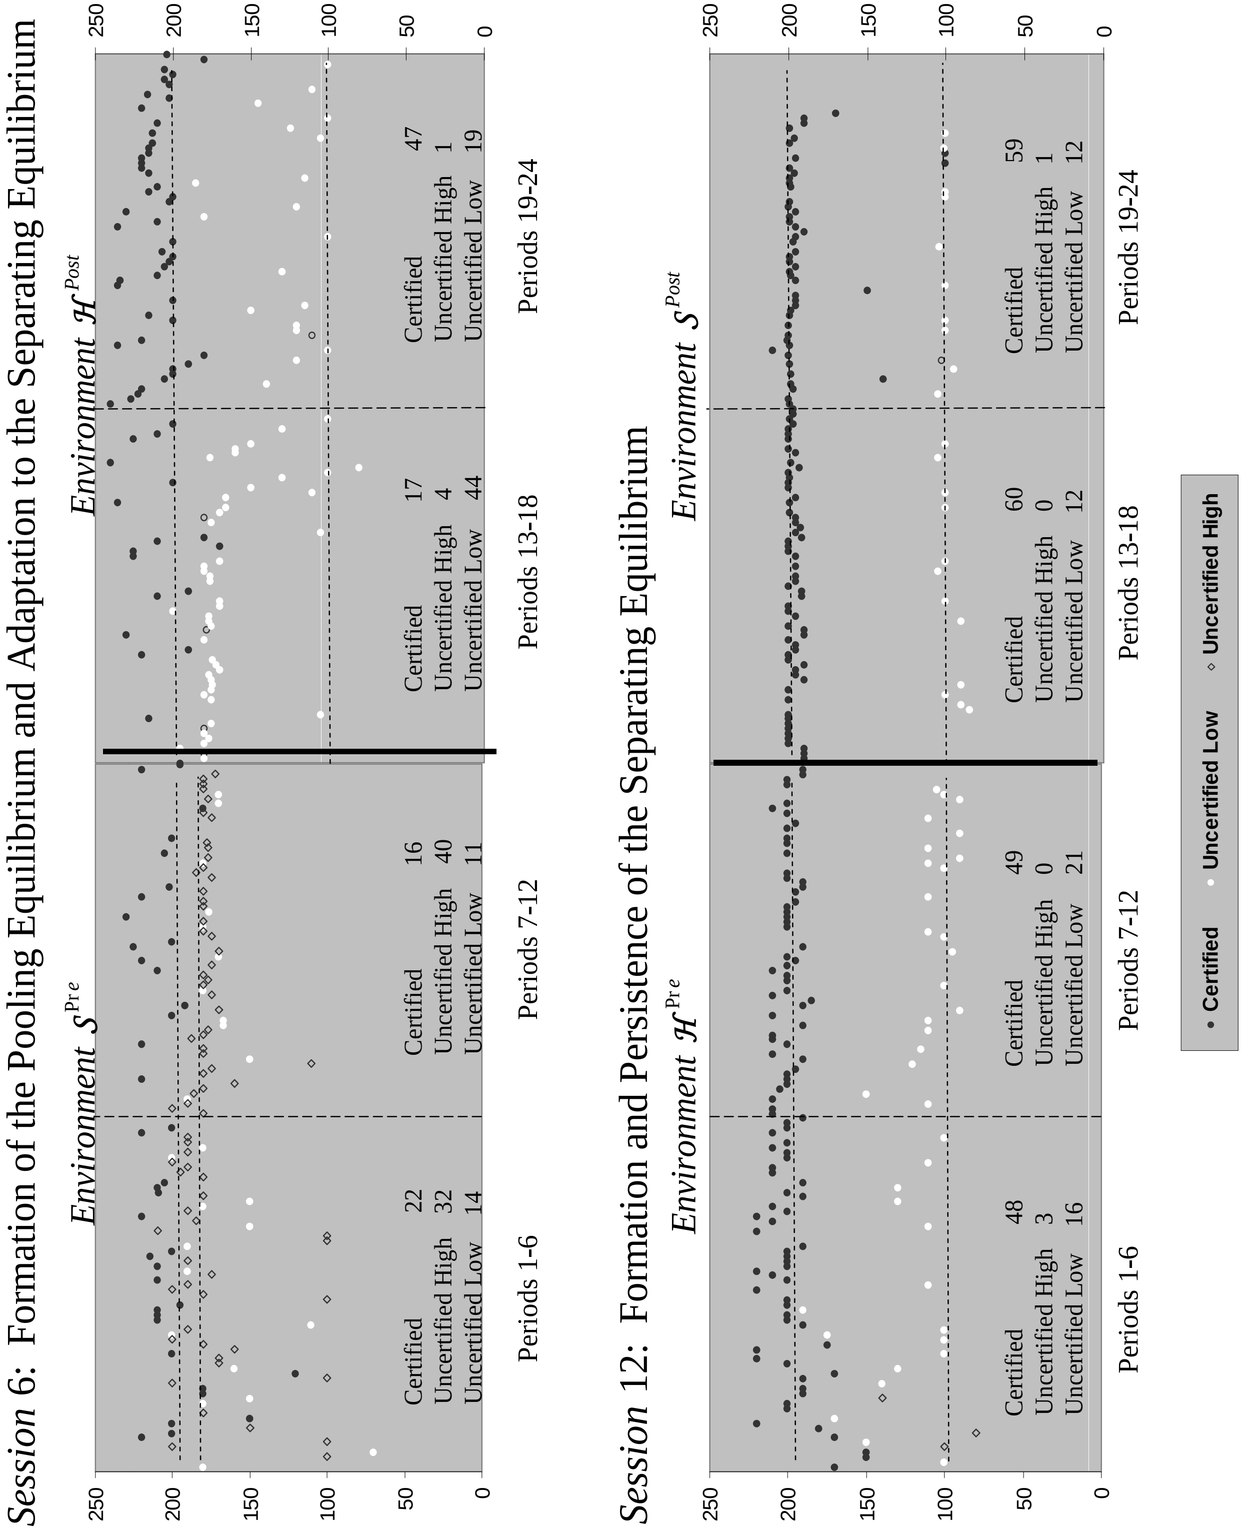

the eventual establishment of a separating equilibrium.

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

The patterns of adaption and persistence evident in this example is typical of most of the sessions.[30] Figure 4 shows average uncertified prices for the last six periods of each environment. Notice that the uncertified price in the $\mathcal{S}^{Post}$ environments is nearly identical to both the $\mathcal{H}^{Pre}$ and $\mathcal{H}^{Post}$ treatments and markedly different from the $\mathcal{S}^{Pre}$ treatment.

Figure 4: Average Uncertified Prices by Environment

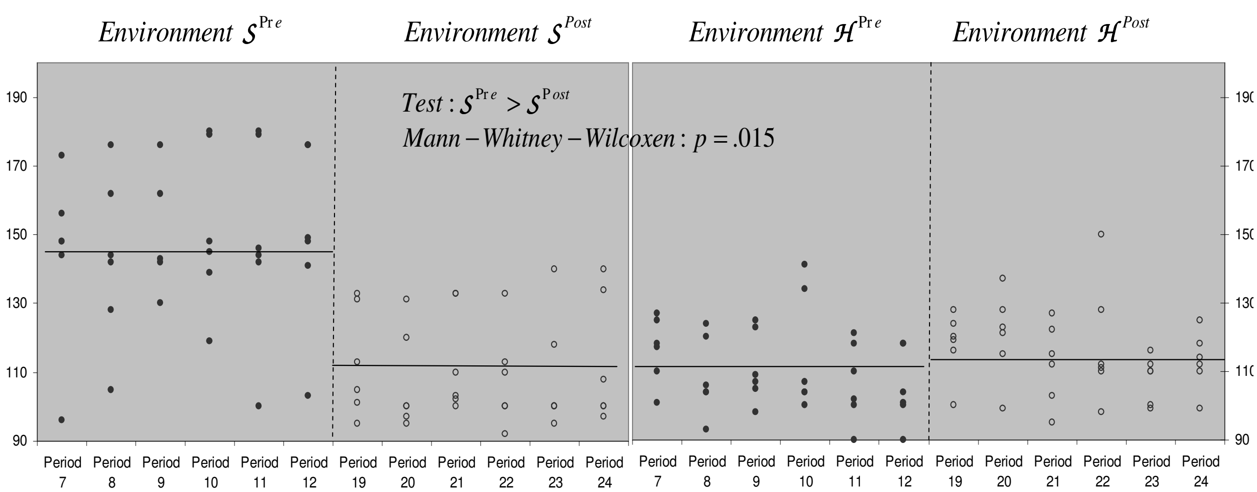

Mann-Whitney-Wilcoxen test conducted on the uncertified price in $\mathcal{S}^{Pre}$ and $\mathcal{S}^{Post}$ averaged at the session level.
Clustered versions of the Rank-Sum Test yield similar results.

Turning to the price regression developed in equation 10, Table 8 extends the original regressions to include periods 18-24 of each session. In support of Hypothesis 2, there is no significant difference between the uncertified prices in the $\mathcal{S}^{Post}$ and $\mathcal{H}^{Post}$ environments relative to the baseline environment of $\mathcal{H}^{Pre}$. Further, the prices in $\mathcal{S}^{Post}$ are significantly lower than those predicted in a pooling or partial-pooling equilibrium.[31] We conclude:

Result 2 *Consistent with hypothesis 2, the price of uncertified trades in the $\mathcal{S}^{Post}$ environment is not statistically significant to those in the $\mathcal{H}^{Pre}$ environment and consistent with*

[30] Time series graphs of all sessions can be found in Appendix B. As noted in the previous section, one of the six markets that began in the Safe environment had the certifying equilibrium form. One of the six markets that began in the Hazardous environment did not appear to converge in the first 12 periods and has a small number of high-quality uncertified trades in the second 12 periods.

[31] Significance based on a Wald test of $\alpha_0 + \beta_{\mathcal{S}^{Post}} = 140$. p-value$< .01$ for regression (1) and p-value$< .01$ for regression (2).

prices predicted in the separating equilibrium. The prices of uncertified trades in the $\mathcal{S}^{Post}$ are also below the prices observed in the $\mathcal{S}^{Pre}$ environment and significantly below the prices which are predicted in the pooling and partial-pooling equilibria.

Table 8: Hypothesis 2: Persistence of the Separating Equilibrium

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

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

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

Figure 5 show the average number of certified and uncertified trades in treatments that start in the Safe environment and the Hazardous environment. Apparent in panel (a), the change in environment from Safe to Hazardous results in an immediate shift from uncertified high-quality units to uncertified low-quality units. Over time, uncertified low-quality units are replaced with certified high-quality units leading to the separating equilibrium in all sessions.[32]

As shown in panel (b), the only significant change in the composition of trades for sessions that began in the Hazardous environment is a shift away from uncertified low-quality units to certified units.[33] This is most likely a result of weaker incentives for type-G sellers to trade uncertified units relative to sellers of type-C. We conclude:

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

(a) Treatments Beginning in the Safe Environment

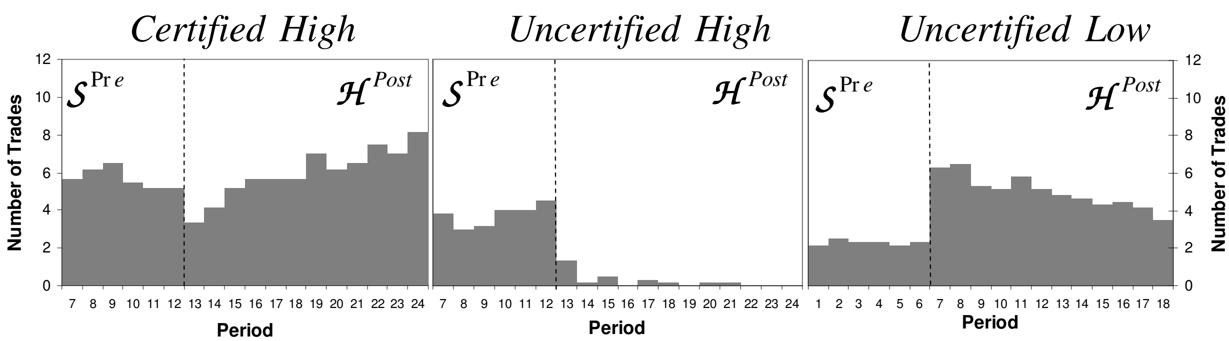

(b) Treatments Beginning in the Hazardous Environment

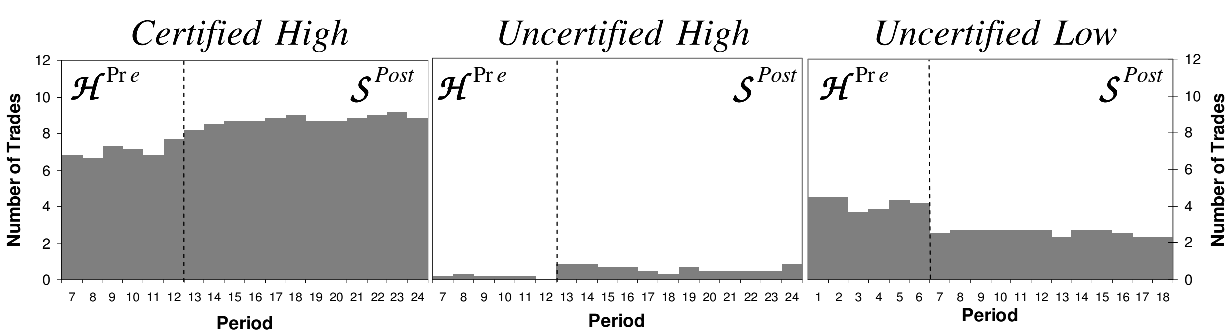

[32]There is also a small but consistent shift of transactions from certified high-quality units to uncertified low-quality units in the two periods following the change in treatment. Recall that in the partial-pooling equilibrium, it may be the case that the type-G sellers are indifferent between trading 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.

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

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

4.3 Are individuals learning from Public Signals or Private Experience?

Thus far we have looked at the aggregate data and seen that the theory model does a good job at rationalizing the observed market data. In this section, we take a more exploratory look at the actions and beliefs of individual buyers and document evidence of individual learning from both publicly observed market signals and private experience.

We begin this section by looking at the beliefs data generated from the bonus game. Recall that buyers and sellers in our experiment were asked to predict the number of type-G sellers in each period. Based on our theory predictions, we would predict that buyers' beliefs are revised downward when the Safe environment is switched to Hazardous environment, but do not change in treatments where the Hazardous environment is changed to Safe.

These predictions have weak support, as can be seen in Table 9, which compares the beliefs of buyers and sellers across treatments. As can be seen in the Buyer's Belief column, beliefs do not increase in the $\mathcal{S}^{Post}$ environment, suggesting that the buyers do not positively increase their beliefs when the environment is switched from Hazardous to Safe. By contrast, there is a significant decrease in beliefs between environments $\mathcal{S}^{Pre}$ and $\mathcal{H}^{Post}$, suggesting increased pessimism when the level of moral hazard in the environment is increased.[34,35] These results are similar but more pronounced for the sellers in the experiment. As four of the six sellers were changed into type-G individuals, the increase in beliefs in the $\mathcal{S}^{Post}$ environment is consistent with the model.

While there is some relation between our beliefs data and the predictions from the model, the buyers' beliefs are extremely noisy at an individual level. 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. This confusion is apparent in the beliefs data, with many subjects guessing randomly over periods. As the beliefs data has

[34] Significance based on a Wald test of $\beta_{\mathcal{S}^{Pre}} = \beta_{\mathcal{H}^{Post}}$. p-value $< .01$.

[35] It should be noted that these results are based on a regression where a linear time trend is removed from the data. While we cannot rule out that this time trend is some sort of learning effect, there are reasons to suspect this is not the case. First, if we split the sample between those who are averse to gambles and those who are not, the time trend is significant only for those who are averse to gambles. As these individuals are the ones who are trading primarily in the certified market and least likely to learn, it is unlikely that the time trend is picking up individual learning. Second, an alternative regression discontinuity design that looks only at the last period of the pre-treatments and the first period of the post-treatments yields results that are similar to the ones shown in Table 9.

Table 9: Hypothesis 3: Beliefs of Buyers and Sellers

	Buyer's Beliefs	Seller's Beliefs
Treatment $\mathcal{S}^{Pre}$	.125	1.726***
	(.242)	(.316)
Treatment $\mathcal{S}^{Post}$	.001	.975**
	(.213)	(.350)
Treatment $\mathcal{H}^{Post}$	-.571	.049
	(.334)	(.257)
Period	.045**	.013
	(.015)	(.232)
Constant	2.11***	110.314***
	(.217)	(3.974)
Fixed Effects[a]	No	No
Adj. R^2	0.025	0.146
Observations	1440	1675

[a]Robust standard errors in parenthesis clustered at the session level. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

the potential for both classical and non-classical measurement error, we take a more direct approach to studying learning in the remainder of this section by looking at the purchase decisions of buyers over time as a function of observable market primitives and their experience in the uncertified market.

In order to study purchase decisions over time, our first step is to generate a Markov transition matrix between (1) actions likely to be taken by individuals with optimistic beliefs about the trade environment and (2) actions likely to be taken when individuals who have pessimistic beliefs about the trade environment. We classify a trade as being made by a buyer with optimistic beliefs if the trade will produce a negative return if a low-quality unit is supplied. These "Risky" trades are those made in the uncertified market with a price greater than an individuals valuation for a blue unit. "Safe" trades are classified as those made in the certified market or trades made in the uncertified market where a profit is guaranteed, as would be the case in the separating equilibrium where the price of uncertified trades is equal to the marginal buyer's valuation.

If price in the market is informative, the Markov transition matrix should have greater switching from Safe trades to Risky trades when the market price for uncertified trades is high. To study this conjecture, we generate two Markov transition matrices: one for trades in a period where the difference in price between the certified trades in a period and the uncertified trades made by other buyers is smaller than the certification cost (i.e., $\Delta P < T$)

and one where the reverse is true. Table 10 shows these two Markov transition matrices over all treatments. As can be seen, when the difference in price is less than the certification cost, individuals who last made a Safe trade have a 23.5% chance of making a Risky trade, while those in an environment where this difference is greater the certification cost, the likelihood of purchasing a Risky asset is only 7.7%.[36] Likewise, individuals who last purchased a Risky trade have a 74.7% chance of continuing to purchase a Risky asset in the next period when the price difference is small relative to a 49.5% chance when the price difference is large.[37]

Table 10: Markov Transition Matrices Between Safe and Risky Trades as a Function of Prices

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

In addition to the role of observable market prices, our data also suggests that the individuals trade experience also play a role in his belief formation. Looking at individuals who made a risky trade last period in a market where the difference in price is less than the certification cost, an individual who receives a high-quality uncertified unit is 20.8% more likely to be willing to trade again, a difference which is significant.[38] Likewise, individuals who make a risky trade when the difference in price is greater than the certification cost are 14.2% more likely to make another risky trade if they receive a high-quality unit[39]

Finally, there is evidence that individuals learn from the composition of trades when the partial-pooling equilibrium forms. The left hand side of Figure 6 shows the proportion of risky trades in the public and private treatments of the $\mathcal{S}^{Pre}$ environment. As can be seen, individuals who are willing to accept actuarially fair gambles dramatically increase the proportion of risky trades they are willing to take, strongly suggesting that they are learning

[36]This difference is significant based on a probit regression which looks at the riskiness of the next trade of an individual following a safe trade with an indicator variable for trades where the difference in average price from other trades is less than the certification cost. Errors clustered at the individual level. p-value $<$ .01

[37]This difference is 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 from other trades is less than the certification cost. p-value $< .01$

[38]Significance based on a probit regression where the left hand side is 1 if a risky trade is made and 0 otherwise, and the right hand side includes the quality of the last risky trade and a dummy variable for the information treatment. Only observations where the last trade was risky and where the difference in average price of other trades is lower than the certification cost are included. Data clustered at the individual level; p-value $< .01$.

[39]p-value $= .098$

from the composition of trades. By contrast, when the separating equilibrium form, as is the case in the $\mathcal{S}^{Post}$, the information treatment appears to reduce experimentation and decrease the number of risky trades which occur in the economy.

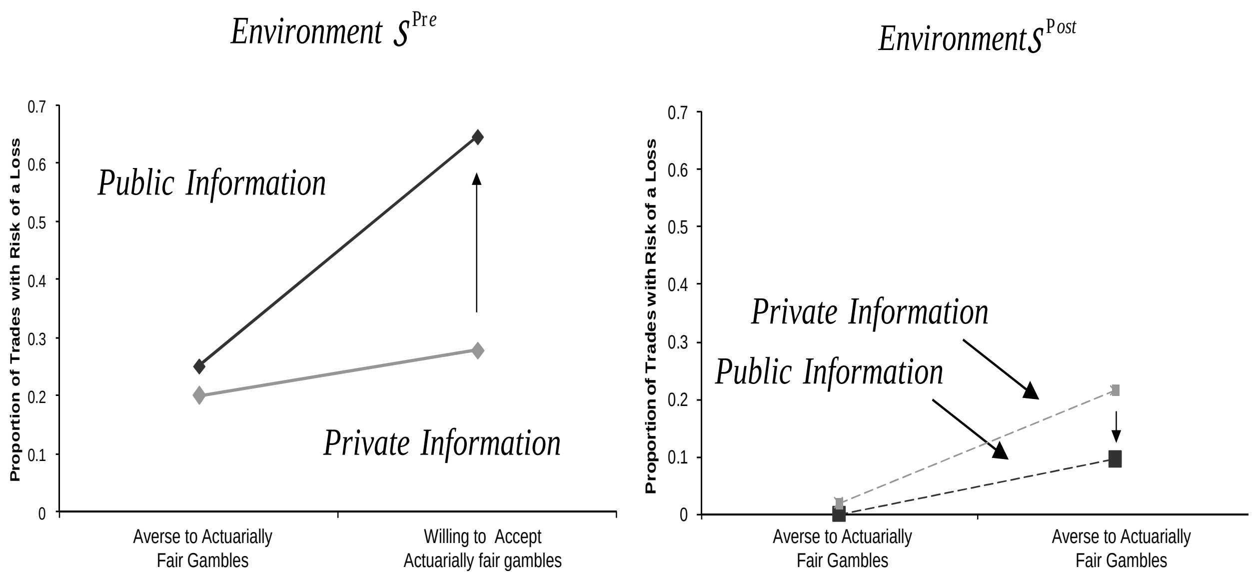

Figure 6: Proportion of risky trades in the information and no information treatments.

Based on the results from the beliefs data, the Markov switching matrices, and the difference in information treatments, we conclude:

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

5 Conclusion

This paper represents a first step in studying the relationship between the organization of markets and the informativeness of publicly observed market signals. We showed formally that, in a market where certification has been adopted endogenously, 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 certification institution. In laboratory experiments, the inefficient persistence of the separating equilibrium was striking. Without exception, markets that adopted certification failed to respond to a change in the underlying distribution of seller types. This failure to adjust led to efficiency losses when compared to sessions which where participants were unhindered by the early adoption of certification.

The experiments described in this paper constitute a stable baseline on which to guide future theoretical and experimental work. We showed that in a double auction environment with anonymity, the benchmark model performed extremely well in predicting both initial convergence and adaptation. We further demonstrated that for some initial distribution of seller types, both the pooling and separating equilibrium were stable. Building on the consistency of these initial experiments, future research will focus on the types of information necessary to adapt away from the separating equilibrium and on the dynamic learning processes that generate persistence.

The information externality highlighted in this paper represents a general phenomenon that extends beyond the simple certification market considered here. Common mechanisms 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

Anderson, L. & Holt, C. (1997), 'Information cascades in the laboratory', *American Economic Review* **87**(3), 797–817.

Arthur, B. (1994), *Increasing Returns and Path Dependence in the Economy*, University of Michigan Press, Ann Arbor, Michigan.

Banerjee, A. (1992), 'A simple model of herd behavior', *The Quarterly Journal of Economics* **107**(3), 797–818.

Benner, M. & Tushman, M. (2003), 'Exploitation, exploration, and process management: The production dilemma revisited.', *Academy of Management Review* **28**(2), 238–256.

Berentsen, A., Bruegger, E. & Loertscher, S. (2008), 'Learning, public good provision, and the information trap', *Journal of Public Economics* **92**(5), 998–1010.

Bertrand, M., Duflo, E. & Mullainathan, S. (2004), 'How much should we trust differences-in-differences estimates?', *The Quarterly Journal of Economics* **119**(1), 249–275.

Bikhchandani, S., Hirshleifer, D. & Welch, I. (1992), 'A thoery of fads, fashion, custom, and cultural change as informational cascades', *Journal of Political Economy* **100**(5), 992–1016.

Brandts & Holt (1992), 'An experimental test of equilibrium dominance in signaling games', *American Economic Review* **82**(5), 1350–1365.

Cachon, G. & Camerer, C. (1996), 'Sunk costs and forward induction in experimental coordination games', *The Quarterly Journal of Economics* **111**(1), 165–194.

Chamley, C. (1999), 'Coordinating regime switches', *The Quarterly Journal of Economics* **114**(3), 869–905.

Cooper, D. & Kagel, J. (2008), 'Learning and transfer in signaling games', *Economic Theory* **34**, 415–439.

Cooper, R., DeJong, D., Forsythe, R. & Ross, T. (1990), 'Selection criteria in coordination games: Some experimental results', *American Economic Review* **80**.

Cooper, R. & John, A. (1988), 'Coordinating coordination failures in keynesian models', *The Quarterly Journal of Economics* **103**(3), 441–63.

Čopič, J. & Galeotti, A. (2012), 'Nash equilibrium, rational expectations, and heterogeneous beliefs', mimeo.

Diamond, P. A. (1982), 'Aggregate demand management in search equilibrium.', *The Journal of Political Economy* **90**(5), 881–894.

Final Report of the Advisory Committee on Smaller Public Companies (2006), Report April 23, 2006, SEC.

Fischbacher, U. (2007), 'z-tree: Zurich toolbox for ready-made economic experiments', *Experimental Economics* **10**(2), 171–178.

Gale, D. (1992), 'A walrasian theory of markets with adverse selection', *Review of Economic Studies* **59**(2), 229–255.

Goeree, J., Palfrey, T., Rogers, B. & McKelvey, R. (2007), 'Self-correcting information cascades', *Review of Economic Studies* **74**(3), 733–762.

Hayek, F. A. (1945), 'The use of knowledge in society', *American Economic Review* **35**(4), 519–530.

Holt, C., Langan, L. & Villamil, A. P. (1986), 'Market power in oral double auctions', *Economic Inquiry* **24**, 107.

Holt, C. & Laury, S. (2002), 'Risk aversion and incentive effects', *American Economic Review* **5**, 1644–1655.

King, A., Lenox, M. & Terlaak, A. (2005), 'The strategic use of decentralized institutions: Exploring certification with the iso 14001 management standard', *Academy of Management Journal* **48**(6), 1091–1106.

Krugman, P. (1991), 'History versus expectations', *The Quarterly Journal of Economics* **106**(2), 651–667.

McLennan, A. (1984), 'Price dispersion and imcomplete learning in the long run', *Journal of Economic Dynamics and Control* **7**(3), 331–347.

Ochs, J. (1990), 'The coordination problem in decentralized markets: An experiment', *The Quarterly Journal of Economics* **105**(2), 545–59.

Piketty, T. (1995), 'Social mobility and redistributive politics.', *The Quarterly Journal of Economics* **110**(3), 551–584.

Smith, L. & Sorensen, P. (2000), 'Pathological outcomes of observational learning', *Econometrica* **68**(2), 371–398.

Smith, V. (1964), 'Effects of market organization on competitive equilibrium', *The Quarterly Journal of Economics* **78**(2), 181–202.

Tirole, J. (1996), 'A theory of collective reputation (with applications to the persistence of corruption and firm quality)', *The Review of Economic Studies* **63**(1), 1–22.

Warren, P. & Wilkening, T. (2011), 'Regulatory fog: The informational origins of regulatory persistence', mimeo.

Welch, I. (1992), 'Sequential sales, learning and cascades', *The Journal of Finance* **47**(2), 695–732.

6 Appendix

6.1 Formal Construction of the Rational Expectations Equilibrium

In this section, we formally define the rational expectation equilibrium and develop the notation necessary for proving Propositions 1-3. 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 $m \in \{\mathcal{C}, \mathcal{NC}, \varnothing\}$ 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,

$$u(m, P^m, \lambda_i, s) = \begin{cases} U^H - P^{\mathcal{C}} & \text{if } m \in \mathcal{C}, s \in \{G, C, B\} \\ U^H - P^{\mathcal{NC}} & \text{if } m \in \mathcal{NC}, s \in \{G\} \\ \lambda_i[U^L - P^{\mathcal{NC}}] & \text{if } m \in \mathcal{NC}, s \in \{C, B\}. \end{cases} \tag{11}$$

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:

$$v(m, P^m, b, s) = \begin{cases} P^{\mathcal{C}} - C^H_s - T & \text{if } m \in \mathcal{C}, s \in \{G, C, B\}, \\ P^{\mathcal{NC}} - C^H_s & \text{if } m \in \mathcal{NC}, s \in \{G\}, \\ P^{\mathcal{NC}} - C^L & \text{if } m \in \mathcal{NC}, s \in \{C, B\}. \end{cases} \tag{12}$$

Note that the sellers value is independent of the buyer type in which she is matched. We leave the parameter b in the left hand side of equation 12 to be clear that both buyer and seller utility are defined over matches.

The description of the rational expectations equilibrium[40] is comprised of three parts: an attainable allocation (D, S), a belief system μ, and a price system P.

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

[40] This 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.

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 $\mu_b(m, s) : \mathcal{M} \times \{G, C, B\} \to \mathbb{R}_+$ satisfies $\Sigma_s \mu_b(m, s) = 1$ for every m and $\mu_s(m, b) : \mathcal{M} \times B \to \mathbb{R}_+$ satisfies $\Sigma_b \mu_s(m, b) = 1$ for every m.

Price System: A price system is a function $P : \mathcal{M} \to \mathbb{R}_+$. For convenience, we define $P^{\mathcal{C}}, P^{\mathcal{NC}}, P^{\varnothing}$ 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

$$\Sigma_s u(m, P^m, b, s)\mu_b(m, s), \tag{13}$$

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 (13). Consequently, an equilibrium allocation must assign all buyers of type b to markets that are in the arg max of (13):

$$D(m^*, b) \neq 0 \Leftrightarrow m^* \in \arg\max_m \Sigma_s u(m, P^m, b, s)\mu_b(m, s) \quad \forall b. \tag{14}$$

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

$$\Sigma_b v(m, P^m, b, s)\mu_s(m, b), \tag{15}$$

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 constrains of the market she has entered. Like the buyer, any competitive equilibrium requires:

$$S(m^*, s) \neq 0 \Leftrightarrow m^* \in \arg\max_m \Sigma_b v(m, P^m, b, s)\mu_s(m, b) \quad \forall s. \tag{16}$$

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

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

As in the main text, we restrict attention to the case where there is exactly one type-B sellers 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 $S^{\mathcal{NC}}$ 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 $q(\hat{g}|S^{\mathcal{C}}, S^{\mathcal{NC}})$ as the prior and posterior distribution regarding the proportion of good types in the economy, which has support over $g \in \{0, \frac{1}{M}, \frac{2}{M}, \ldots, \frac{M-1}{M}, \}$. Finally, let $\mathbb{E}_q\hat{S}(m, s|S^{\mathcal{C}}, S^{\mathcal{NC}})$ be the expected number of sellers of of type s in market m based on the posterior $q(\hat{g}|S^{\mathcal{C}}, S^{\mathcal{NC}})$ 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:*

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

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

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

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

Analysis of the rational expectation 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

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

which is the requirement that all sellers enter the market where the difference between price and the cost of their constrained optimal production choice is largest. Second, since all buyers share the same utility function given in equation (11), only beliefs about $\mu_b(\mathcal{NC}, G)$, the probability of matching with a type-G seller in the uncertified market, affect utility. Since seller's actions only depend on prices, we define a function $\pi^H(\Delta P, \mathbb{E}(\hat{g}))$ where $\pi^H : P \to [0,1]$ is a buyer's belief about the proportion of high-quality units in the uncertified market for a difference in prices of $\Delta P \equiv P^{\mathcal{C}} - P^{\mathcal{NC}}$. Note that $\pi^H(\Delta P, \mathbb{E}(\hat{g})) = \mu_b(\mathcal{NC}, G)$, which is given by:

$$\mu_b(\mathcal{NC}, G) = \begin{cases} \mathbb{E}(\hat{g}) & \text{if } \Delta P < T \\ \frac{M\mathbb{E}(\hat{g}|S^{\mathcal{C}}) - S^{\mathcal{C}}}{M - S^{\mathcal{C}}} & \text{if } \Delta P = T \\ 0 & \text{if } \Delta P > T \end{cases}. \tag{17}$$

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

6.2 Proofs

Proof. Lemma 1: By the definition of $s \in \{G, C, B\}$, $C_B^H \geq C_L + U^H - U^L - T \geq C_C^H \geq C_L \geq C_G^H$. Thus, in the uncertified market, only type-G sellers will produce high-quality goods. Writing out the utility of the seller:

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

By Definition 1,

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

Finding the points where each seller type is indifferent between the certified and uncertified markets lead directly to Lemma 1. ■

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

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

It follows:

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

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

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

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

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

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

■

Proof. Proposition 1:

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

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

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

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

2. When $\Delta P > T$:

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

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

If $P^{\mathcal{C}} - P^{\mathcal{NC}} = (1 - \mathbb{E}(\hat{g}))(U^H - U^L) > T$, then $D(\mathcal{NC}, \lambda_0) = M, D(\varnothing, \lambda_0) = N_{\lambda_0} - M$ 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. Proposition 2: When a pooling equilibrium exists, $P^{\mathcal{NC}} = U^H - (1 - \mathbb{E}(\hat{g}))(U^H - U^L)$. Thus

$$\mathbb{E}(\hat{g}) = \frac{P^{NC} - U^L}{U^H - U^L} \tag{18}$$

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

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

$$\Sigma_x q(x|\hat{g}_i) log \left[\frac{q(x|\hat{g}_i)}{q(x|\hat{g}_j)} \right] > 0, \tag{19}$$

where $q(x|\hat{g}_i)$ is the posterior of receiving a good of quality x given the true parameter is $\hat{g}_i$.[41] Expanding condition (19) yields:

$$\hat{g}_i log \left(\frac{\hat{g}_i}{\hat{g}_j} \right) + (1 - \hat{g}_i) log \left(\frac{1 - \hat{g}_i}{1 - \hat{g}_j} \right). \tag{20}$$

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 (19) holds. Since $\mathfrak{g} \in \{0, \frac{1}{M}, \ldots, \frac{M-1}{M}\}$, convergence is guaranteed as $t \to \infty$.

Returning to the original problem, M buyers purchase each period. Thus, there must

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

be at least M individuals whose individual observations T go to infinity as the number of periods goes to infinity. ■

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

6.3 Appendix B: Time Series Graphs for All Treatments

Session 1: Safe/Hazardous Treatment. No Information

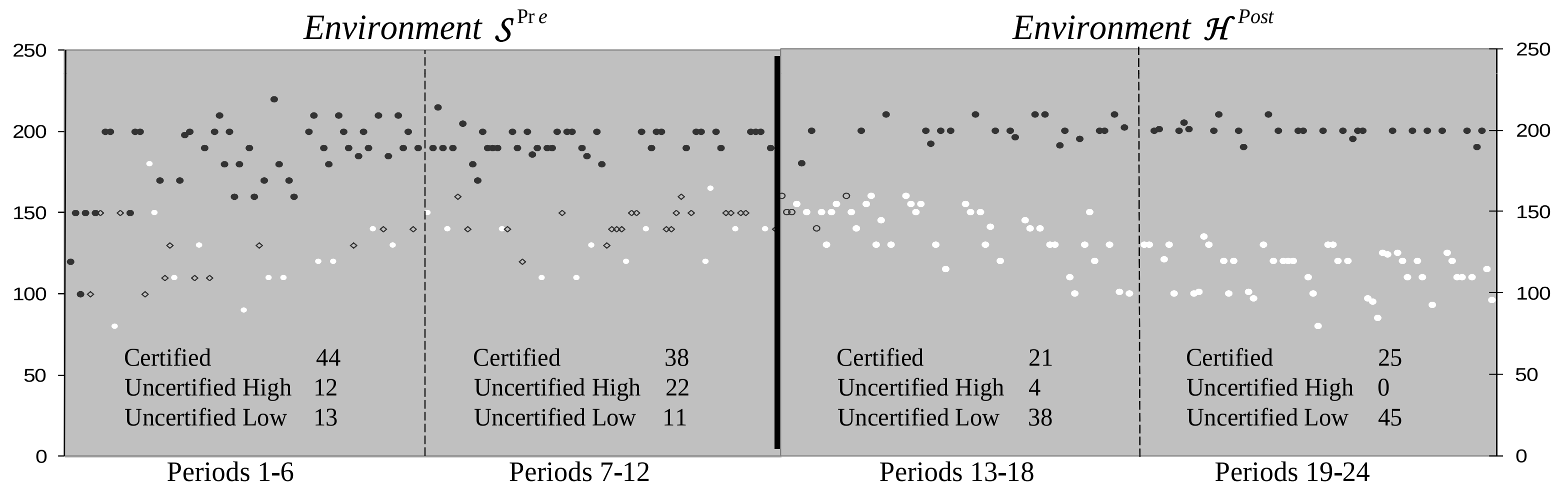

Session 2: Safe/Hazardous Treatment. No Information

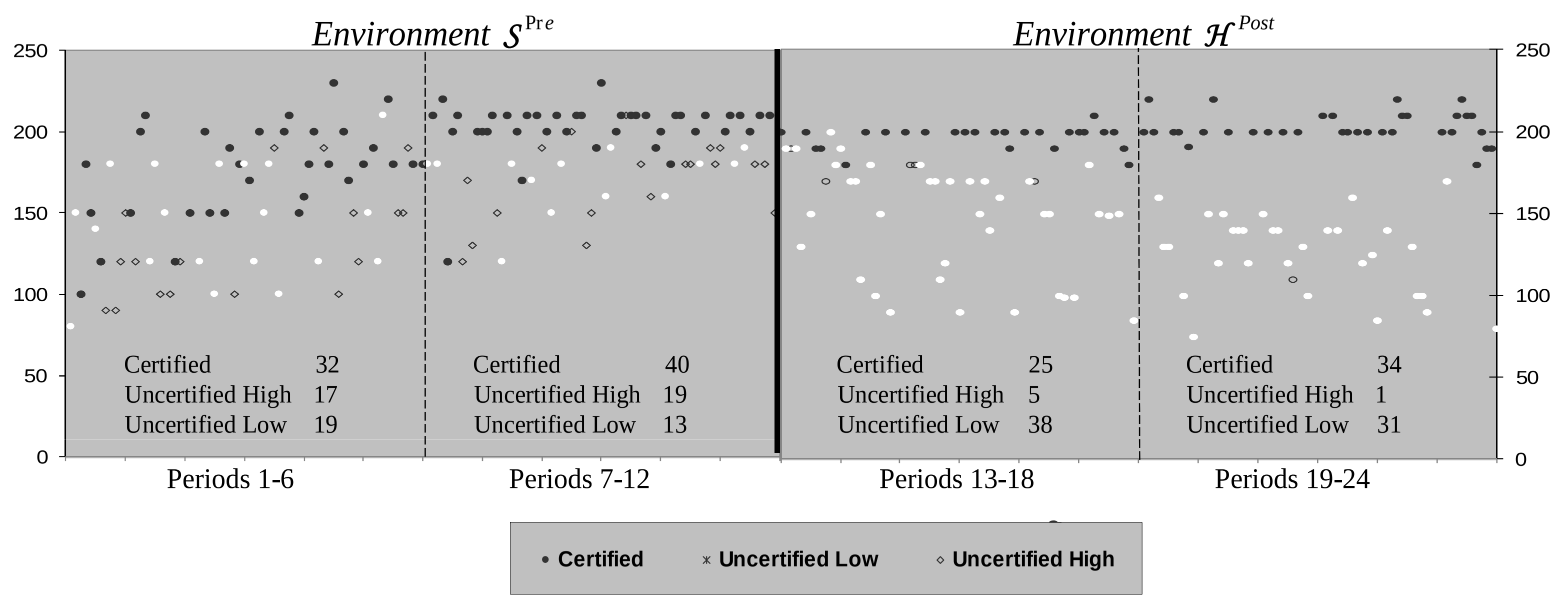

Session 3: Safe/Hazardous Treatment. No Information

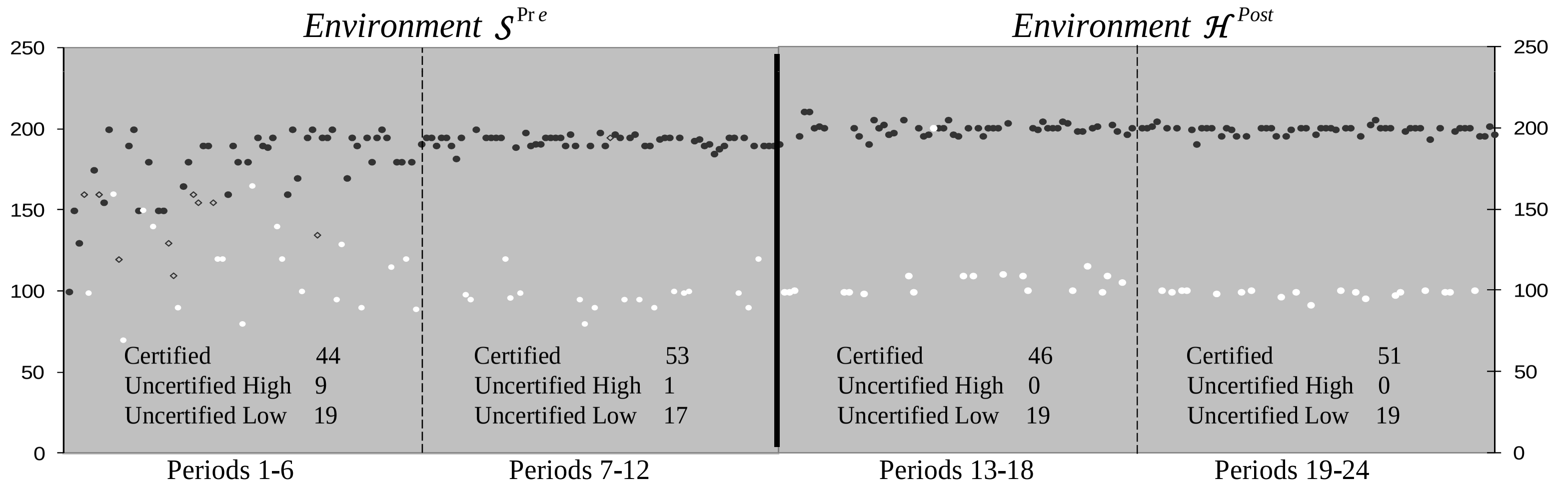

Session 4: Safe/Hazardous Treatment. Public Information

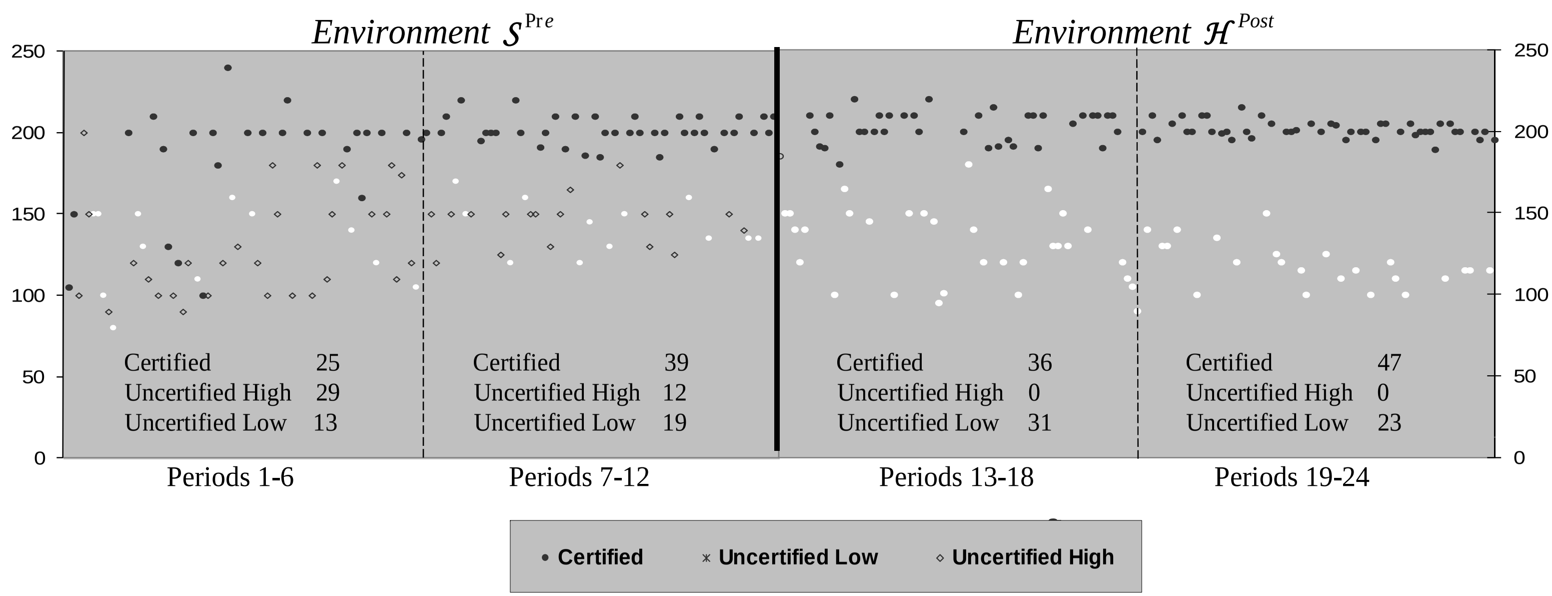

Session 5: Safe/Hazardous Treatment. Public Information

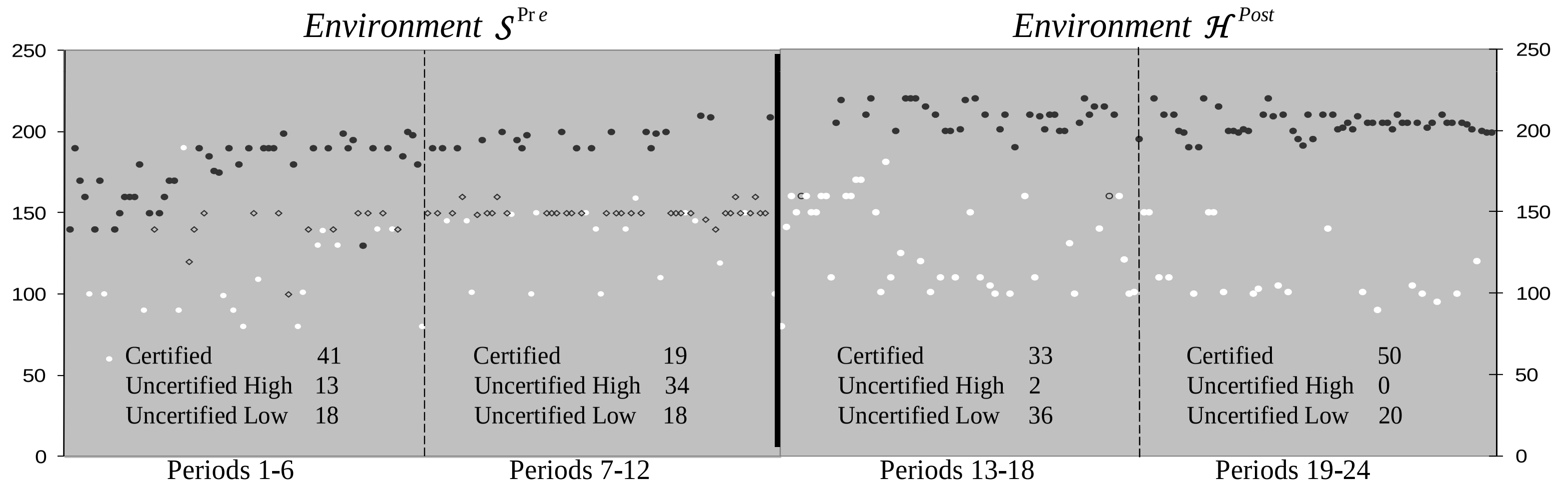

Session 6: Safe/Hazardous Treatment. Public Information

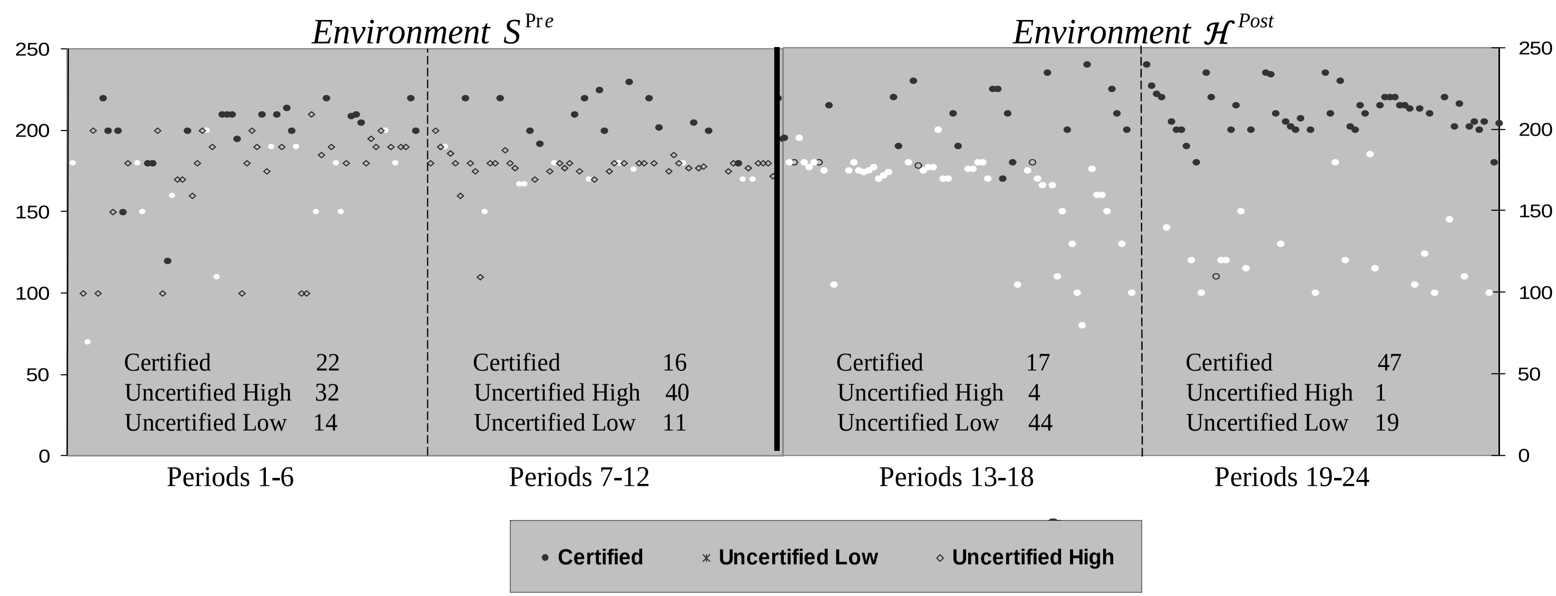

Session 7: Hazardous/Safe Treatment. Private Information

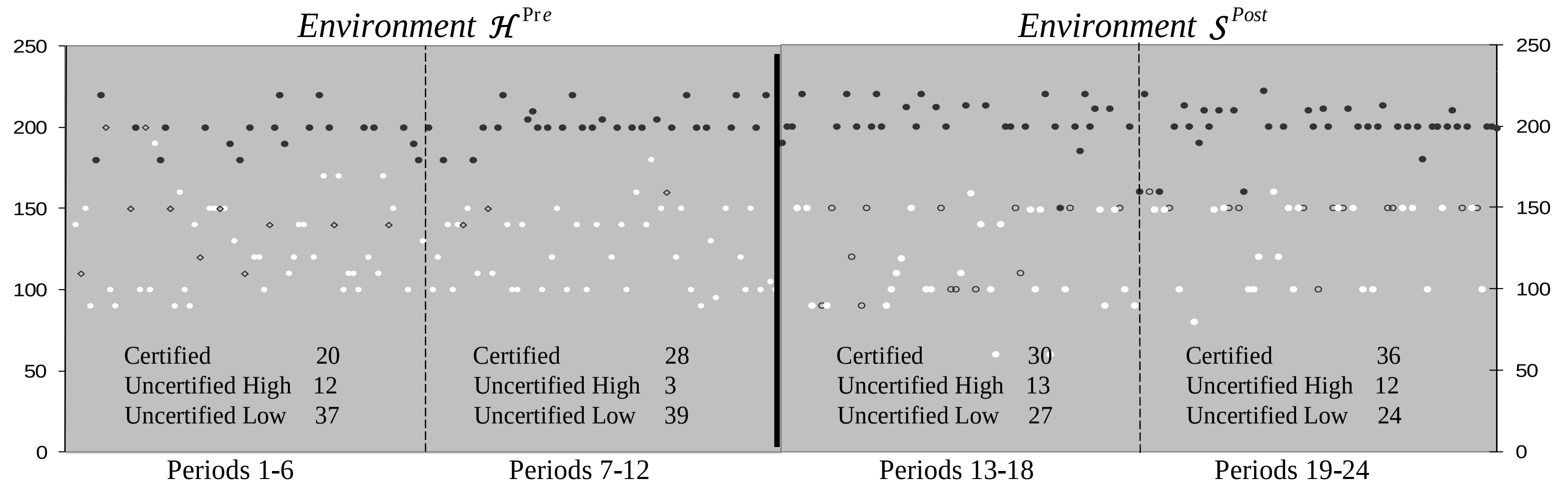

Session 8: Hazardous/Safe Treatment. Private Information

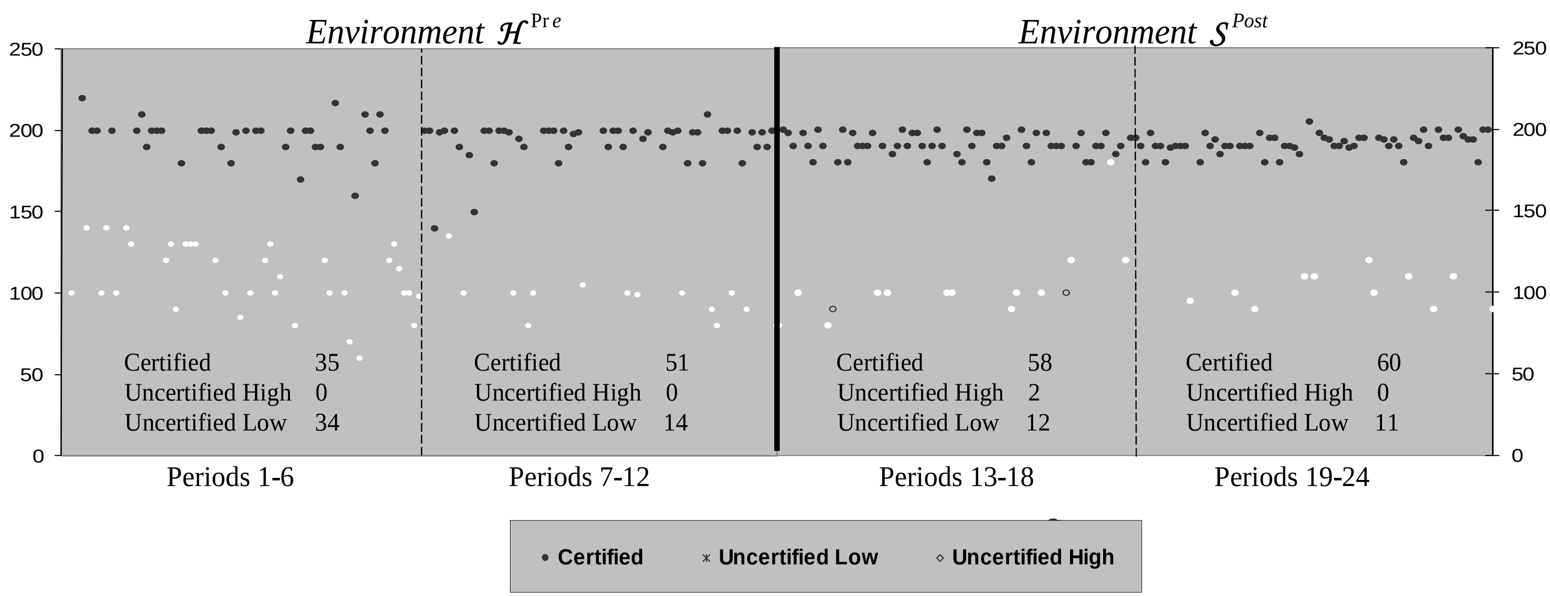

Session 9: Hazardous/Safe Treatment. Private Information

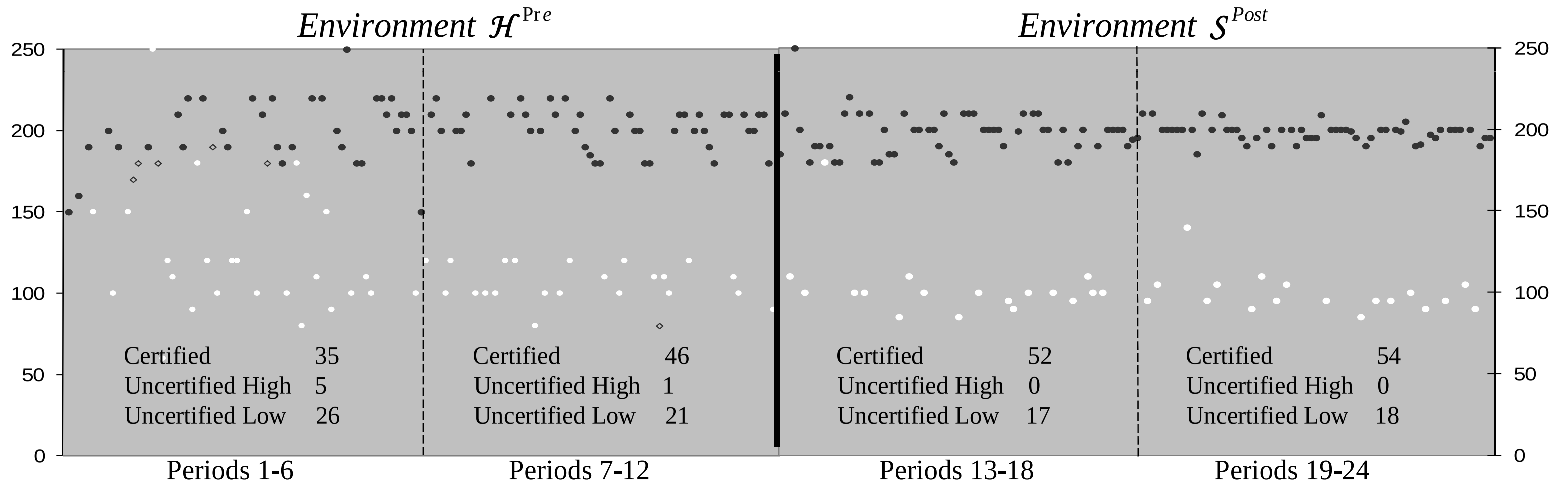

Session 10: Hazardous/Safe Treatment. Public Information

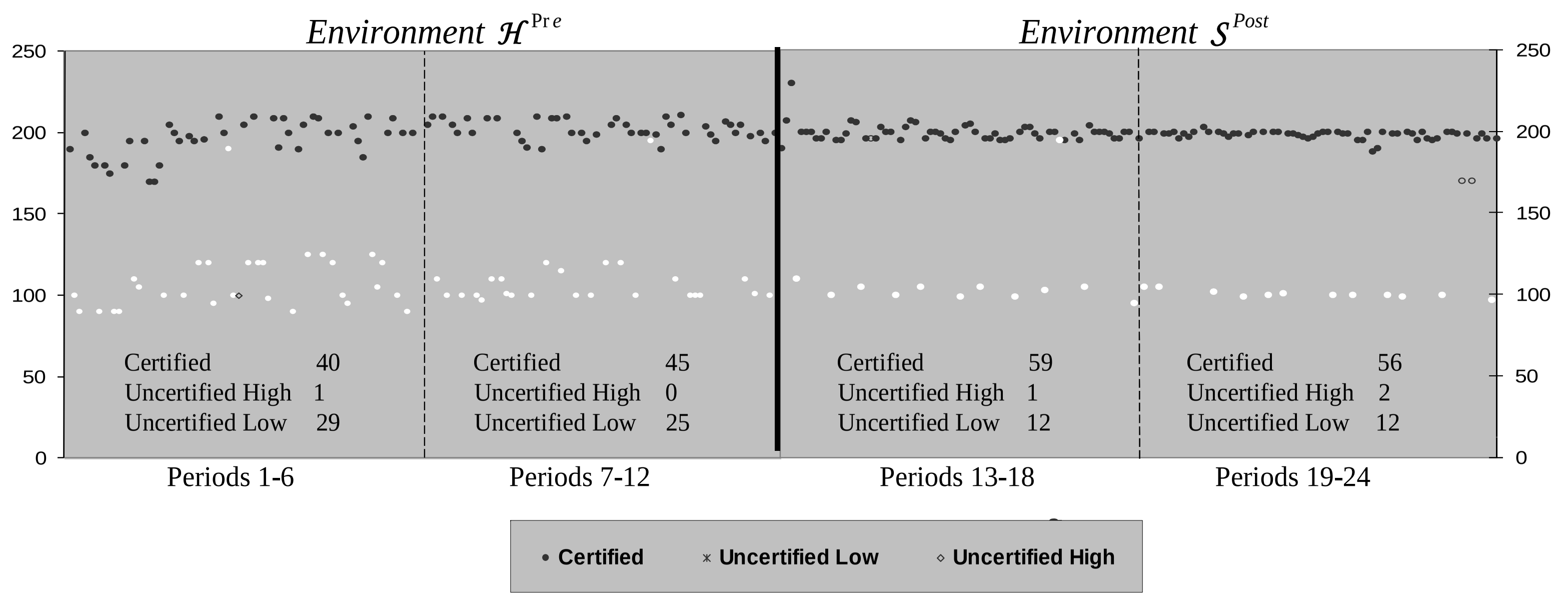

Session 11: Hazardous/Safe Treatment. Public Information

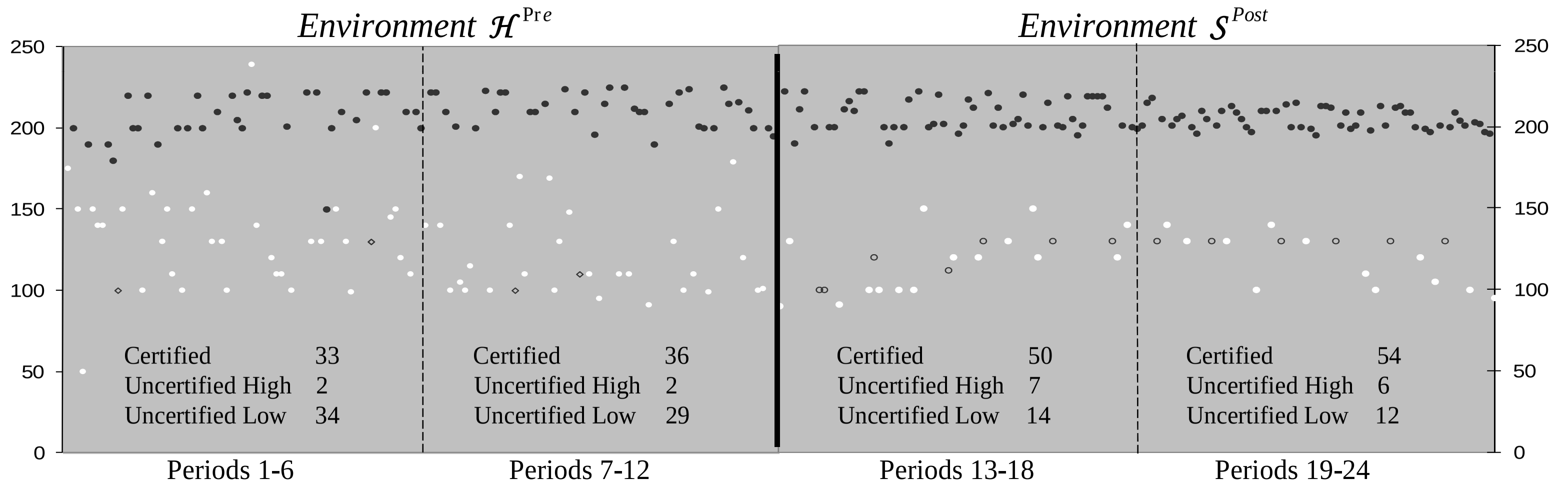

Session 12: Hazardous/Safe Treatment. Public Information

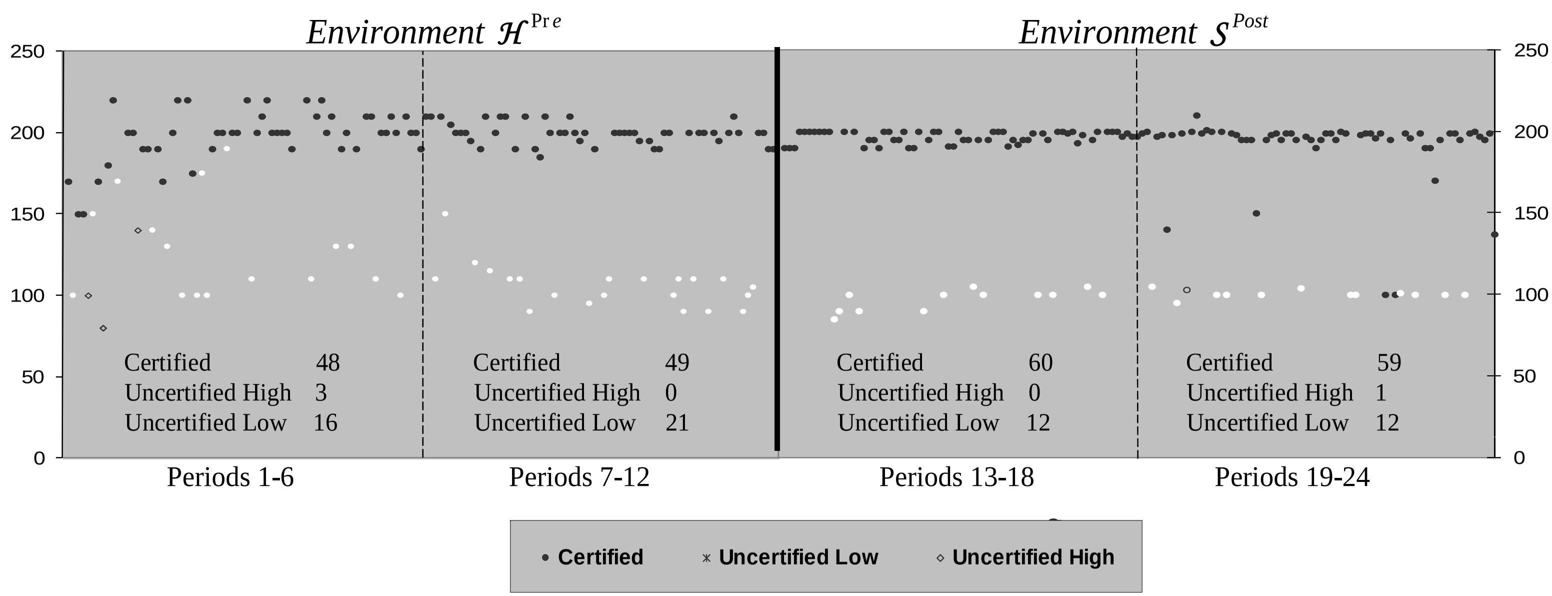

Sellers Instructions

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

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

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

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

30 Points = 1 Swiss Franc

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

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

Overview of the course of the experiment

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

Your Earnings = Price – Production Cost – Certification Fee

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

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

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

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

<u>Your offer to sell:</u>

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

<u>The offers from buyers:</u>

- A buyer's bid to buy consists of the following specifications:
  1) the price he is willing to pay for a unit of the product
  2) the desired quality of the product
  3) whether the buyer requires a certificate or not
- If a buyer requests a certificate you **must** sell the buyer his desired quality. If the buyer doesn't request a certificate you can sell either quality.

2. **The Bonus Phase:** The next phase is the bonus phase. In this phase you have to guess how many of the sellers had lower cost producing the RED quality than producing the BLUE quality during the respective period. If your guess is correct you will earn 20 points.

3. **The Earnings Screen:** At the end of each period you will see the earnings screen. Each participant is informed how much he has earned during the last trading period.

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

Detailed course of the experiment

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

<u>1. The Trading Phase</u>

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

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

During each trading phase you will see the following screen:

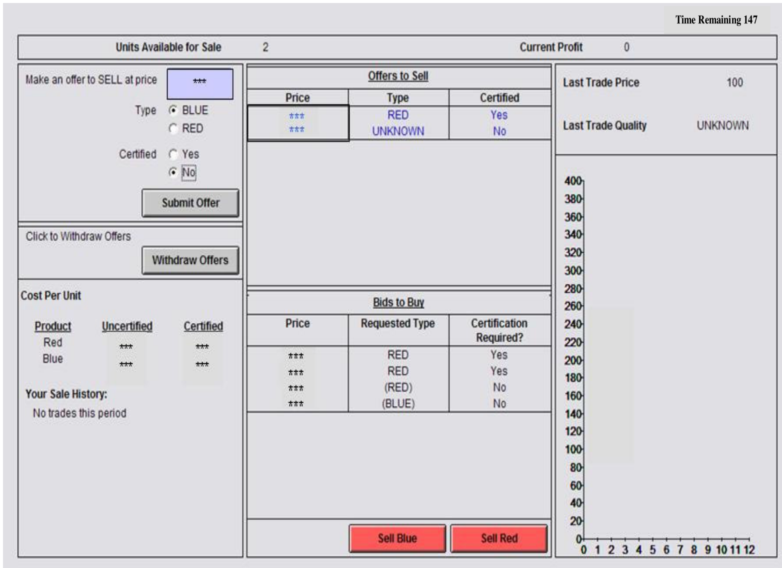

Product Quality

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

Sellers Production Costs

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

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

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

Case 1, RED Quality is cheaper to produce:

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

Case 2a, BLUE Quality is cheaper to produce:

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

Case 2b, BLUE Quality is cheaper to produce

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

Certification

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

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

Your offers to buyers

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

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

$0 \leq$ Price ≤ 400

- You have to specify the quality:

Quality = RED or BLUE

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

Certificate = Yes or No
Costs of certification = 60

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

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

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

As long as you can sell at least one unit you can have two standing offers, one that is certified and one that is not certified. After your second sale all of your standing offers will be deleted.
If you have a standing offer, and you enter a new offer, the new offer replaces the old one, if both offers have the same certification status.

Example:
You have the following standing offers:

Quality	Price	Certified
RED	400	Yes
BLUE	50	No

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

Quality	Price	Certified
RED	**350**	Yes
BLUE	50	No

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

Quality	Price	Certified
RED	350	Yes
RED	**250**	No

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

Accepting offers from buyers

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

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

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

History

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

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

On the right side of the screen you will see the market history. On the top you will find the information of the last traded good. Below you find a chart with all the trades of the period.
On the axis to the right you will find the amount of products traded. On the other axis you will find the price that has been paid for the product. Depending on the quality and certification of the product, the entry is of a different color:

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

2. The bonus phase

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

3. The earning screen

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

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

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

Exercises

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

Please also write down intermediary steps.

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

For these exercises please use the following cost structure:

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

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

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

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

Earnings with certificate =
Earnings without certificate =

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

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

Through which sale can you make the higher earnings?

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