Nº 327
Enero 2008

Documento de Trabajo

ISSN (edición impresa) **0716-7334**
ISSN (edición electrónica) **0717-7593**

Competitive Equilibrium and Reputation Under Imperfect Public Monitoring

Bernardita Vial

www.economia.puc.cl

INDEX

Competitive equilibrium and reputation under imperfect public monitoring

Bernardita Vial*
Pontificia Universidad Católica de Chile

Abstract

In this paper we analyze a reputation-based mechanism that sustains the provision of high quality in a market for an experience good. In contrast to existing models of reputation, however, we consider a competitive market: there is a continuum of firms, each serving at most one consumer each period. We assume a perpetual probability of type replacement and imperfect public monitoring, and we analyze the evolution of firms' reputations in the high quality equilibrium. We find that there is an invariant long run distribution of firms' reputations: each firm's reputation changes every period even in the long run, but the population distribution of reputations remains constant. We consider the long run distribution of firms' reputations to further characterize the steady-state high quality equilibrium. In the equilibrium of the stage game firms with a higher reputation charge a higher price. Furthermore, we show that if the cost of high quality is decreasing in some consumer's characteristic, then buyers pay personalized prices in equilibrium.

When the quality of the good or service is not observable prior to its acquisition, the provision of high quality can only be sustained as an equilibrium outcome if there is a credible and low cost mechanism to enforce agreements. Formal mechanisms include appealing to courts of law or arbitrators, explicit warranties and bonuses.[1] In many cases, however, even after purchasing the good or service it is difficult to verify its quality.[2] In those cases, formal mechanisms become very costly. If high quality is costly and non-contractible, the only possible equilibrium in the absence of repeated interactions is one in which firms sell a low quality good. If, however, firms are long lived and consumers can learn the behavior of each firm in previous periods, the provision of high

*I am very grateful to Felipe Zurita, David K. Levine, Rodrigo Harrison, José Miguel Sánchez and Federico Weinschelbaum for very helpful comments and suggestions.

[1] See MacLeod (2007) for a review of the literature on formal and informal mechanisms to enforce agreements.

[2] Consider, for example, the provision of education by schools; can you verify that the school provided low quality education to your child if he/she fails in a test?

quality may emerge as an equilibrium outcome. In this case, reputation would act as an informal mechanism to enforce contracts.

For the reputation-based mechanism to be able to solve the moral hazard problem there must be a reward to high reputation. Two simple rewards are the possibility of charging higher prices, and the possibility of attracting more clients. Another possible reward is the advantage of avoiding (costly) monitoring, as Diamond (1991) analyzes in the context of the debt market. In any case, the high quality equilibrium is sustained by the fact that reputation is a valuable asset.

In this paper we examine a reputation-based mechanism, where the provision of high quality can be sustained by a reputation premium: firms with higher reputation charge higher prices. In contrast to previous literature, however, we focus on a competitive market. In the context of a competitive market with heterogeneous buyers, three new questions arise relative to the monopolistic case studied in the literature: How does the distribution of firm reputations evolve across time in the repeated game? Which consumers buy from the providers with higher reputation in the stage game? And how large is the reputation premium in equilibrium? This paper addresses those questions, analyzing the evolution of firms' reputations in the repeated game –the "supply side"–, and its interaction with consumers valuations –the "demand side"– to determine the equilibrium assignment of consumers and reputation premiums in the stage game.

The model considers a sequence of generations of short-lived consumers, each of unit measure, and a continuum of firms (long run players). Following Mailath and Samuelson (2001) we assume that there are two possible types of firms: competent firms, which may or may not choose high effort, and inept firms, which always exert low effort. Each producer can serve at most one consumer per period. Therefore, the only decision of a competent firm is the effort to be exerted at each stage. In a "high quality equilibrium", the competent type *actually* chooses high effort. The firm's reputation refers to the consumers' beliefs about the firm type –or the probability consumers assign to the firm choosing high effort in the high quality equilibrium. Consumers have common prior beliefs, and they observe the same public signal: although realized quality is privately observed, there is an indirect, imperfect, public signal of it.[3] Hence, they all assign the same probability to a particular firm being competent. However, as each firm has a different history of public signals, firms also have different reputations in the high quality equilibrium.

In Section 2 we prove that in the high quality equilibrium there is a non degenerate invariant, long run distribution of firms' reputations: each firm's reputation changes every period even in the long run, but the population distribution of reputations remains constant. The existence of a steady-state distribution of reputations guarantees the existence of a steady-state high-quality equilibrium, where the price function and the assignment of consumers to firms with different reputations are the same at any period. This is so because the

[3] We ignore the possibility of (credible) communication among consumers.

equilibrium assignment and the price function depend on the belief distribution (and on the distribution of income, which we assume constant).

We consider the long run distribution of firms' reputations to further characterize the steady-state high-quality equilibrium in Section 3. In the stage game equilibrium, firms with higher reputations charge higher prices, and richer clients purchase from firms with better reputation. Hörner (2002) conjectures that in competitive markets with reputation premiums, prices are determined by consumers' indifference among offers. This paper shows, however, that if consumers differ in their valuations –e.g. due to an income effect–, then equilibrium prices are not determined by consumers' indifference among offers, but by the interaction of reputations' demand and supply. Moreover, even though the existence of capacity constraints prevents the provision of high quality from being sustained by attracting more clients after a good outcome, it can still be sustained by charging higher prices. Therefore, the shape of the equilibrium price function, and its relation to the distribution of firms' reputations, are key elements for the analysis of the high quality equilibrium.

Furthermore, in Section 4 we provide an extension of the model to a market where the cost of providing high quality depends on some consumer characteristic. A salient example of such a market is the school market: the cost of providing high educational achievement and of obtaining good scores in a standardized test –the public signal– is lower for high ability students. We show that if the cost of providing high quality is decreasing in some consumer's characteristic, then consumers with a higher endowment of this characteristic receive price discounts. Therefore, buyers pay personalized prices in equilibrium, as in Mailath, Postlewaite and Samuelson (2006). The equilibrium allocation is characterized by stratification by income and also by stratification by this consumer's characteristic: consumers with a higher income and consumers with a higher endowment of this characteristic purchase from firms with better reputation.

Related Literature

Klein and Leffler (1981) first analyzed a competitive market in which the provision of high quality is sustained by the threat of losing all future trades if the firm provides low quality. The authors emphasized the need for a reward (pricing over cost) to those firms that do provide high quality, so that they could overcome the temptation of a short-run profit.

For Klein and Leffler's mechanism to operate, it is essential that consumers can recognize the name of the firm in order to keep track of its behavior. For this reason, this mechanism is sometimes referred to as a "reputation-based" mechanism. Game-theoretic models of reputation, however, have as an important ingredient the idea that some characteristic of the firm is unobservable to consumers (i.e., there is incomplete information), and that they can learn about this characteristic upon observation of firm's past play. If instead all consumers know the firm's type –as in Klein and Leffler (1981)–, then past play is only informative of past stage-game strategies.

In a game-theoretic reputation model, the provision of high quality can be sustained by the threat of being perceived as incapable of producing high quality (or losing the credibility of being competent) if the firm provides low quality. As in Klein and Leffler (1981), for this mechanism to operate it is essential that consumers keep track of the firm's behavior. This monitoring can be perfect or imperfect: in the first case, consumers perfectly observe past stage-game firm's strategies, while in the latter consumers merely observe a signal imperfectly related to the firm's strategy. If, on the other hand, all consumers observed the same signal, this monitoring would be public, in contrast to the case where different consumers observe different signals (private monitoring). Fudenberg and Levine (1989) provide a lower bound on the payoff to the long run player –the firm, in the product quality example– in the perfect public monitoring case: as the discount factor approaches one, this lower bound converges to the payoff the firm could obtain by providing high quality if quality were observable by buyers. Fudenberg and Levine (1992) extend this result to imperfect monitoring, as long as the history of past signals allows the buyers to statistically identify the firm's strategy.

The game-theoretic reputation literature on reputation developed in the context of a monopolistic market. In the model of Mailath and Samuelson (2001), for example, there is a unique firm characterized by its type (competent or inept) and its reputation (as measured by the probability of being competent), and a continuum of consumers who repeatedly purchase an experience good from the firm. There is imperfect public monitoring, in the sense that if the firm were to choose high effort, consumers would receive a good utility outcome with high probability. Therefore, consumers can learn about firm's type upon observation of the realization of utility outcomes in the high quality equilibrium. In Mailath and Samuelson (1998), in contrast, monitoring is private: each consumer receives a different realization of the utility outcome, and updates beliefs accordingly. Therefore, there is a (non-degenerate) distribution of consumers' posteriors –where the posterior refers to the probability that the unique firm is competent–, since each consumer faces a different history of realizations of the utility outcome.

Hörner (2002) studies a reputation model with imperfect monitoring in a competitive market, and provides an incomplete information model for Klein and Leffler's envisioned mechanism. As Mailath and Samuelson (1998 and 2001) do, he assumes that consumers experience a utility outcome stochastically related to firm effort: the probability of a good outcome increases with effort. Even though all the clients of a given firm experience the same outcome, other consumers do not observe it; hence, it is not public. Moreover, he focuses on nonrevealing high-effort equilibria, where all operating firms charge identical prices in each period. In his model, the incentives to produce high quality come from the fact that customers stop buying from the firm after a bad realization, and the firm goes out of business. Therefore, the only surviving firms in a given period are those whose customers never experienced a bad outcome, and the distribution of firms' reputations is degenerate: all operating firms in each period share the same reputation. As the author points out, however, the market for some services –such as those provided by lawyers, physicians or consultants– are better characterized by a revealing equilibrium, where the prices charged for the service depend on the recent performance of each provider.

An important issue in a model with imperfect public monitoring, is that there would be no room for permanent reputations if the uncertainty about types were not continually replenished, as Cripps, Mailath and Samuelson (2004) proved in a monopolistic context. This occurs because as the firm's reputation converges to one (i.e., consumers become convinced that the firm is competent), a bad signal has no effect on the posterior reputation, and hence the cost of choosing low effort goes to zero. In Hörner's (2002) model, the existence of competition among firms is a sufficient condition to assure that there is always a cost associated to low effort: as all firms charge the same price, if the customers of a given firm experience a bad outcome they would switch to another firm. In contrast to Hörner's (2002) model, we consider the case of a public record in the analysis of competitive markets with reputation. Therefore, a firm that obtained a low signal can still attract customers in subsequent periods by charging lower prices. We follow Mailath and Samuelson (2001) in introducing an exogenous probability of replacements that assures that the uncertainty about types is continually replenished, since there is always a positive probability that a firm changes its type.

We further assume that the probability of being replaced by a competent firm is exogenous. That is, we rule out the possibility of trading "names". Instead, Tadelis (1999) and Mailath and Samuelson (2001) analyze such market for reputations.

1 The Basic Model

1.1 Firms

There is a unit measure of long-lived, profit-maximizing private firms. Each firm can serve at most one consumer. We will refer to the product provided by those firms as a *service*, which could be of low or high quality. Firms are heterogeneous in their types $\tau \in \{C, I\}$ (competent –type C– or inept –type I–, private information) and their prior reputations ($\mu = \Pr(C|\cdot)$, public information). The fraction of competent firms is $\theta \in (0,1)$, constant across time. G_t denotes the cumulative distribution function (henceforth, cdf) of μ at time t, with support $U_t \subset [0,1]$. G_t^C and G_t^I are the cdf of the subpopulations of each private firm type, so that $G_t = \theta G_t^C + (1-\theta) G_t^I$.

We follow Mailath and Samuelson (2001) in introducing an exogenous probability of replacements. Consumers are never completely sure about any firm's type: there is an exogenous probability $\lambda \in (0,1)$ that a firm leaves the market each period and is replaced by another firm. The probability of being replaced by a competent firm is θ, also exogenous, and the new firm inherits the reputation of the old firm it replaces. In other words, λ is the probability that an existing firm changes its owner or management, and θ is the probability that the new manager is competent. Consumers do not observe the replacement, but they know it may occur (and they take this into account when updating beliefs).

Competent firms must choose whether to exert high or low effort. The effort level deterministically leads to a quality level $q \in \{0, 1\}$. For this reason we refer to low/high effort and low/high quality interchangeably. $\alpha_t(\mu)$ denotes the fraction of type C firms with reputation μ that chooses $q = 1$ at time t, and it is referred to as a **production strategy** for it summarizes the choice of production of high and low quality (effort). α_t is the average over the whole population of type C firms. The cost of providing high quality is $c > 0$. Without effort (at zero cost), the quality is 0. Inept firms always provide $q = 0$.[4]

At the end of the period a public signal of the service quality is realized. It so becomes commonly known which firms got a high signal (H) and which ones got a low signal (L). The probability of obtaining a high signal in t depends on the quality provided by the firm, in the following sense:

$$\pi(H|q) = \begin{cases} \pi_1 & \text{if } q = 1 \\ \pi_0 & \text{if } q = 0 \end{cases} \tag{1}$$

where $1 > \pi_1 > \pi_0 > 0$. Therefore, competent firms can increase the probability of obtaining a high signal by exerting sufficient effort in order to provide a high quality service. Consumers beliefs are updated according to the public signal: the posterior belief after a high signal is denoted μ_H, and after a low signal μ_L. We abuse terms and refer to a particular firm with reputation μ (or group of them) as "the firm μ", although the firm's reputation is generally not constant across time, and despite the fact that the set of firms is not countable (atomless economy).[5]

Firms and consumers are price takers. However, in general they do not face a single price for the service, but a price function defined as:

$$\begin{aligned} p_t^* &: U_t \to \mathbb{R} \\ &: \mu \to p_t^*(\mu) \end{aligned}$$

Generally speaking, each firm could make a decision contingent on the past history, including: its own public record of signals h_t, its own reputation μ_t, the history of equilibrium price functions $\{p_k^*(\mu)\}_{k=0}^{t-1}$, the history of distributions of reputations $\{G_k\}_{k=0}^{t}$, the history of equilibrium assignments $\{\mu_k^*(y)\}_{k=0}^{t-1}$ and the history of equilibrium production of quality $\{\alpha_k\}_{k=0}^{t-1}$. We focus, however, on strong Markov strategies where firms make decisions taking as the sole state variable their own reputation μ and type τ (as well as the price function p^*).

[4] They cannot choose higher effort, or the productivity of their effort is zero: they don't know how to improve their quality.

[5] Because there are uncountable many firms and consumers, we should qualify almost every phrase with an "a.e." provision, specifying that what we say is true only in positive-measure sets; this could become annoying, so we avoid every explicit comment of this sort, in the understanding that when we say "all" we mean "almost all".

This is justified by the perfect competition assumption, which rules out any sort of strategic behavior among firms, and the fact that consumers only care about the probability of receiving a high quality service (which is linked to the firm type) and can only base their judgments about firm types on their signals.

We thus restrict attention to perfect Markov strategies σ of the form:

$$\begin{aligned} \sigma^C &: U \to \{0,1\} \\ &: \sigma^C(\mu) = q, \end{aligned} \tag{2}$$

for competent firms.

1.2 Consumers

There is a sequence of (non-overlapping) generations of short-lived consumers, each of unit measure. Consumers are heterogeneous in income $y \in [\underline{y}, \overline{y}] \equiv Y$. F denotes the cdf over y, identical every period. We assume that F is strictly increasing and continuous in its support Y.

Consumers face prices $p^*(\mu)$, and they choose a provider from the set U. They spend what is left of their income $y - p^*(\mu)$ on the consumption of the composite good z. All consumers have the same Bernoulli function $u(q, z)$ defined for $q = \{0, 1\}$ and $z \geq 0$. To shorten notation a bit, write $u_z(\cdot, z)$ for $\frac{\partial u(\cdot,z)}{\partial z}$ and $\Delta u(\cdot, z) = u(1, z) - u(0, z)$. We make the following assumptions about u:

Assumption 1:

u is twice-differenciable in z with a continuous second derivative.

Assumption 2:

u is strictly increasing in z and weakly increasing in q, in the following sense:

$$\begin{aligned} u_z(\cdot, z) &> 0 \text{ for } q = \{0, 1\}; \\ u(1, z) &> u(0, z) \quad \text{ for all } z > 0, \text{ and} \\ u(1, 0) &= u(0, 0). \end{aligned}$$

Consumers do not observe the quality of the service provided by each firm in advance. Since consumers are short run players, they play static best replies, regardless of the history of the game. The only elements of the (public) history that may take into account are those that are judged to be useful in ascertaining firm types if the two types behave differently. Consumers' prior beliefs are common, and since they observe the same history of signals, posterior beliefs are common too.

Hence, a consumer of income y chooses μ (the only firm attribute that matters to him) so as to:

$$\begin{aligned} \max_{\mu \in U} E[u(q, z) | \cdot] &= \Pr(q = 1 | \cdot) u(1, z) + \Pr(q = 0 | \cdot) u(0, z) \\ s/t \quad z &= y - p^*(\mu). \end{aligned} \tag{3}$$

Replacing $\Pr(q=1|\cdot) = \Pr(q=1|\tau=C,\mu)\Pr(\tau=C|\cdot) = \alpha(\mu)\mu$ leaves:

$$\max_{\mu\in U} \mu\alpha(\mu)\,u(1, y-p^*(\mu)) + (1-\mu\alpha(\mu))\,u(0, y-p^*(\mu)). \tag{4}$$

Imposing $\mathrm{d}E[u]=0$, we obtain the marginal rate of substitution between perceived quality (i.e., $\Pr(q=1|\cdot)$) and consumption:

$$MRS = \left|\frac{\mathrm{d}z}{\mathrm{d}\Pr(q=1|\cdot)}\right|_{\mathrm{d}E[u]=0} = \frac{\Delta u(\cdot,z)}{E[u_z(\cdot,z)]}. \tag{5}$$

Assumption 3:
Perceived quality is a normal good, i.e., the MRS is increasing in z:

$$\frac{\partial MRS}{\partial z} = \frac{\Delta u_z(\cdot,z)}{E[u_z(\cdot,z)]} - \frac{(\Delta u(\cdot,z))\,E[u_{zz}(\cdot,z)]}{(E[u_z(\cdot,z)])^2} > 0.$$

Summarizing, the timeline for the stage game at time t is the following:

1. **Initial conditions.** Each firm is endowed with a (public) history of signals h_t and a (private) history of own types, which determines $\mu_t \in U_t \subset [0,1]$, the time-$t$ **prior probability of being competent** (simply called **reputation)**.

2. **Nature's moves**.

   (a) Nature announces privately its type to each firm, drawing from the conditional distributions:

	$\tau_t = C$	$\tau_t = I$
$\Pr(\tau_{t+1}=C\|\tau_t=\cdot)$	$\lambda\theta + (1-\lambda)$	$\lambda\theta$
$\Pr(\tau_{t+1}=I\|\tau_t=\cdot)$	$\lambda(1-\theta)$	$1-\lambda\theta$

   (6)

   The draws are independent within firms that were of the same type in $t-1$.

   (b) At the same time, Nature announces publicly to each (recently born) consumer his/her income, $y \in Y$.

3. **Matching Stage**. A Walrasian market for the service opens.

   (a) Consumers choose firms taking as given the price list $p^*(\mu)$.

   (b) Consumers and firms are correspondingly matched (one consumer per firm); the equilibrium assignment is described by $\mu = \mu^*(y)$.

4. **Production Stage**. Each competent firm chooses whether to spend c to offer a high quality service to its consumer or not. Note that quality is a dichotomous variable, so it wouldn't make sense to spend anything different from c or 0.

5. **Information Stage**. A public signal about the service quality is realized. The (common) beliefs about firm types are the Bayesian update of prior reputations. These posteriors are next period's priors.

These choices result collectively in:

- An **equilibrium price function** p_t^*
- An **equilibrium assignment function** or matching of consumers and firms μ_t^*
- An equilibrium production strategy for each firm $\alpha_t(\mu) \in [0,1]$, the fraction of competent firms with reputation μ that chooses $q = 1$. This fraction determines the **equilibrium aggregate production of high quality services** at t. Because of the continuum assumption, exactly a fraction $\theta\alpha_t$ of consumers receives a high quality service, and a fraction $\pi_1\theta\alpha_t + \pi_0((1-\alpha_t)\theta + (1-\theta))$ of firms gets an H signal at time t.

2 Equilibrium

The equilibrium on a single stage of the game cannot be analyzed in complete isolation. Since firms are long run players, they take into consideration current and future benefits when choosing their effort level at any stage t. For instance, as providing high quality is costly and firm's effort is not observable, the production decision of the firm at date t will be $q = 1$ only if there is a future reward for this effort. The continuation value depends on the way consumers update their beliefs with the public signals, and on the sequence of price lists $\{p_t^*\}$. Therefore, the equilibrium production decision of the firms can only be understood in the repeated game.

In the extremes, we have two types of equilibria:

1. **Low quality equilibrium** (LQE). An equilibrium with all competent firms choosing null effort: all competent firms provide $q = 0$ to their consumers, so that $\alpha_t = 0$ for all t. In this equilibrium, all consumers –by their knowledge of the equilibrium strategies– assign zero probability to the event $\{q = 1\}$. Therefore, firms' reputations are completely irrelevant, for consumers see no advantage in purchasing from a competent firm that offers the same quality than inept firms do. Hence all firms μ must charge the same p, and consequently there is no incentive to provide high quality. This equilibrium would be the only possible outcome of the game if, for instance, the cost of $q = 1$, c, were too high relative to the consumer's willingness to pay for it.
2. **High quality equilibrium** (HQE). An equilibrium where the industry operates at full capacity, that is, all competent firms choose effort so as to provide $q = 1$ to their consumers: $\alpha_t = 1$ for all t. Observe that for this to be true, $q = 1$ must be optimal for competent firms regardless of the current reputation μ_t.

2.1 Equilibrium in the stage game

Consider the stage game at date t, where each firm serves one consumer, and the production decision of competent firms is α_t. Given this production strategy, the equilibrium in the stage game is defined as follows.

Definition 1 *A **Walrasian equilibrium** (WE) for the non-atomic economy* $\left\langle F, G_t^C, G_t^I, \alpha_t \right\rangle$ *of the stage t of this game is a price function* p_t^* *and an assignment function*[6] μ_t^* *such that:*

WE1 *the assignment function* μ_t^* *is optimal for the consumers, given the price function* p_t^* *and the production strategy* α_t*. That is, no consumer with income* y' *wants to purchase from a firm with a reputation different than* $\mu_t^*(y')$*.*

WE2 *the assignment function* μ_t^* *is optimal for the firms, given the price function* p_t^* *and the production strategy* α_t*. That is, no firm with reputation* μ' *wants to serve a consumer such that* $\mu_t^*(y) \neq \mu'$*.*

WE3 *Market clearing. This is to say, given the assignment* μ_t^**, the proportion of consumers that are served by firms with* $\mu \leq \mu'$ *equals* $G_t(\mu')$*.*

If $\langle p_t^*, \mu_t^* \rangle$ is a WE for the production strategy α_t, p_t^* is called a **supporting price function** for μ_t^*.

2.2 Equilibria in the Repeated Game

Definition 2 *A **Markov sequential equilibrium** (MSE) for the repeated game of the economy* $\left\langle F, G_0^C, G_0^I \right\rangle$ *is a sequence of price lists* $\{p_t^*\}$*, a sequence of Markov production strategies* $\{\alpha_t^*\}$*, a sequence of assignment correspondences* $\{\mu_t^*\}$ *and a sequence of reputation distribution pairs* $\left\{\left(G_t^C, G_t^I\right)\right\}$ *such that for all t and* μ *:*

MSE1 Each tuple $\langle p_t^*, \mu_t^* \rangle$ *conforms a WE of the stage game given* α_t^**.*

MSE2 $\alpha_t^*(\mu)$ *is optimal for competent firms with reputation* μ *for all* $\mu \in U_t$*.*

MSE3 Firms' reputation and the population cdfs of reputations $\left(G_t^C, G_t^I\right)$ *evolve in accordance to Bayes' rule and equilibrium strategies, defining the dynamical system* $\left(G_{t+1}^C, G_{t+1}^I\right) = T\left(G_t^C, G_t^I\right)$*.*

When in a MSE the industry works at all times at full capacity (i.e., the equilibrium production strategy is $\alpha_t^*(\mu) = 1$ for all $\mu \in U_t$ and all t) we call such equilibrium a **high-quality Markov sequential equilibrium** (HQ-MSE).

In the next Section, we analyze the dynamics of the distribution of firms' reputations in a HQ-MSE. The main theorem of this Section guarantees the

[6] Strictly speaking, $\mu^*(y)$ could be a correspondence, but in the cases we are interested in, it is not, and it greatly simplifies the presentation to treat it like a function.

existence of a pair of distributions $\left(G^C, G^I\right)$ such that $\left(G^C, G^I\right) = T\left(G^C, G^I\right)$ –provided the existence of a supporting price function p^* is granted.

Observe that if there were such a pair of distributions $\left(G^C, G^I\right)$, and a tuple $\langle p^*, \mu^* \rangle$ that constituted a WE of the economy $\left\langle F, G^C, G^I, \alpha^* \right\rangle$, then the constant sequence $\left\langle \{p^*\}, \{\mu^*\}, \{\alpha^*\}, \left\{G^C\right\}, \left\{G^I\right\} \right\rangle$ would be a MSE of the repeated game for the economy with initial reputations $\left(G^C, G^I\right)$. In this case, we would say that $\langle p^*, \mu^*, \alpha^* \rangle$ is a **steady-state Walrasian equilibrium** (SS-WE) of $\left\langle F, G^C, G^I \right\rangle$.

2.3 Belief Updating in a HQ-MSE

2.3.1 Individual reputations

Consumers form beliefs about a particular firm's current type τ_t in a Bayesian fashion taking into account the information contained in its history h_t, the equilibrium production strategies $\{\alpha_t\}$ and the firm's date-0 reputation μ_0. Notice that prices and assignments do not intervene directly in the evolution of reputations, but only through α_t.

In the high quality equilibrium where the industry always works at full capacity ($\forall t, \quad \alpha_t = 1$), the firm's reputation follows a Markov process, where $\Pr\left(\tau_t = C | h_t, \{\alpha_k = 1\}_{k=0}^{t}, \mu_0\right)$ is defined recursively as follows:

$$\begin{aligned}
\mu_0 &= \Pr(\tau_0 = C) \quad \text{(given as an initial condition), and} \\
\mu_1 &= \Pr(\tau_1 = C | h_1, \alpha_0 = 1) \\
&= \begin{cases} \lambda\theta + (1-\lambda)\frac{\mu_0 \pi_1}{\mu_0 \pi_1 + (1-\mu_0)\pi_0} & \text{if } h_1 = H, \\ \lambda\theta + (1-\lambda)\frac{\mu_0(1-\pi_1)}{\mu_0(1-\pi_1)+(1-\mu_0)(1-\pi_0)} & \text{if } h_1 = L, \end{cases}
\end{aligned} \tag{7}$$

and from then on, is governed by the stochastic difference equation

$$\begin{aligned}
\mu_{t+1} &= \Pr\left(\tau_t = C | h_t, \{\alpha_k = 1\}_{k=0}^{t}\right) \\
&= \begin{cases} \lambda\theta + (1-\lambda)\frac{\mu_t \pi_1}{\mu_t \pi_1 + (1-\mu_t)\pi_0} \equiv \mu_H(\mu_t) & \text{if } h_{t+1} = h_t H, \\ \lambda\theta + (1-\lambda)\frac{\mu_t(1-\pi_1)}{\mu_t(1-\pi_1)+(1-\mu_t)(1-\pi_0)} \equiv \mu_L(\mu_t) & \text{if } h_{t+1} = h_t L. \end{cases}
\end{aligned} \tag{8}$$

where $h_t H$ and $h_t L$ denote a sequence h_t followed by a signal H and L respectively.

$\mu_L(\mu_t)$ and $\mu_H(\mu_t)$ are illustrated in Figure 1. It is apparent in it that both functions are continuous, one-to-one, and have continuous inverse functions.

Observe that the image of each function is strictly smaller than the interval $[0,1]$:

$$\begin{aligned}
\mu_L([0,1]) &= \mu_H([0,1]) \\
&= [\lambda\theta, \lambda\theta + (1-\lambda)] \subset [0,1].
\end{aligned} \tag{9}$$

Moreover, the fact that μ_L lies above the identity for small values of μ means that for too low reputations even after the worst results next period's reputation

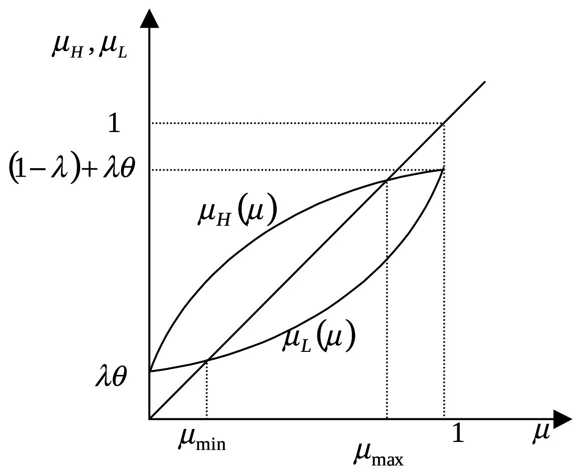

Figure 1: μ_H and μ_L as a function of μ.

would be better, because of the probability of type change. A similar comment applies to very high reputations. Define $\mu_{\min}$ and $\mu_{\max}$ as the values that solve the following equations:

$$\mu_{\min} = \lambda\theta + (1-\lambda)\frac{(1-\pi_1)\,\mu_{\min}}{(1-\pi_1)\,\mu_{\min} + (1-\pi_0)\,(1-\mu_{\min})}, \text{ and} \tag{10}$$

$$\mu_{\max} = \lambda\theta + (1-\lambda)\frac{\pi_1\mu_{\max}}{\pi_1\mu_{\max} + \pi_0\,(1-\mu_{\max})}. \tag{11}$$

Then, $\mu_{\min}$ and $\mu_{\max}$ are bounds for the support of G_t^τ if t is sufficiently large. Indeed, $\lim_{n\to\infty} \Pr\left(\mu_{t+n} < \mu_{\min}\right) = 0$ and $\lim_{n\to\infty} \Pr\left(\mu_{t+n} > \mu_{\max}\right) = 0$. Therefore, if U_t^τ is the support of G_t^τ, then $\lim_{t\to\infty} U_t^\tau \subseteq [\mu_{\min}, \mu_{\max}] \subset [\lambda\theta, \lambda\theta + (1-\lambda)]$.

Observe that in this case, since $\alpha_t(\mu) = 1$ (and unlike any other conceivable equilibrium), choosing $q = 1$ and being of type C are the same event.

2.3.2 Population reputations in a sequence of HQE

Looking at the population as a whole, unlike individual reputations and because of the continuum assumption, the evolution of the cdf of reputations is deterministic. From t to $t+1$ beliefs are updated as the public signals are received, and take into account the fact that a fraction λ of existing firms leave the market and with probability θ are replaced by competent firms.

The pool of competent firms at $t+1$ includes previously existing competent firms, and new competent firms that replaced either old competent or old inept firms. Thus, for competent firms, when $\alpha_t(\mu) = 1$ for all μ, t, we have:

$$\begin{aligned} G^C_{t+1}(x|\mu_t) &= \Pr\left(\mu_{t+1} < x|\mu_t, \tau_{t+1} = C\right) \\ &= (1-\lambda+\lambda\theta)\Pr\left(\mu_{t+1} < x|\mu_t, \tau_t = C\right) \\ &\quad + \frac{\lambda\theta(1-\theta)}{\theta}\Pr\left(\mu_{t+1} < x|\mu_t, \tau_t = I\right). \end{aligned}$$

Given a prior μ_t, the posterior will be less than a given level x if either $\mu_H(\mu_t) < x$ or $\mu_L(\mu_t) < x$, or both. Since H occurs with probability π_1 if $\tau_t = C$ and with probability π_0 if $\tau_t = I$ in the high quality equilibrium, we obtain:

$$\begin{aligned} G^C_{t+1}(x|\mu_t) &= (1-\lambda+\lambda\theta)\left[I_{\{\mu_H < x|\mu_t\}}\pi_1 + I_{\{\mu_L < x|\mu_t\}}(1-\pi_1)\right] \\ &\quad + \lambda(1-\theta)\left[I_{\{\mu_H < x|\mu_t\}}\pi_0 + I_{\{\mu_L < x|\mu_t\}}(1-\pi_0)\right], \end{aligned} \tag{12}$$

where $I_{\{\cdot\}}$ is an indicator function, taking the value 1 if the event $\{\cdot\}$ occurs, and 0 otherwise.

By the law of iterated expectations, we obtain the (unconditional) distribution function of reputations for competent firms in $t+1$ as:

$$\begin{aligned} G^C_{t+1}(x) &\equiv \Pr\left(\mu_{t+1} < x\right) \\ &= (1-\lambda+\lambda\theta)\left[\pi_1\int_0^{\mu_H^{-1}(x)} \mathrm{d}G^C_t + (1-\pi_1)\int_0^{\mu_L^{-1}(x)} \mathrm{d}G^C_t\right] \\ &\quad + \lambda(1-\theta)\left[\pi_0\int_0^{\mu_H^{-1}(x)} \mathrm{d}G^I_t + (1-\pi_0)\int_0^{\mu_L^{-1}(x)} \mathrm{d}G^I_t d\mu\right] \\ &= (1-\lambda+\lambda\theta)\left[\pi_1 G^C_t\left(\mu_H^{-1}(x)\right) + (1-\pi_1)G^C_t\left(\mu_L^{-1}(x)\right)\right] \\ &\quad + \lambda(1-\theta)\left[\pi_0 G^I_t\left(\mu_H^{-1}(x)\right) + (1-\pi_0)G^I_t\left(\mu_L^{-1}(x)\right)\right]. \end{aligned} \tag{13}$$

A similar procedure yields the cdf for inept firms:

$$\begin{aligned} G^I_{t+1}(x) &\equiv \Pr\left(\mu_{t+1} < x\right) \\ &= (1-\lambda\theta)\left[\pi_0 G^I_t\left(\mu_H^{-1}(x)\right) + (1-\pi_0)G^I_t\left(\mu_L^{-1}(x)\right)\right] \\ &\quad + \lambda\theta\left[\pi_1 G^C_t\left(\mu_H^{-1}(x)\right) + (1-\pi_1)G^C_t\left(\mu_L^{-1}(x)\right)\right]. \end{aligned} \tag{14}$$

The overall cdf of μ in $t+1$ is then:

$$\begin{aligned} G_{t+1}(x) &= \theta G^C_{t+1}(x) + (1-\theta)G^I_{t+1}(x) \\ &= \theta\left[\pi_1 G^C_t\left(\mu_H^{-1}(x)\right) + (1-\pi_1)G^C_t\left(\mu_L^{-1}(x)\right)\right] \\ &\quad + (1-\theta)\left[\pi_0 G^I_t\left(\mu_H^{-1}(x)\right) + (1-\pi_0)G^I_t\left(\mu_L^{-1}(x)\right)\right]. \end{aligned} \tag{15}$$

Let X^τ be the set of distribution functions G^τ with support contained in the interval $[\lambda\theta, \lambda\theta + (1-\lambda)]$, where $\tau \in \{C, I\}$ denotes the firm type. We define the distance between two distribution functions G^τ and H^τ with the sup metric, that is:

$$\rho_\infty \left(G^\tau, H^\tau\right) = \sup_{x\in[\lambda\theta,\lambda\theta+(1-\lambda)]} \left|G^\tau(x) - H^\tau(x)\right|. \tag{16}$$

$X^C \times X^I$ is the set of all pairs of distribution functions $\left(G^C, G^I\right)$ with support contained in the interval $[\lambda\theta, \lambda\theta + (1-\lambda)]$. We define the distance between two elements of $X^C \times X^I$ as follows:

$$\rho\left(\left(G^C, G^I\right), \left(H^C, H^I\right)\right) = \max\left\{\rho_\infty\left(G^C, H^C\right), \rho_\infty\left(G^I, H^I\right)\right\}. \tag{17}$$

Therefore, in a sequence of HQE for a given pair of distributions $\left(G_t^C, G_t^I\right)$, next period's distributions $\left(G_{t+1}^C, G_{t+1}^I\right)$ are given by

$$T\left(G_t^C, G_t^I\right) = \left(G_{t+1}^C, G_{t+1}^I\right), \tag{18}$$

with G_{t+1}^C and G_{t+1}^I defined as in Equations (13) and (14) respectively, and where $T : X^C \times X^I \to X^C \times X^I$.

Hence, the operator T defines a first-order difference equation on the product of two cdf spaces. Notice that T is constructed under the assumption that $\alpha_t(\mu_t) = 1$ at all times, so it defines the dynamics of population reputations in an infinite sequence of HQE.

2.4 Long run distribution of firms' reputations

We know from Banach's fixed point theorem, which we reproduce below for convenience, that if an operator in a complete metric space is a contraction, it has a unique fixed point in this space. Unfortunately, the operator T defined in (18) fails to be a contraction. We will prove that Banach's fixed point theorem can be extended to operators that satisfy a weaker condition –a condition for an $N-$contraction, as we define it below. Furthermore, we will prove that T satisfies this condition and that $X^C \times X^I$ is a complete metric space. This allows us to conclude that the operator T has a unique fixed point in $X^C \times X^I$.

The significance of this result is twofold. On the one hand, it implies that μ has a unique and globally stable ergodic distribution function G, which we refer to as the "long run distribution of firms' reputations" or the "long run distribution of beliefs".[7]

On the other hand, the existence of a steady-state distribution of reputations guarantees the existence of a steady-state HQE, where the price function and assignment are the same at any t. This is so because the equilibrium assignment of consumers to firms with different reputations and the price function in the HQE depend on the distribution of beliefs (and the distribution of income, which we assume constant). We characterize this price function and assignment in Section 3.

[7] Similarly, we will refer to the distribution of reputations for competent firms as to the "long run distribution of reputations for competent firms", and likewise for inept firms.

2.4.1 Fixed Point Theorem

Definition 3 *An operator T on a metric space (X,ρ) is said to be an $N-$contraction if*

$$\rho(Tx,Ty) \leq \rho(x,y)$$

for all $x,y \in X$ and there exists a finite $N \in \mathbb{N}$ such that its N-th iterate is a contraction, i.e., if $f : X \to X$ as defined by:

$$fx = T^N x$$

is a contraction in X. That is, $\rho(fx,fy) \leq \beta\rho(x,y)$ for all $x,y \in X$, with $\beta \in (0,1)$.

Hence, the definition of an $N-$contraction is more general than the definition of a contraction, in the sense that a contraction is also a $1-$contraction.

Banach's Fixed Point Theorem *Let (X,ρ) be a complete metric space and $f : X \to X$ a contraction. Then the sequence $\{x_n\}_{n=0}^{\infty}$ defined according to $x_n = f^n x$ is Cauchy, and f has a unique fixed point in X, where f^n is the nth iterate of the operator. Furthermore, if x_0 is any point in X, then the sequence $\{x_n\}_{n=0}^{\infty}$ converges to the fixed point.*[8]

The following Theorem generalizes Banach's fixed point theorem to $N-$contractions:

Proposition 1 (generalization of Banach's Fixed Point Theorem) *Let (X,ρ) be a complete metric space and $T : X \to X$ an $N-$contraction. Then the sequence $\{x_n\}_{n=0}^{\infty}$ defined according to $x_n = T^n x$ is Cauchy, and T has a unique fixed point in X. Furthermore, if x_0 is any point in X, then the sequence $\{x_n\}_{n=0}^{\infty}$ converges to the fixed point.*

Proof. The formal proof is in Appendix A; we sketch here the main ideas. If we define two different sequences $\{x_n\}_{n=0}^{\infty}$ and $\{y_n\}_{n=0}^{\infty}$ according to $x_n = T^n x$ and $y_n = T^n y$ respectively, with $x,y \in X$, we know that after every N periods the distance between x_n and y_n is reduced. Therefore, taking n sufficiently large the distance between the elements of the sequences can be made arbitrarily small. We use this result to prove that any sequence $\{x_n\}_{n=0}^{\infty}$ defined according to $x_n = T^n x$ is Cauchy. Since (X,ρ) is a complete metric space, the sequence converges to a unique point in X, and this point is in fact a fixed point of T. ■

The following three lemmas establish that the operator T defined in Equation (18) is an $N-$contraction.

Lemma 1 *The operator T defined in Equation (18) satisfies:*

$$\rho\left(T\left(G_t^C,G_t^I\right),T\left(H_t^C,H_t^I\right)\right) \leq \rho\left(\left(G_t^C,G_t^I\right),\left(H_t^C,H_t^I\right)\right),$$

for any $\left(G_t^C,G_t^I\right),\left(H_t^C,H_t^I\right) \in X^C \times X^I$, with the metric defined as in Equation (17).

[8] See for example Burkill and Burkill (1970, pp. 52).

Proof. The proof is in Appendix B, but the idea is at follows. $G_{t+1}^{\tau}(x)$ is a weighted average of G_t^C and G_t^I, both evaluated at two different values: $\mu_H^{-1}(x)$ and $\mu_L^{-1}(x)$ -and the same is true for H_{t+1}^{τ}. Therefore, $\left(G_{t+1}^{\tau}(x) - H_{t+1}^{\tau}(x)\right)$ is a weighted average of the difference between G_t^C and H_t^C and the difference between G_t^I and H_t^I, both evaluated at $\mu_H^{-1}(x)$ and $\mu_L^{-1}(x)$. But the distance between G_{t+1}^{τ} and H_{t+1}^{τ} is the supremum in $[\lambda\theta, \lambda\theta + (1-\lambda)]$ of this magnitude. If each addend in the average is maximized at a different level, then the supremum of the average is lower than the average of the supremum of each addend, and therefore the distance between G_{t+1}^{τ} and H_{t+1}^{τ} is lower than the distance between G_t^{τ} and H_t^{τ}. Otherwise, the distance between G_{t+1}^{τ} and H_{t+1}^{τ} may be the same as (but never higher than) the distance between G_t^{τ} and H_t^{τ}. ■

A necessary condition for the distance between G_{t+1}^{τ} and H_{t+1}^{τ} to be a weighted average of $\rho_{\infty}\left(G_t^C, H_t^C\right)$ and $\rho_{\infty}\left(G_t^I, H_t^I\right)$ is the following:

$$\begin{aligned}
&\arg\max_{\{x\}} \left|G_t^C\left(\mu_H^{-1}(x)\right) - H_t^C\left(\mu_H^{-1}(x)\right)\right| \\
&\qquad \cap \arg\max_{\{x\}} \left|G_t^C\left(\mu_L^{-1}(x)\right) - H_t^C\left(\mu_L^{-1}(x)\right)\right| \\
&\qquad \cap \arg\max_{\{x\}} \left|G_t^I\left(\mu_H^{-1}(x)\right) - H_t^I\left(\mu_H^{-1}(x)\right)\right| \\
&\qquad\qquad \cap \arg\max_{\{x\}} \left|G_t^I\left(\mu_L^{-1}(x)\right) - H_t^I\left(\mu_L^{-1}(x)\right)\right| \neq \phi, \quad (19)
\end{aligned}$$

As the support of G_t^{τ} and H_t^{τ} is contained in the interval $[\lambda\theta, \lambda\theta + (1-\lambda)]$, we obtain:

$$\begin{aligned}
G_t^{\tau}\left(\mu_H^{-1}(x)\right) &= H_t^{\tau}\left(\mu_H^{-1}(x)\right) = 0 \quad \text{for } x < \mu_H(\lambda\theta), \\
G_t^{\tau}\left(\mu_L^{-1}(x)\right) &= H_t^{\tau}\left(\mu_L^{-1}(x)\right) = 0 \quad \text{for } x < \mu_L(\lambda\theta), \\
G_t^{\tau}\left(\mu_H^{-1}(x)\right) &= H_t^{\tau}\left(\mu_H^{-1}(x)\right) = 1 \quad \text{for } x > \mu_H(\lambda\theta + (1-\lambda)) \text{ and} \\
G_t^{\tau}\left(\mu_L^{-1}(x)\right) &= H_t^{\tau}\left(\mu_L^{-1}(x)\right) = 1 \quad \text{for } x > \mu_L(\lambda\theta + (1-\lambda)).
\end{aligned}$$

Therefore, if $\rho\left(\left(G_t^C, G_t^I\right), \left(H_t^C, H_t^I\right)\right) > 0$, we obtain:

$$\arg\max_{\{x\}} \left|G_t^{\tau}\left(\mu_H^{-1}(x)\right) - H_t^{\tau}\left(\mu_H^{-1}(x)\right)\right| \subseteq \left[\mu_H(\lambda\theta), \mu_H(\lambda\theta + (1-\lambda))\right] \text{ and} \quad (20)$$

$$\arg\max_{\{x\}} \left|G_t^{\tau}\left(\mu_L^{-1}(x)\right) - H_t^{\tau}\left(\mu_L^{-1}(x)\right)\right| \subseteq \left[\mu_L(\lambda\theta), \mu_L(\lambda\theta + (1-\lambda))\right], \quad (21)$$

for some $\tau \in \{C, I\}$. It follows that Condition (19) cannot hold if $\mu_L(\lambda\theta + (1-\lambda)) < \mu_H(\lambda\theta)$. In this case, the distance between G_{t+1}^{τ} and H_{t+1}^{τ} is strictly lower than a weighted average of $\rho_{\infty}\left(G_t^C, H_t^C\right)$ and $\rho_{\infty}\left(G_t^I, H_t^I\right)$ for both $\tau \in \{C, I\}$. Therefore,

$$\begin{aligned}
&\rho\left(T\left(G_t^C, G_t^I\right), T\left(H_t^C, H_t^I\right)\right) < \rho\left(\left(G_t^C, G_t^I\right), \left(H_t^C, H_t^I\right)\right) \\
&\qquad = \max\left\{\rho_{\infty}\left(G_t^C, H_t^C\right), \rho_{\infty}\left(G_t^I, H_t^I\right)\right\} \\
&\qquad\qquad \forall\ \left(G_t^C, G_t^I\right), \left(H_t^C, H_t^I\right) \in X^C \times X^I, \quad (22)
\end{aligned}$$

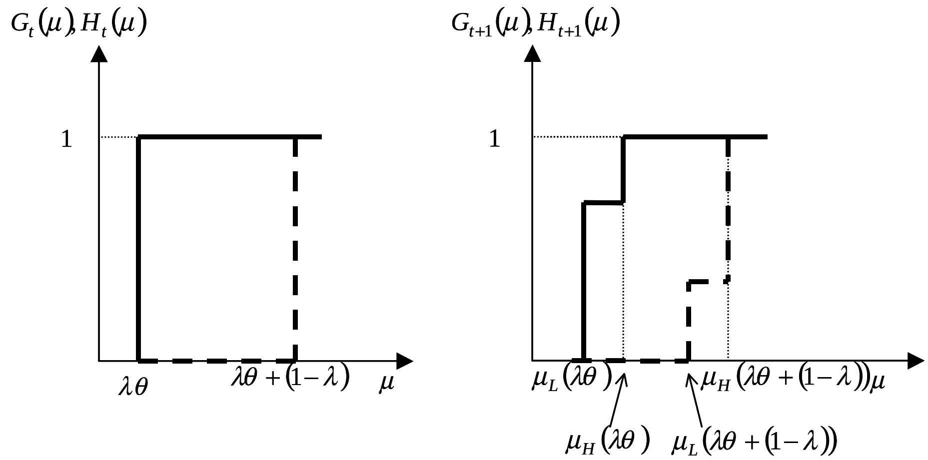

Figure 2: Example where $\rho\left(\left(G_{t+1}^C, G_{t+1}^I\right), T\left(H_{t+1}^C, H_{t+1}^I\right)\right) = \rho\left(\left(G_t^C, G_t^I\right), \left(H_t^C, H_t^I\right)\right)$.

and we can say that T is a contraction mapping.

Now consider the case where $\mu_L(\lambda\theta + (1-\lambda)) > \mu_H(\lambda\theta)$. If we take $G_t^C = G_t^I$ degenerate at $\lambda\theta$ and $H_t^C = H_t^I$ degenerate at $\lambda\theta + (1-\lambda)$, then Condition (19) holds, and $\rho\left(\left(G_{t+1}^C, G_{t+1}^I\right), \left(H_{t+1}^C, H_{t+1}^I\right)\right) = \rho\left(\left(G_t^C, G_t^I\right), \left(H_t^C, H_t^I\right)\right) = 1$. This case is illustrated in Figure 2. Therefore, if $\mu_L(\lambda\theta) > \mu_H(\lambda\theta + (1-\lambda))$, T is not a contraction.

Iterating the operator T and using Lemma 1, we obtain:

$$\rho\left(T^n\left(G_t^C, G_t^I\right), T^n\left(H_t^C, H_t^I\right)\right) \leq \rho\left(\left(G_t^C, G_t^I\right), \left(H_t^C, H_t^I\right)\right). \tag{23}$$

A necessary condition to obtain (23) with equality is that:

$$\begin{aligned} &\arg\max_{\{x\}} \left| G_t^\tau\left(\mu_H^{-1} \ldots \left(\mu_H^{-1}(x)\right)\right) - H_t^\tau\left(\mu_H^{-1} \ldots \left(\mu_H^{-1}(x)\right)\right)\right| \cap \\ &\qquad \arg\max_{\{x\}} \left| G_t^\tau\left(\mu_L^{-1} \ldots \left(\mu_L^{-1}(x)\right)\right) - H_t^\tau\left(\mu_L^{-1} \ldots \left(\mu_L^{-1}(x)\right)\right)\right| \neq \phi, \end{aligned} \tag{24}$$

for $\tau \in \{C, I\}$.

We denote by $\mu_H^n(\lambda\theta)$ the posterior belief after n periods with H starting from $\lambda\theta$, and by $\mu_L^n(\lambda\theta + (1-\lambda))$ the posterior belief after n periods with L

starting from $(\lambda\theta + (1-\lambda))$. Since there are no firms with $\mu < \lambda\theta$, we obtain:[9]

$$\begin{aligned}
G_t^{\tau}\left(\mu_H^{-1}...\left(\mu_H^{-1}(x)\right)\right) &= H_t^{\tau}\left(\mu_H^{-1}...\left(\mu_H^{-1}(x)\right)\right) = 0 \quad \text{for } x < \mu_H^n(\lambda\theta), \\
G_t^{\tau}\left(\mu_L^{-1}...\left(\mu_L^{-1}(x)\right)\right) &= H_t^{\tau}\left(\mu_L^{-1}...\left(\mu_L^{-1}(x)\right)\right) = 0 \quad \text{for } x < \mu_L^n(\lambda\theta), \\
G_t^{\tau}\left(\mu_H^{-1}...\left(\mu_H^{-1}(x)\right)\right) &= H_t^{\tau}\left(\mu_H^{-1}...\left(\mu_H^{-1}(x)\right)\right) = 1 \quad \text{for } x > \mu_H^n(\lambda\theta + (1-\lambda)) \text{ and} \\
G_t^{\tau}\left(\mu_L^{-1}...\left(\mu_L^{-1}(x)\right)\right) &= H_t^{\tau}\left(\mu_L^{-1}...\left(\mu_L^{-1}(x)\right)\right) = 1 \quad \text{for } x > \mu_L^n(\lambda\theta + (1-\lambda)).
\end{aligned}$$

Hence, if $\rho\left(\left(G_t^C, G_t^I\right), \left(H_t^C, H_t^I\right)\right) > 0$, then:

$$\arg\max_{\{x\}} \left|G_t^{\tau}\left(\mu_H^{-1}...\left(\mu_H^{-1}(x)\right)\right) - H_t^{\tau}\left(\mu_H^{-1}...\left(\mu_H^{-1}(x)\right)\right)\right| \subseteq \left[\mu_H^n(\lambda\theta), \mu_H^n(\lambda\theta + (1-\lambda))\right] \text{ and} \tag{25}$$

$$\arg\max_{\{x\}} \left|G_t^{\tau}\left(\mu_L^{-1}...\left(\mu_L^{-1}(x)\right)\right) - H_t^{\tau}\left(\mu_L^{-1}...\left(\mu_L^{-1}(x)\right)\right)\right| \subseteq \left[\mu_L^n(\lambda\theta), \mu_L^n(\lambda\theta + (1-\lambda))\right], \tag{26}$$

for some $\tau \in \{C, I\}$. It follows that Condition (24) cannot hold if $\mu_L^n(\lambda\theta + (1-\lambda)) < \mu_H^n(\lambda\theta)$.

In the following Lemma we prove that, even if $\mu_L(\lambda\theta + (1-\lambda)) > \mu_H(\lambda\theta)$, there is a finite number of periods N such that $\mu_L^N(\lambda\theta + (1-\lambda)) < \mu_H^N(\lambda\theta)$.

Lemma 2 *There is $N < \infty$ such that $\mu_L^N(\lambda\theta + (1-\lambda)) < \mu_H^N(\lambda\theta)$*

Proof. Since $\mu_H^n(x)$ converges to $\mu_{\max}$ and $\mu_L^n(x)$ converges to $\mu_{\min} < \mu_{\max}$, we know that there is $N < \infty$ such that $\mu_L^N(\lambda\theta + (1-\lambda)) < \mu_H^N(\lambda\theta)$. ■

Now we can prove that the operator T is an $N-$contraction.

Lemma 3 *The operator T defined in Equation (18) is an $N-$contraction, that is, $\exists\ N < \infty$:*

$$\rho\left(T^N\left(G_t^C, G_t^I\right), T^N\left(H_t^C, H_t^I\right)\right) \leq \beta\rho\left(\left(G_t^C, G_t^I\right), \left(H_t^C, H_t^I\right)\right),$$

for any pair $\left(G_t^C, G_t^I\right), \left(H_t^C, H_t^I\right) \in X^C \times X^I$, with $\beta \in (0,1)$ and the metric defined in Equation (17).

Proof. From Lemma 2 we know that there is $N < \infty$ such that $\mu_L^N(\lambda\theta + (1-\lambda)) < \mu_H^N(\lambda\theta)$. Since Condition (24) cannot hold if $\mu_L^N(\lambda\theta + (1-\lambda)) < \mu_H^N(\lambda\theta)$, we know that:

$$\rho\left(T^N\left(G_t^C, G_t^I\right), T^N\left(H_t^C, H_t^I\right)\right) < \rho\left(\left(G_t^C, G_t^I\right), \left(H_t^C, H_t^I\right)\right)$$

for any pair $\left(G_t^C, G_t^I\right), \left(H_t^C, H_t^I\right) \in X^C \times X^I$. Furthermore, there is $\beta \in (0,1)$ such that:

$$\rho\left(T^N\left(G_t^C, G_t^I\right), T^N\left(H_t^C, H_t^I\right)\right) \leq \beta\rho\left(\left(G_t^C, G_t^I\right), \left(H_t^C, H_t^I\right)\right).$$

Therefore, the operator T is an $N-$contraction. ■

[9] Notice that if $x < \mu_H^n(\lambda\theta)$, we know that $\mu_H^{-1}...\left(\mu_H^{-1}(x)\right) < \lambda\theta$. Similarly, if $x > \mu_L^n(\lambda\theta + (1-\lambda))$, we know that $\mu_L^{-1}...\left(\mu_L^{-1}(x)\right) > \lambda\theta + (1-\lambda)$.

Proposition 2 *The set of distribution functions G with support contained in the interval $[\lambda\theta, \lambda\theta + (1-\lambda)]$ endowed with the sup metric, (X, ρ_∞), is a complete metric space. Therefore, $\left(X^C \times X^I, \rho\right)$ is complete.*

See Appendix C for proof. As $\left(X^C \times X^I, \rho\right)$ is a complete metric space, we can apply the result established in Proposition 1, and we obtain the following theorem.

Theorem 1 *The operator T has a unique fixed point in $X^C \times X^I$, that can be reached beginning from any pair of distribution functions $\left(G^C, G^I\right)$ in $X^C \times X^I$.*

2.4.2 Characterization of the long run distribution of beliefs

Since $\left(G^C, G^I\right)$ is a fixed point of the operator T, we know that if the distribution of reputations in t is the long run distribution, G, then $G_{t+1} = G_t = G$. We use this fact to prove that G has the following property:

Lemma 4 *The support of the long run distribution of reputations, U, is contained in $[\mu_{\min}, \mu_{\max}]$. Furthermore, $\mu_{\min}$ and $\mu_{\max}$ are limit points of U.*

Proof. Consider a distribution G_t with support U_t, such that $U_t \nsubseteq [\mu_{\min}, \mu_{\max}]$. Then, either $\inf U_t < \mu_{\min}$ or $\sup U_t > \mu_{\max}$ (or both). But if $\mu_t < \mu_{\min}$, then $\mu_{t+1} > \mu_t$, and if $\mu_t > \mu_{\max}$, then $\mu_{t+1} < \mu_t$. Therefore, either $\inf U_{t+1} > \inf U_t$ or $\sup U_{t+1} < \sup U_t$ (or both), hence $U_{t+1} \neq U_t$, and G_t is not the long run distribution of reputations G. We conclude that $U \subseteq [\mu_{\min}, \mu_{\max}]$.

Now consider a distribution G_t with support U_t, such that $\mu_{\min}$ is not a limit point of U_t. Let define $U_t' = U_t \backslash \{\mu_{\min}\}$. As $\mu_{\min} < \mu_L(\mu_t) < \mu_t$ for all $\mu_t > \mu_{\min}$, then $\inf U_{t+1}' < \inf U_t'$, and G_t is not the long run distribution of reputations G. We conclude that $\mu_{\min}$ is a limit point of U. The same reasoning applies to $\mu_{\max}$. ■

Lemma 5 *The support of the long run distribution of reputations is infinite. In particular, U is non degenerate.*

Proof. G cannot be degenerate because there is no μ such that $\mu_L(\mu) = \mu_H(\mu)$. To see that it must have an infinite support, consider a non-degenerate distribution G_t with finite support U_t and check that it cannot be the long-run distribution G as follows. Define m as:

$$m = \min\{\mu \in U_t : \mu > \mu_{\min}\}. \tag{27}$$

We know m exists because G_t is non degenerate and U_t is finite. Since $m > \mu_{\min}$, we know that $\mu_{\min} < \mu_L(m) < m$, and therefore G_{t+1} differs from G_t because $\mu_L(m) \in U_{t+1}$ and $\mu_L(m) \notin U_t$. Hence, G_t is not the long run distribution of reputations G. We conclude that U must be infinite. ■

As we pointed out before, firms' types are not even eventually revealed. However, in the long run a good-type firm has a better reputation than a bad-type firm, in the sense that G^C first-order stochastically dominates G^I. We prove this result by exploiting the fact that a pair $\left(G^C, G^I\right)$ of long run distribution functions for types C and I is the limit point of the sequence defined from the iteration of the operator T starting from any pair of initial distribution functions in $X^C \times X^I$.

Lemma 6 *G^C first-order stochastically dominates G^I.*

Proof. From Equations (13) and (14) we obtain:

$$G_{t+1}^I(x) - G_{t+1}^C(x) = (1-\lambda)\left[\pi_0 G_t^I\left(\mu_H^{-1}(x)\right) + (1-\pi_0) G_t^I\left(\mu_L^{-1}(x)\right) - \left(\pi_1 G_t^C\left(\mu_H^{-1}(x)\right) + (1-\pi_1) G_t^C\left(\mu_L^{-1}(x)\right)\right)\right]. \quad (28)$$

But $\mu_H^{-1}(x) < \mu_L^{-1}(x)$ for all $x \in [\lambda\theta, \lambda\theta + (1-\lambda)]$, so $G_t^\tau\left(\mu_H^{-1}(x)\right) \leq G_t^\tau\left(\mu_L^{-1}(x)\right)$ for $\tau \in \{C, I\}$. Hence, if $G_t^I(x) \geq G_t^C(x)$ for all $x \in [\lambda\theta, \lambda\theta + (1-\lambda)]$, we obtain that $G_{t+1}^I(x) \geq G_{t+1}^C(x)$ for all $x \in [\lambda\theta, \lambda\theta + (1-\lambda)]$, since $\pi_1 > \pi_0$. Therefore, starting with the same distribution functions for competent and inept firms, $G_t^C = G_t^I$, we obtain that $G_{t+n}^I(x) \geq G_{t+n}^C(x)$ for all $x \in [\lambda\theta, \lambda\theta + (1-\lambda)]$ and for all $n > 0$. Taking the limit as n goes to infinity, we get:

$$G^I(x) \geq G^C(x) \text{ for all } x \in U.$$

■

Example 1 *Consider the following parameters: $\lambda = 0.1$; $\theta = 0.5$; $\pi_1 = 0.6$; $\pi_0 = 0.4$. Starting from a degenerate distribution of prior beliefs at θ, and with N=10000 firms, the histograms for μ after 1000 periods for firms C and I are illustrated in Figures 3 and 4 respectively.*

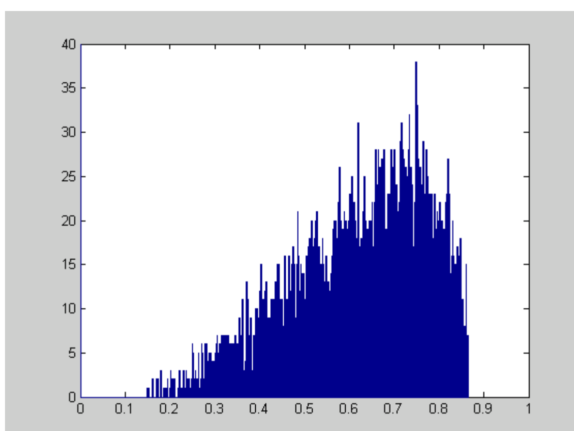

Figure 3: Histogram of μ for competent firms.

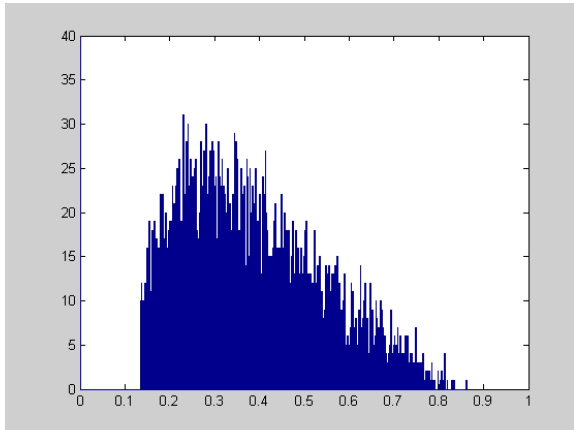

Figure 4: Histogram of μ for inept firms.

3 The Steady-State High-Quality Equilibrium

This Section analyzes a steady-state Walrasian equilibrium (SS-WE) of the economy $\langle F, G^C, G^I \rangle$ in which the constant sequence $\langle \{p^*\}, \{\mu^*\}, \{\alpha^*\}, \{G^C\}, \{G^I\} \rangle$ is a HQ-MSE such that $\alpha^* = 1$ for all $\mu \in U$. For short, we will refer to it as a **steady-state high quality equilibrium** (SS-HQE).

The solution to the Bellman equations for competent and inept firms in the SS-HQE will be as follows:

$$v_C(\mu) = p(\mu) - c + \delta(1-\lambda)(\pi_1 v_C(\mu_H(\mu)) + (1-\pi_1) v_C(\mu_L(\mu))), \text{ and} \quad (29)$$

$$v_I(\mu) = p(\mu) + \delta(1-\lambda)(\pi_0 v_I(\mu_H(\mu)) + (1-\pi_0) v_I(\mu_L(\mu))). \quad (30)$$

The next three lemmas follow from incentive compatibility constraints, and establish that there is an upper bound to the cost of effort for the HQE to exist (Lemma 7), that any given reputation is more valuable for competent than for inept firms (Lemma 8), and that reputations are positively priced in equilibrium (Lemma 9).

Lemma 7 *The cost of high effort, c, is bounded above by:*

$$\delta(1-\lambda)(\pi_1 - \pi_0)(v_C(\mu_H) - v_C(\mu_L)) \text{ for all } \mu \in U.$$

Proof. For spending in high achievement to be preferred, it must be the case that

$$p(\mu) - c + \delta(1-\lambda)(\pi_1 v_C(\mu_H) + (1-\pi_1) v_C(\mu_L)) \\ \geq p(\mu) + \delta(1-\lambda)(\pi_0 v_C(\mu_H) + (1-\pi_0) v_C(\mu_L)).$$

Solving for c yields the desired result. ■

Lemma 8 $v_C(\mu) \geq v_I(\mu) \geq 0$ *for all* $\mu \in U$.

Proof. $v_I(\mu)$ is not larger than the payoff that a competent firm could get through the policy of never investing in $q = 1$. It could be smaller, because inept firms cannot choose $q = 1$. $v_I(\mu) \geq 0$ is a participation constraint. ■

Notice that this is not to say that an inept firm will not be willing to incur in a loss at a particular moment, but simply that it will be willing to do so only if the expected future profits make it worthwhile.

Lemma 9 p^* *is not decreasing.*

Proof. Since

$$\frac{\partial v_C(\mu)}{\partial \mu} = \frac{\partial p}{\partial \mu} + \delta(1-\lambda)\left(\pi_1 \frac{\partial v_C(\mu_H)}{\partial \mu_H}\frac{\partial \mu_H}{\partial \mu} + (1-\pi_1)\frac{\partial v_C(\mu_L)}{\partial \mu_L}\frac{\partial \mu_L}{\partial \mu}\right),$$

and $\frac{\partial \mu_H}{\partial \mu}$, $\frac{\partial \mu_L}{\partial \mu}$ are both positive, then:

$$\text{sign}\left(\frac{\partial p}{\partial \mu}\right) = \text{sign}\left(\frac{\partial v_C}{\partial \mu}\right).$$

Finally, if $\frac{\partial v_C}{\partial \mu} < 0$, the expected value would be decreasing in the probability of a high score:

$$(\pi_1 v_C(\mu_H) + (1-\pi_1) v_C(\mu_L)) < (\pi_0 v_C(\mu_H) + (1-\pi_0) v_C(\mu_L)),$$

resulting in a negative gross benefit of investing in $q = 1$. ■

We now turn to the characterization of the equilibrium assignment and price function. We prove that in any SS-HQE there results a unique equilibrium price function p^* after imposing the condition that $p(\mu_{\min}) = \underline{p} < \underline{y}$. Furthermore, this price function is strictly increasing (see Theorem 2 bellow), hence the condition of Lemma 9 is satisfied. If $\underline{p} > 0$ it turns out that the condition of Lemma 8 is also satisfied, and therefore we know that there exist $c > 0$ such that for them SS-HQE exists, as we conclude in Theorem 3.

3.1 Assignment and pricing

Since $\alpha = 1$ in the HQE, the consumer's expected utility becomes:

$$E[u] = \mu u(1,z) + (1-\mu) u(0,z). \tag{31}$$

Totally differentiating (31) with respect to z and μ we get:

$$\mathrm{d}E[u] = (u(1,z) - u(0,z))\,\mathrm{d}\mu + \left(\mu \frac{\partial u(1,z)}{\partial z} + (1-\mu)\frac{\partial u(0,z)}{\partial z}\right)\mathrm{d}z. \tag{32}$$

The consumer's choice is determined by:[10]

$$\max_{\mu \in U} E[u] = (\mu u(1, y - p^*(\mu)) + (1-\mu) u(0, y - p^*(\mu))). \tag{33}$$

[10] In the analysis that follows we assume that U is a convex set; that is, $U = [\mu_{\min}, \mu_{\max}]$. For this to happen, a restriction on the parameters λ, π_0, π_1 and θ must be imposed in order to obtain $\mu_H(\mu_{\min}) < \mu_L(\mu_{\max})$.

This yields a "demand curve" that coincides with the equilibrium assignment $\mu^*(y)$ because it is feasible.

An interior μ^* satisfies the following first order condition (henceforth, FOC):

$$\frac{\partial E[u]}{\partial \mu} = E[u_z(\cdot, z)]\left(-\frac{\partial p^*}{\partial \mu}\right) + \Delta u(\cdot, z) = 0. \tag{34}$$

Therefore, for an interior solution the price function must be increasing in μ:

$$\frac{\partial p^*}{\partial \mu} = \frac{\triangle u(\cdot, z)}{E[u_z(\cdot, z)]} = \frac{\triangle u(\cdot, y - p(\mu))}{E[u_z(\cdot, y - p(\mu))]} > 0, \tag{35}$$

provided that $y > p(\mu)$. Observe that this is a sufficient condition for Lemma 9.

The second order condition (henceforth, SOC) is:

$$\begin{aligned}\frac{\partial^2 E[u]}{\partial \mu^2} = &-2\frac{\partial p(\mu)}{\partial \mu}\Delta u_z(\cdot, z) \\ &+ E[u_{zz}(\cdot, z)]\left(\frac{\partial p(\mu)}{\partial \mu}\right)^2 - E[u_z(\cdot, z)]\frac{\partial^2 p(\mu)}{\partial \mu^2} < 0.\end{aligned} \tag{36}$$

Lemma 10 *In a SS-HQE μ^* is strictly increasing in y under assumption 3. This is to say, there is stratification by income.*

Proof. The stratification result is obtained from comparative statics on the solution to the consumers' problem. Taking the derivative of the FOC with respect to y, we obtain:

$$\frac{\partial \mu}{\partial y} = \frac{\Delta u_z(\cdot, z) - E[u_{zz}(\cdot, z)]\frac{\partial p}{\partial \mu}}{2\frac{\partial p}{\partial \mu}\Delta u_z(\cdot, z) - E[u_{zz}(\cdot, z)]\left(\frac{\partial p}{\partial \mu}\right)^2 + E[u_z(\cdot, z)]\frac{\partial^2 p}{\partial \mu^2}}. \tag{37}$$

Replacing $\frac{\partial p}{\partial \mu}$ with $\frac{\Delta u(\cdot, z)}{E[u_z(\cdot, z)]}$ in the numerator, we get:

$$\frac{\partial \mu}{\partial y} = \frac{\Delta u_z(\cdot, z) - E[u_{zz}(\cdot, z)]\frac{\Delta u(\cdot, z)}{E[u_z(\cdot, z)]}}{2\frac{\partial p}{\partial \mu}\Delta u_z(\cdot, z) - E[u_{zz}(\cdot, z)]\left(\frac{\partial p}{\partial \mu}\right)^2 + E[u_z(\cdot, z)]\frac{\partial^2 p}{\partial \mu^2}}. \tag{38}$$

Under the assumption that perceived quality is a normal good, the numerator in (38) is positive. On the other hand, the denominator of (38) is the negative of the expression in Equation (36). Therefore, $\frac{\partial \mu}{\partial y}$ is positive if and only if the SOC is met. ■

Remark 1 *The consumers' problem has an interior solution if and only if there is stratification by income. Stratification by income is the consequence of the combination of quality being a "normal good", and prices being increasing in μ, so it will be present in any equilibrium with those characteristics. In particular, it has nothing to do with the fact that we are looking at a steady state.*

As the assignment function μ^* is increasing, Condition WE3 implies:

$$G(\mu^*(y)) = F(y),$$

where F denotes the cummulative distribution of income, and G the cummulative distribution of firms' reputations. Therefore, the equilibrium assignment function is uniquely determined as:

$$\mu^*(y) = G^{-1}(F(y)).$$

That is, the allocation of consumers among firms is uniquely determined by the distribution of consumers valuations (the demand side), and of firms reputations (the supply side).

Example 2 *Assume that the distribution of μ is uniform in $\left(\frac{1}{4}, \frac{3}{4}\right)$ and that the distribution of y is uniform in $[5, 10]$. Then,*

$$\begin{aligned} \mu^*(y) &= G^{-1}\left(\tfrac{1}{5}y - 1\right) \\ &= \frac{\left(\frac{1}{5}y - 1\right)}{2} + \frac{1}{4} \\ &= \frac{y}{10} - \frac{1}{4}. \end{aligned}$$

The equilibrium price function $p^*(\mu)$ must satisfy the FOC of the consumers maximization problem at the equilibrium assignment $\mu^*(y)$. Since μ^* is increasing, the inverse function $y^*(\mu) \equiv \mu^{*-1}(\mu)$ is well defined. Thus replacing $\mu^{*-1}(\mu)$ on the FOC we obtain a differential equation that any supporting price function for $\mu^*(y)$ must satisfy:

$$\begin{aligned} \frac{\partial p^*(\mu)}{\partial \mu} &= \frac{\triangle u(\cdot, z)}{E[u_z(\cdot, z)]} \\ &= \frac{\triangle u(\cdot, y^*(\mu) - p^*(\mu))}{E[u_z(\cdot, y^*(\mu) - p^*(\mu))]}. \end{aligned} \tag{39}$$

If we impose a boundary condition $p(\mu_{\min}) = \underline{p} < \underline{y}$, we obtain the following result:

Lemma 11 *Under Assumptions 1 and 2 the solution to (39) with $p(\mu_{\min}) = \underline{p} < \underline{y}$ exits and is unique. Moreover, the solution meets the condition $y^*(\mu) > p^*(\mu)$ for all $\mu \in U$.*[11]

Proof. Since u_z is differenciable in z with a continuous derivative, the conditions for the existence and uniqueness of solutions for ordinary differential equations are met (see De la Fuente (2000)). ■

[11] To see this, notice that $\frac{\partial p^*(\mu)}{\partial \mu} = 0$ if and only if $y^*(\mu) = p^*(\mu)$. As $y^*(\mu)$ is strictly increasing, this implies that $p^*(\mu)$ cannot cut $y^*(\mu)$ from below.

Example 3 *If $u(z,q) = z(q+1)$, from the FOC we obtain:*

$$\frac{\partial p(\mu)}{\partial \mu} = \frac{y - p(\mu)}{\mu + 1}. \tag{40}$$

Evaluating (40) at the equilibrium assignment obtained before yields:

$$\frac{\partial p^*(\mu)}{\partial \mu} = \frac{10\mu + 2.5 - p(\mu)}{\mu + 1}. \tag{41}$$

We can find a particular solution to the differential equation above by imposing an initial condition. As an example, if we impose that $p^*(0) = 0$, we obtain the following price function:

$$p^*(\mu) = \frac{2.5}{\mu + 1} - 2.5 + 5\mu$$

as illustrated in figure 5.

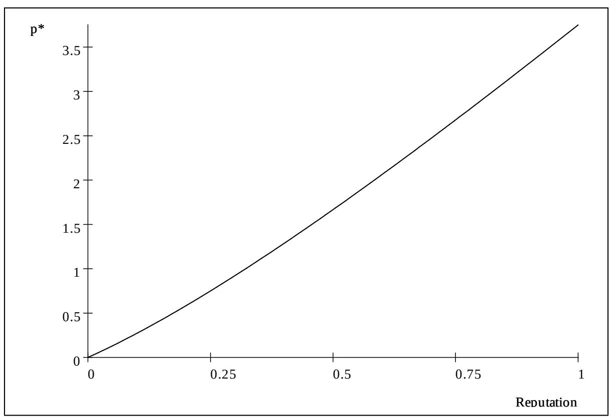

Figure 5: Equilibrium price function $p^*(\mu)$

As $y^*(\mu) > p^*(\mu)$, then $\frac{\partial p^*(\mu)}{\partial \mu} > 0$ and we conclude:

Theorem 2 *In any SS-HQE μ^* and p^* are strictly increasing under Assumptions 1, 2 and 3.*

Hence, under the assumptions stated in Theorem 2 the condition of Lemma 9 is satisfied in any SS-HQE. Moreover, if $\underline{p} > 0$, we also know that Lemma 8 is also satisfied. The question of existence of a SS-HQE bears then on the level of the cost c. The game will not have a SS-HQE if c is too high, i.e., if it does not respect the bound of Lemma 7:

$$c \leq \delta(1-\lambda)(\pi_1 - \pi_0)(v_C(\mu_H(\mu)) - v_C(\mu_L(\mu))), \tag{42}$$

for all $\mu \in U$. Since $(\mu_H(\mu) - \mu_L(\mu)) > 0$ for all $\mu \in (0,1)$, and

$$\lim_{\mu \to 0} (\mu_H(\mu) - \mu_L(\mu)) = 0 \text{ and} \tag{43}$$
$$\lim_{\mu \to 1} (\mu_H(\mu) - \mu_L(\mu)) = 0, \tag{44}$$

the set of possible reputations U –which we already know to be bounded away from 0 and 1– must be *sufficiently* far away from those borders.

Theorem 3 *There exist $c > 0$ such that for them SS-HQE exists.*

Proof. The long run reputation distribution G has a support $U \subseteq [\mu_{\min}, \mu_{\max}]$, where $\mu_{\min}$ and $\mu_{\max}$ are as defined in Equations (10) and (11) respectively. Therefore, $\mu_H(\mu) > \mu_L(\mu)$ for all $\mu \in U$. As $p^*(\mu)$ is increasing in μ, so is $v_C(\mu)$. Hence,

$$(v_C(\mu_H(\mu)) - v_C(\mu_L(\mu))) > 0 \quad \text{for all } \mu \in U.$$

and it follows that there exist $c > 0$ such that for them the condition in Lemma 7 is satisfied. ■

We have just proved that if the cost of providing high quality is sufficiently low, the moral hazard problem can be solved with a reputation-based mechanism. Under this mechanism, firms with a better reputation benefit from it by charging high prices. Therefore, the high quality equilibrium is sustained by a reputation premium. As perceived quality is assumed to be a normal good, those firms sell to richer clients –i.e., clients with a higher willingness to pay. Therefore, the equilibrium assignment shows stratification by income.

4 Personalized cost

In this Section we extend the previous model to the case where consumers are heterogeneous in income $y \in Y$ and some characteristic $b \in B$, where Y and B are closed intervals, and the cost of serving a consumer depends on b. This extension is motivated by the observation that in the market for education, the cost of providing any given educational achievement level is decreasing on the student's ability. In this example, the result in standardized tests can act as imperfect signals of the school quality.[12] Other examples include the health care market, where the cost of providing a successful procedure depends on the condition of the patient. In this example, the public records of physician performance and lawsuits can act as imperfect signals of the physician quality.[13]

We assume that the cost of serving a consumer is decreasing in b. The competent firm that chooses high quality must spend an extra cost $\overline{c}$. Thus, $c(b)$ is the cost of producing a low quality service for a b-customer, and $c(b) + \overline{c}$ is the corresponding high quality cost. As before, the probability of obtaining

[12] For a detailed analysis of this extension of the model, see Vial (2006).
[13] See, for example, Quinn (1998).

a high signal at t depends on the quality provided by the firm. That is, if a competent firm spends $c(b)+\bar{c}$, the probability of obtaining a high signal is $\pi_1 > \pi_0$.

The joint distribution of b and y is denoted by F, and it is assumed continuous in its (convex) support $S \subseteq B \times Y$ and also constant across time. Each consumer observes his own characteristics. All consumers have the same bounded Bernoulli function $u(q,z)$. We maintain Assumptions 1 to 3 about consumers' preferences.

Firms observe consumers' characteristics and face a price function $p^* : U \times B \to \mathbb{R}$, and they choose which consumer to serve. Therefore, each school chooses an admission policy (a set of "admitted" consumers, $P \subset B$). We restrict the sets P to belong to the family $\mathcal{B}(B)$ of compact, convex sets in B (i.e., intervals). That is, admission policies can be summarized by a function defined as:

$$P : U \to \mathcal{B}(B).$$

The timeline of the stage game is the same as before, except for the **matching stage**, which is replaced by:

(a) Firms first choose admission policies taking as given the price list p^*. Admission policies are observable.[14]

(b) Consumers choose firms taking as given the price list p^* and the admission policy of each firm. That is, consumers choose only among those firms that admitted them as clients.

(c) Consumers and firms are correspondingly matched (at most one consumer per firm); the equilibrium assignment is described by $\mu = \mu^*(y,b)$.

As before, we restrict attention to perfect Markov strategies. We define the equilibrium of the stage game for the production strategy α_t as a tuple $\langle p_t^*, \mu_t^* \rangle$ such that $\mu_t^*(b,y)$ is optimal for all $(b,y) \in S$ and for all $\mu \in U$ given the price function p_t^* and the production strategy α_t, and such that the market clears.

The equilibrium of the repeated game is defined in exactly the same way as in the basic model. Since the only information consumers can use to update their beliefs about firm types is the public signal, the evolution of individual reputations and of the distribution of firms reputations is also exactly the same as in the basic model. Therefore, we can analyze the **steady-state high quality equilibrium** (SS-HQE) in which the constant sequence $\langle \{p^*\}, \{\mu^*\}, \{\alpha^*\}, \{G^C\}, \{G^I\} \rangle$ is an equilibrium of the repeated game such that $\alpha^* = 1$ for all $\mu \in U$.

If $\langle p^*, \mu^* \rangle$ is an equilibrium of the stage game, then p^* is called a **supporting price function** for μ^*. Optimality of μ_t^* for firms requires that $\mu_t^*(b,y) \in P^{-1}(b)$ for all $(b,y) \in S$. Moreover, for P to be optimal, the firm with

[14] As the extra cost of providing high quality does not depend on b, we assume that admission policies chosen by firms that provide low and high quality are the same.

reputation μ must be indifferent among all consumers in $P(\mu)$ for all $\mu \in U$. Therefore, the admission policy must satisfy:

$$P(\mu) = \{b \in B : p(\mu, b) - c(b) = k(\mu)\}, \tag{45}$$

where k is any real function that depends on μ but does not depend on b. Therefore, we obtain the following Lemma:

Lemma 12 (Pricing) *Every SS-HQE supporting price function satisfies that for all $\mu \in U$ and for all b such that $\mu = \mu^*(b, y)$:*

$$p^*(\mu, b) - c(b) = k(\mu).$$

Observe that the equilibrium price function p^* defines a price quote for every pair $(\mu, b) \in U \times B$, yet in equilibrium transactions occur only in a strict subset of $U \times B$. There is a great deal of redundancy outside this subset, in the sense that there could be infinitely many supporting price functions for a given assignment function μ^*. That is, there could be many supporting price functions p^* such that the resulting admission policy $P(\mu)$ satisfy that $\mu^*(b, y) \in P^{-1}(b)$ $\forall (b, y) \in B \times Y$. Consider for instance a supporting price function for μ^* such that a firm with reputation μ' does not admit a customer with characteristic b' as a client under p^* (that is, $b' \notin P(\mu')$). Then, any price function p such that:

$$\begin{aligned} p(\mu, b) &< p^*(\mu, b) \text{ if } (\mu, b) = (\mu', b') \text{ and} \\ p(\mu, b) &= p^*(\mu, b) \text{ otherwise} \end{aligned}$$

is also a supporting price function for μ^*.

In what follows we consider the following refinement: $\langle p^*, \mu^* \rangle$ is a refined equilibrium if there is no alternative contract preferred by some consumer and some firm to their equilibrium contract, with strict preference for at least one of them.[15] It can be shown that with this refinement we restrict our attention to equilibria with the property that $p^*(\mu, b)$ is a supporting price function for μ^* under the admission policy of accepting everyone: $P(\mu) = B$ $\forall \mu \in U$. In this case, $\mu^*(b, y)$ is solely determined by the choice of the consumers among all $\mu \in U$.

Therefore, in every refined equilibrium the restricted price function characterized in Lemma 12 can be extended to all $U \times B$ if the admission policies take the simple form "admit everyone", i.e., $P(\mu) = B$ for all $\mu \in U$. We will assume that this extension is the actual equilibrium price function in $U \times B$. In this case, the supporting price function takes the following additively separable form:

$$p^*(\mu, b) = k(\mu) + c(b), \tag{46}$$

where k is any real function that depends on μ but does not depend on b as before. In what follows we will characterize k and μ^*.

[15] In a standard model this is indeed a characterization of the walrasian equilibrium, not a refinement. This is not the case here, because consumers only choose among those firms that admitted them as clients in the first stage.

The solutions to the Bellman equations in the HQE for competent and inept firms are respectively:

$$v_C(\mu) = k(\mu) - \bar{c} + \delta(1-\lambda)(\pi_1 v_C(\mu_H(\mu)) + (1-\pi_1) v_C(\mu_L(\mu))) \text{ and} \tag{47}$$
$$v_I(\mu) = k(\mu) + \delta(1-\lambda)(\pi_0 v_I(\mu_H(\mu)) + (1-\pi_0) v_I(\mu_L(\mu))). \tag{48}$$

As in the basic model, we can characterize the cost and price functions in the SS-HQE from incentive compatibility constraints. The parallel Lemmas are:

Lemma 13 *k is not decreasing.*

Proof. Since

$$\frac{\partial v_C(\mu)}{\partial \mu} = \frac{\partial k}{\partial \mu} + \delta(1-\lambda)\left(\pi_1 \frac{\partial v_C(\mu_H)}{\partial \mu_H}\frac{\partial \mu_H}{\partial \mu} + (1-\pi_1)\frac{\partial v_C(\mu_L)}{\partial \mu_L}\frac{\partial \mu_L}{\partial \mu}\right).$$

and since $\frac{\partial \mu_H}{\partial \mu}$ and $\frac{\partial \mu_L}{\partial \mu}$ are both positive, we obtain:

$$\text{sign}\left(\frac{\partial p}{\partial \mu}\right) = \text{sign}\left(\frac{\partial k}{\partial \mu}\right) = \text{sign}\left(\frac{\partial v_C}{\partial \mu}\right).$$

Finally, if $\frac{\partial v_C}{\partial \mu} < 0$, the expected value would be decreasing in the probability of a high signal, resulting in a negative gross benefit of investing in $q = 1$. ■

Lemma 14 *The cost of high effort, $\bar{c}$, is bounded above by:*

$$\delta(1-\lambda)(\pi_1 - \pi_0)(v_C(\mu_H) - v_C(\mu_L)).$$

Proof. For choosing high quality be optimal for competent schools, it must be the case that for $b \in P^*(\mu)$:

$$\begin{aligned} &k(\mu) - \bar{c} + \delta(1-\lambda)(\pi_1 v_C(\mu_H) + (1-\pi_1) v_C(\mu_L)) \\ &\qquad \geq k(\mu) + \delta(1-\lambda)(\pi_0 v_C(\mu_H) + (1-\pi_0) v_C(\mu_L)). \end{aligned}$$

■

From the consumer's problem we can also characterize the equilibrium assignment and pricing functions.

Theorem 4 *In a SS-HQE $\mu^*(b, y)$ and $k(\mu)$ are strictly increasing under Assumptions 1, 2 and 3. Therefore, there is stratification by income and by b, and there is a positive reputation premium. Moreover, as $p^*(\mu, b) = k(\mu) + c(b)$, consumers pay personalized prices.*

Proof. See Appendix D. ■

As in the basic model, the existence of the HQE bears on the level of $\bar{c}$. Noticing that $U \subseteq [\mu_{\min}, \mu_{\max}]$, we know that $(v_C(\mu_H(\mu)) - v_C(\mu_L(\mu))) > 0$ for all $\mu \in U$, and it follows that:

Theorem 5 *There exist $\bar{c} > 0$ such that for them SS-HQE exists.*

Therefore, in a SS-HQE and under the trivial admission policy $P(\mu) = B$ the pricing function p^* is additively separable, strictly increasing in μ and strictly decreasing in b, for all $\mu \in U$ and $b \in B$. Holding b fixed, consumers with a higher income are served by firms with a higher reputation, and holding income fixed, consumers with a higher b are served by firms with a higher reputation in the SS-HQE.

The fact that more able students attend schools with better results (holding income fixed) is called "stratification by ability" in the economics of education literature. Epple and Romano (1998) obtain stratification by income and stratification by ability as a prediction of a theoretical model with peer effects. Epple, Figlio and Romano (2004) test the central predictions of the Epple and Romano's (1998) model, and they find evidence of stratification by income and ability, and also of discounting to ability in the private sector: higher ability students pay lower tuitions. Therefore, students pay personalized prices in equilibrium, as in Mailath, Postlewaite and Samuelson (2006). Regarding the health care market example, the possibility of charging higher prices for riskier procedures is limited by the fact that insurance companies and other related institutions set maximum payments for each procedure. Therefore, in this case some physicians simply do not admit riskier patients as their clients. Quinn (1998) finds some evidence that supports the hypothesis that reputational factors increase the use of this kind of defensive medicine by physicians.

These are just two examples where the existence of personalized costs engenders an allocation of consumers that depends not only on their income (or preferences, in a more general context), but also on other observable characteristics of the customers. The stratification result is particularly interesting in the educational market, since it determines the distribution of educational achievement –and subsequent income– in the population: the probability of receiving high-quality education is higher for students that are served by schools with a higher reputation. In the health care market example, on the other hand, is the impossibility of charging personalized prices –that is, of adjusting prices according to the patient condition or the riskiness of the procedure demanded– what causes that some physicians are not willing to serve those patients.

5 Concluding Remarks

In this paper we extended game-theoretic reputation models developed in the context of a monopolistic firm to a competitive setting. We considered a market with many long-lived firms that sell an experience good or service, and many short-lived consumers who observe imperfect public signals of the quality provided by each firm in past periods. We analyzed the "high quality equilibrium", where all competent firms provide high quality –while inept firms do not–, and consumers learn about firm types upon observation of the history of public signals.

The concept of "perfect competition" is more subtle in a market with asymmetric information, however. If consumers update their beliefs about the quality of the service provided by each firm according to its history, the firm's reputation can be considered as a (valuable) attribute of the service itself. Therefore, the price-taking condition must accomodate the fact that firms with different reputations charge different prices for the same service, even if the actual quality of the service is the same. That is, the price-taking condition means that consumers and firms take a "price list" –or a price function that depends on the provider's reputation– instead of just one price for the service.

Moreover, the reward to high quality in our model is the price differential: firms with a better reputation charge a higher price in equilibrium. But as reputation can be considered as an attribute of the service provided by each firm, the market clearing condition must take into consideration the supply of this attribute in the economy. Therefore, the shape of the equilibrium price function depends crucially on the distribution of consumers valuations (the demand side, that we assume constant over time), and of firms' reputations (the supply side). The main result of this paper is that even though each firm's reputation changes every period even in the long run, the population distribution of reputations converges to a steady state distribution in the high quality equilibrium. That is, we proved that the distribution of firms' reputations converges to a "long run distribution", the fixed point of the dynamical system that describes the evolution of the distribution of reputations.

The existence of a steady-state distribution of reputations allowed us to focus on the stage game of the steady-state high quality equilibrium, where the price function and the assignment of consumers to firms with different reputations are the same at any period. We showed that firms with a higher reputation charge a higher price in equilibrium, as expected, and consumers with a higher income purchase from firms with better reputation. We also showed that if the cost of providing the service is decreasing on a consumer's characteristic, then consumers with a higher endowment of this characteristic receive price discounts. In this case the equilibrium allocation is also characterized by stratification by this consumer's characteristic: consumers with a higher endowment of this characteristic purchase from firms with better reputation.

References

[1] Burkill, J.C., and H. Burkill, 1970. *A Second Course in Mathematical Analysis.* Cambridge University Press.

[2] Cripps, M., Mailath, G. and L. Samuelson, 2004. "Imperfect Monitoring and Impermanent Reputations". Econometrica, Vol. 72, No. 2 (March), 407–432.

[3] Diamond, D. W., 1991. "Monitoring and Reputation: The Choice between Bank Loans and Directly Placed Debt". The Journal of Political Economy, Vol. 99, No. 4., 689-721.

[4] Epple, D. and R. Romano, 1998. “Competition between Private and Public schools, Vouchers and Peer Group Effects”, The American Economic Review, Vol. 88, No.1 (Mar.), 32-62.

[5] Epple, D. and R. Romano, 2002. “Educational Vouchers and Cream Skimming”, NBER Working Paper No.9354.

[6] Epple, D., D. Figlio and R. Romano, 2004. “Competition between private and public schools: testing stratification and pricing predictions”. Journal of Public Economics 88, 1215-1245.

[7] Fudenberg, D. and D. K. Levine, 1989. “Reputation and Equilibrium Selection in Games with a Single Long-Run Player”. Econometrica, Vol. 57, 759-778.

[8] Fudenberg, D. and D. K. Levine, 1992. “Maintaining a Reputation when Strategies are Imperfectly Observed”. Review of Economic Studies, 59: 561-579.

[9] Quinn, R., 1998. “Medical Malpractice Insurance: The Reputation Effect and Defensive Medicine”. The Journal of Risk and Insurance, Vol. 65, No. 3., 467-484.

[10] Hörner, J., 2002. “Reputation and Competition”, The American Economic Review, Vol. 92, No. 3., 644-663.

[11] Klein, B., and K. Leffler, 1981. “The Role of Market Forces in Assuring Contractual Performance.” Journal of Political Economy 89, 615-41.

[12] Mailath, G. and L. Samuelson, 1998. “Your Reputation Is Who You’re Not, Not Who You’d like To Be”, CARESS Working Paper 98-11.

[13] Mailath, G. and L. Samuelson, 2001. “Who Wants a Good Reputation?”, Review of Economic Studies 68, 415-441.

[14] Mailath, G., A. Postlewaite and L. Samuelson, 2006. “Pricing in Matching Markets”, paper presented at the European Meeting of the Econometric Society, Vienna, August.

[15] Nelson, P., 1970. “Information and Consumer Behavior”, The Journal of Political Economy, Vol. 78, No. 2, 311-329.

[16] Tadelis, S., 1999. “What’s in a Name? Reputation as a Tradeable Asset,” American Economic Review 89, 548–563.

A Proof of Proposition 1

Take x_n and x_m such that $n > m \geq N$, and define:

$$a = \max \{a' \in \mathbb{N} : (a' \leq m) \wedge (a' = k * N \text{ for some } k \in \mathbb{N})\}. \qquad (49)$$

Then, $0 \leq m-a < N$ and $a \geq N$. We can find bounds on the distance between x_n and x_m as follows:

$$\begin{aligned}
\rho\left(x_n, x_m\right) &= \rho\left(T^a x_{n-a}, T^a x_{m-a}\right) \\
&\leq \beta \rho\left(T^{a-N} x_{n-a}, T^{a-N} x_{m-a}\right) \\
&\leq \beta^2 \rho\left(T^{a-2N} x_{n-a}, T^{a-2N} x_{m-a}\right) \\
&\ldots \\
&\leq \beta^{a/N} \rho\left(x_{n-a}, x_{m-a}\right). \qquad (50)
\end{aligned}$$

Now, let:

$$d = \max\left\{\rho\left(x_i, x_j\right) : i, j \in [0, N]\right\}. \qquad (51)$$

If $n - a > N$, define:

$$b = \max\left\{b' \in \mathbb{N} : (b' \leq n-a) \wedge (b' = k * N \text{ for some } k \in \mathbb{N})\right\}, \qquad (52)$$

then $0 \leq n - a - b < N$.

Using the triangular inequality, we obtain:

$$\rho\left(x_{n-a}, x_{m-a}\right) \leq \rho\left(x_{n-a}, x_b\right) + \rho\left(x_b, x_{b-N}\right) + \ldots + \rho\left(x_N, x_{m-a}\right). \qquad (53)$$

Applying (50) to each addend but the last one, we obtain:

$$\begin{aligned}
\rho\left(x_{n-a}, x_{m-a}\right) &\leq \beta^{b/N} \rho\left(x_{n-a-b}, x_0\right) + \beta^{(b-N)/N} \rho\left(x_N, x_0\right) + \ldots + \rho\left(x_N, x_{m-a}\right) \\
&\leq \left(\beta^{b/N} + \beta^{(b-N)/N} + \beta^{(b-2N)/N} + \ldots + 1\right) d \\
&= \sum_{i=0}^{b/N} \left(\beta^i\right) d \\
&\leq \sum_{i=0}^{\infty} \left(\beta^i\right) d \\
&= \frac{d}{(1-\beta)}. \qquad (54)
\end{aligned}$$

If $n - a \leq N$,

$$\rho\left(x_{n-a}, x_{m-a}\right) \leq d < \frac{d}{(1-\beta)}. \qquad (55)$$

Therefore, from (50), (54) and (55) we obtain:

$$\rho\left(x_n, x_m\right) \leq \beta^{a/N} \frac{d}{(1-\beta)}. \qquad (56)$$

Since $\beta < 1$ we can take m sufficiently large in order to make $\beta^{a/N}$ arbitrarily small. Therefore, $\rho\left(x_n, x_m\right)$ can be made arbitrarily small taking m sufficiently large. That is, $(\forall \varepsilon > 0)\left(\exists N(\varepsilon)\right) : \rho\left(x_n, x_m\right) < \varepsilon$ if $n, m \geq N(\varepsilon)$ and so the sequence is Cauchy. As X is complete, the sequence $\{x_n\}$ converges to some element $x^*$of X.

As $\rho(Tx, Ty) \leq \rho(x, y)$, we know that T is continuous: taking $\delta = \varepsilon$, we obtain that $(\forall \varepsilon > 0)(\exists \delta > 0) : \rho(x, y) < \delta \Rightarrow \rho(Tx, Ty) \leq \rho(x, y) < \varepsilon$. Continuity of T implies that:

$$\lim_{n\to\infty} T(x_n) = T\left(\lim_{n\to\infty} x_n\right) = T(x^*), \tag{57}$$

where $x^* \equiv \lim_{n\to\infty} x_n$. But $\lim_{n\to\infty} T(x_n) = \lim_{n\to\infty} x_{n+1} = x^*$. Therefore, x^* is a fixed point of T in X. The fixed point of T is unique, since if x and x' were two fixed points of T, we would know that $x = Tx = T^N x$ and $x' = Tx' = T^N x'$. Therefore,

$$\begin{aligned} \rho(x, x') &= \rho\left(T^N x, T^N x'\right) \leq \beta\rho(x, x') \qquad \text{with } \beta < 1 \\ &\Leftrightarrow \rho(x, x') = 0, \end{aligned} \tag{58}$$

so that $x = x'$.

B Proof of Lemma 1

From the definitions of G^C_{t+1} and G^I_{t+1} in equations (13) and (14) we get:

$$\begin{aligned} \sup_x \left|G^C_{t+1}(x) - H^C_{t+1}(x)\right| = \sup_x \big|&(1-\lambda+\lambda\theta)\left[\pi_1\left(G^C_t\left(\mu_H^{-1}(x)\right) - H^C_t\left(\mu_H^{-1}(x)\right)\right)\right. \\ &\left.+(1-\pi_1)\left(G^C_t\left(\mu_L^{-1}(x)\right) - H^C_t\left(\mu_L^{-1}(x)\right)\right)\right] \\ &+\lambda(1-\theta)\left[\pi_0\left(G^I_t\left(\mu_H^{-1}(x)\right) - H^I_t\left(\mu_H^{-1}(x)\right)\right)\right. \\ &\left.+(1-\pi_0)\left(G^I_t\left(\mu_L^{-1}(x)\right) - H^I_t\left(\mu_L^{-1}(x)\right)\right)\right]\big|, \end{aligned} \tag{59}$$

and

$$\begin{aligned} \sup_x \left|G^I_{t+1}(x) - H^I_{t+1}(x)\right| = \sup_x \big|&(1-\lambda\theta)\left[\pi_0\left(G^I_t\left(\mu_H^{-1}(x)\right) - H^I_t\left(\mu_H^{-1}(x)\right)\right)\right. \\ &\left.+(1-\pi_0)\left(G^I_t\left(\mu_L^{-1}(x)\right) - H^I_t\left(\mu_L^{-1}(x)\right)\right)\right] \\ &+\lambda\theta\left[\pi_1\left(G^C_t\left(\mu_H^{-1}(x)\right) - H^C_t\left(\mu_H^{-1}(x)\right)\right)\right. \\ &\left.+(1-\pi_1)\left(G^C_t\left(\mu_L^{-1}(x)\right) - H^C_t\left(\mu_L^{-1}(x)\right)\right)\right]\big|. \end{aligned} \tag{60}$$

Noticing that the supremum of the sum is weakly smaller than the sum of the supremum of each addend, and using it with the definition of ρ_∞, we obtain:

$$\begin{aligned} \sup_x \left|G^C_{t+1}(x) - H^C_{t+1}(x)\right| \leq &(1-\lambda+\lambda\theta)\left[\pi_1\rho_\infty\left(G^C_t, H^C_t\right) + (1-\pi_1)\rho_\infty\left(G^C_t, H^C_t\right)\right] \\ &+\lambda(1-\theta)\left[\pi_0\rho_\infty\left(G^I_t, H^I_t\right) + (1-\pi_0)\rho_\infty\left(G^I_t, H^I_t\right)\right], \end{aligned} \tag{61}$$

and

$$\begin{aligned} \sup_x \left|G^I_{t+1}(x) - H^I_{t+1}(x)\right| \leq &(1-\lambda\theta)\left[\pi_0\rho_\infty\left(G^I_t, H^I_t\right) + (1-\pi_0)\rho_\infty\left(G^I_t, H^I_t\right)\right] \\ &+\lambda\theta\left[\pi_1\rho_\infty\left(G^C_t, H^C_t\right) + (1-\pi_1)\rho_\infty\left(G^C_t, H^C_t\right)\right]. \end{aligned} \tag{62}$$

Hence we obtain the desired result:

$$\begin{aligned}\rho\left(\left(G_{t+1}^{C},G_{t+1}^{I}\right),\left(H_{t+1}^{C},H_{t+1}^{I}\right)\right) \leq \max\{&(1-\lambda+\lambda\theta)\,\rho_{\infty}\left(G_{t}^{C},H_{t}^{C}\right)+\lambda\,(1-\theta)\,\rho_{\infty}\left(G_{t}^{I},H_{t}^{I}\right)\\ &,(1-\lambda\theta)\,\rho_{\infty}\left(G_{t}^{I},H_{t}^{I}\right)+\lambda\theta\rho_{\infty}\left(G_{t}^{C},H_{t}^{C}\right)\}\\ &\leq \max\left\{\rho_{\infty}\left(G_{t}^{C},H_{t}^{C}\right),\rho_{\infty}\left(G_{t}^{I},H_{t}^{I}\right)\right\}\\ &=\rho\left(\left(G_{t}^{C},G_{t}^{I}\right),\left(H_{t}^{C},H_{t}^{I}\right)\right). \quad (63)\end{aligned}$$

If the inequality in (63) is not strict for some distribution functions $\left(G_{t}^{C},G_{t}^{I}\right),\left(H_{t}^{C},H_{t}^{I}\right)\in X^{C}\times X^{I}$ such that $\rho\left(\left(G_{t}^{C},G_{t}^{I}\right),\left(H_{t}^{C},H_{t}^{I}\right)\right)>0$, then T is not a contraction. Two necessary conditions to obtain (63) with equality are:

Equality 1 (E1)

$$\begin{aligned}\arg\max_{\{x\}}&\left|G_{t}^{C}\left(\mu_{H}^{-1}(x)\right)-H_{t}^{C}\left(\mu_{H}^{-1}(x)\right)\right|\\ &\cap\arg\max_{\{x\}}\left|G_{t}^{C}\left(\mu_{L}^{-1}(x)\right)-H_{t}^{C}\left(\mu_{L}^{-1}(x)\right)\right|\\ &\cap\arg\max_{\{x\}}\left|G_{t}^{I}\left(\mu_{H}^{-1}(x)\right)-H_{t}^{I}\left(\mu_{H}^{-1}(x)\right)\right|\\ &\cap\arg\max_{\{x\}}\left|G_{t}^{I}\left(\mu_{L}^{-1}(x)\right)-H_{t}^{I}\left(\mu_{L}^{-1}(x)\right)\right|\neq\phi. \quad (64)\end{aligned}$$

If this condition is not met, it follows that the supremum of the sum is strictly smaller than the sum of the supremum of each addend in (59) and (60), because each addend is maximized in a different level of x.

Equality 2 (E2)

$$\rho_{\infty}\left(G_{t}^{C},H_{t}^{C}\right)=\rho_{\infty}\left(G_{t}^{I},H_{t}^{I}\right). \quad (65)$$

If this condition is not met, it follows that the average of $\rho_{\infty}\left(G_{t}^{C},H_{t}^{C}\right)$ and $\rho_{\infty}\left(G_{t}^{I},H_{t}^{I}\right)$ is strictly lower than $\max\left\{\rho_{\infty}\left(G_{t}^{C},H_{t}^{C}\right),\rho_{\infty}\left(G_{t}^{I},H_{t}^{I}\right)\right\}$ in (63).

C Proof of Proposition 2

Remember that we defined X^{τ} as the set of distribution functions G^{τ} with support contained in the interval $[\lambda\theta,\lambda\theta+(1-\lambda)]$, where $\tau\in\{C,I\}$ denotes the firm type, and with the distance between two distribution functions G^{τ} and H^{τ} as:

$$\rho_{\infty}\left(G^{\tau},H^{\tau}\right)=\sup_{x\in[\lambda\theta,\lambda\theta+(1-\lambda)]}\left|G^{\tau}(x)-H^{\tau}(x)\right|. \quad (66)$$

As distribution functions are bounded, if $\{G_n\}_{n=0}^{\infty}$ is a Cauchy sequence in X^{τ} with $\tau\in\{C,I\}$, it is also a Cauchy sequence in $\mathcal{B}(\mathbb{R},[0,1])$, the space of bounded functions from $\mathbb{R}$ into $[0,1]$ equipped with the same metric ρ_{∞} as X^{τ}. As $\mathcal{B}(\mathbb{R},[0,1])$ is a complete metric space, then there is a bounded function $G\in\mathcal{B}(\mathbb{R},[0,1])$ such that $G_n\to G$. We must show that $G\in X^{\tau}$, that is,

Condition 1 G is non-decreasing, i.e., $G(x_1) \geq G(x_0)$ whenever $x_1 > x_0$

Condition 2 G is right continuous, i.e., $\lim_{x\downarrow a} G(x) = G(a)$

Condition 3 G is normalized, i.e., $G(\lambda\theta) = 0$ and $G(\lambda\theta + (1-\lambda)) = 1$

Assume that G does not satisfy Condition 1. Then, there are two points $x_0, x_1 \in [\lambda\theta, \lambda\theta + (1-\lambda)]$ such that $x_1 > x_0$ and $G(x_1) < G(x_0)$. Define d as follows:

$$d \equiv G(x_0) - G(x_1) > 0, \tag{67}$$

and take $\varepsilon < \frac{d}{2}$. Then, if we consider any $G_n \in \{G_k\}_{k=0}^{\infty}$ such that $|G_n(x_0) - G(x_0)| < \varepsilon$, we know that

$$-\frac{d}{2} < G_n(x_0) - G(x_0) < \frac{d}{2}. \tag{68}$$

Since $G_n \in X^\tau$, G_n is non-decreasing, and therefore $G_n(x_1) \geq G_n(x_0)$. Then,

$$\begin{aligned}
G_n(x_1) - G(x_1) &= G_n(x_1) - G(x_0) + d & (69)\\
&\geq G_n(x_0) - G(x_0) + d & (70)\\
&> \frac{d}{2} & (71)\\
&> \varepsilon. & (72)
\end{aligned}$$

Therefore $|G_n(x_1) - G(x_1)| > \varepsilon$, and we conclude that $\rho_\infty(G_n, G) > \varepsilon$ for all $G_n \in \{G_n\}_{n=0}^{\infty}$. Hence, $G_n \nrightarrow G$, contradicting our hypothesis. We conclude that G is non-decreasing.

Now let x_0 be an arbitrary point of $[\lambda\theta, \lambda\theta + (1-\lambda)]$ and take $\varepsilon > 0$. Since $G_n \to G$, there is an n such that:

$$\rho_\infty(G_n, G) = \sup_{x \in [\lambda\theta, \lambda\theta + (1-\lambda)]} |G_n(x) - G(x)| < \varepsilon, \tag{73}$$

and since $G_n \in X$, G_n is right continuous and there is $\delta > 0$ such that if $0 < x - x_0 < \delta$, then:

$$|G_n(x_0) - G_n(x)| < \varepsilon. \tag{74}$$

Using the triangular inequality, we obtain that if $0 < x - x_0 < \delta$, then:

$$\begin{aligned}
&|G(x_0) - G(x)| \\
&\quad \leq |G(x_0) - G_n(x_0)| + |G_n(x_0) - G_n(x)| + |G_n(x) - G(x)| \\
&\quad < 3\varepsilon.
\end{aligned} \tag{75}$$

Therefore, G is right continuous at x_0 and, since x_0 was chosen as an arbitrary point of $[\lambda\theta, \lambda\theta + (1-\lambda)]$, we conclude that G is right continuous.

Finally, assume that $G(\lambda\theta) \neq 0$ or $G(\lambda\theta + (1-\lambda)) \neq 1$. Let's define d as follows:

$$d \equiv \max\{|G(\lambda\theta)|, |1 - G(\lambda\theta + (1-\lambda))|\}, \tag{76}$$

and take $\varepsilon < d$. As $G_n \in X$, G_n is normalized, we have:

$$\rho_\infty (G_n, G) = \sup_{x\in[\lambda\theta,\lambda\theta+(1-\lambda)]} |G_n (x) - G (x)| \geq d > \varepsilon \tag{77}$$

for all $G_n \in \{G_n\}_{n=0}^{\infty}$. Hence, $G_n \nrightarrow G$. We conclude that G is normalized.

Summarizing, G is non-decreasing, right continuous and normalized, therefore $G \in X^\tau$, and (X^τ, ρ_∞) is a complete metric space.

Now we must show that $\left(X^C \times X^I, \rho\right)$ is also a complete metric space. Since the distance between two elements of $X^C \times X^I$ is defined as

$$\rho\left(\left(G^C, G^I\right), \left(H^C, H^I\right)\right) = \max\left\{\rho_\infty\left(G^C, H^C\right), \rho_\infty\left(G^I, H^I\right)\right\}, \tag{78}$$

then a sequence in $X^C \times X^I$ is Cauchy if and only if the sequences defined for each component (G_n^C and G_n^I) are Cauchy. Therefore, for any Cauchy sequence in $X^C \times X^I$ we know that the sequences defined for each component (G_n^C and G_n^I) converge in X^C and X^I respectively, and this in turn implies that the Cauchy sequence of pairs $\left(G_n^C, G_n^I\right)$ converges in $X^C \times X^I$. Therefore, $X^C \times X^I$ is a complete metric space.

D Proof of Theorem 4

The consumers' choice is determined by:

$$\max_{\mu\in P^{-1}(b)} E[u] = \left(\mu u\left(1, y - p^*(\mu, b)\right) + (1-\mu)\, u\left(0, y - p^*(\mu, b)\right)\right), \tag{79}$$

where in the HQE under consideration, $P^{-1}(B) = U$ for all b. Shortening notation,

$$\max_{\mu\in U} E[u] = u(0, z) + \mu \triangle u(\cdot, z), \tag{80}$$

where $\triangle u(\cdot, z) = u(1, z) - u(0, z)$ and $z = y - p^*(\mu, b)$. This yields a "demand curve" (in fact, an optimal application policy for consumers) $\mu^d = \mu^d(b, y|p^*)$ that coincides with the equilibrium assignment $\mu^*(b, y)$ because it is feasible under the equilibrium admission policies. By the maximum theorem, $\mu^*(b, y)$ is a continuous function. An interior μ^* satisfies the FOC:

$$\frac{\partial E[u]}{\partial \mu} = E[u_z(\cdot, z)]\left(-\frac{\partial p^*}{\partial \mu}\right) + \triangle u(\cdot, z) = 0. \tag{81}$$

Therefore, for an interior solution the price function must be increasing in μ:

$$\frac{\partial p^*}{\partial \mu} = \frac{\triangle u(\cdot, z)}{E[u_z(\cdot, z)]} > 0, \tag{82}$$

provided that $p < y$.

The SOC is:

$$\frac{\partial^2 E[u]}{\partial \mu^2} = -2\frac{\partial p(\mu,b)}{\partial \mu}\Delta u_z(\cdot,z) + E[u_{zz}(\cdot,z)]\left(\frac{\partial p(\mu,b)}{\partial \mu}\right)^2 - E[u_z(\cdot,z)]\frac{\partial^2 p(\mu,b)}{\partial \mu^2} < 0. \quad (83)$$

The result is obtained from comparative statics. Taking the derivative of the FOC with respect to y and b respectively, we obtain:

$$\frac{\partial \mu}{\partial y} = \frac{\left(\Delta u_z(\cdot,z) - E[u_{zz}(\cdot,z)]\right)\frac{\partial p(\mu,b)}{\partial \mu}}{2\frac{\partial p(\mu,b)}{\partial \mu}\Delta u_z(\cdot,z) - E[u_{zz}(\cdot,z)]\left(\frac{\partial p(\mu,b)}{\partial \mu}\right)^2 + E[u_z(\cdot,z)]\frac{\partial^2 p(\mu,b)}{\partial \mu^2}},$$

$$\frac{\partial \mu}{\partial b} = \frac{\left(-\Delta u_z(\cdot,z) + E[u_{zz}(\cdot,z)]\frac{\partial p(\mu,b)}{\partial \mu}\right)\frac{\partial p(\mu,b)}{\partial b}}{2\frac{\partial p(\mu,b)}{\partial \mu}\Delta u_z(\cdot,z) - E[u_{zz}(\cdot,z)]\left(\frac{\partial p(\mu,b)}{\partial \mu}\right)^2 + E[u_z(\cdot,z)]\frac{\partial^2 p(\mu,b)}{\partial \mu^2}} \quad (84)$$

Replacing $\frac{\partial p(\mu,b)}{\partial \mu}$ with $\frac{\Delta u(\cdot,z)}{E[u_z(\cdot,z)]}$, and $\frac{\partial p(\mu,b)}{\partial b}$ with $\frac{\partial c(b)}{\partial b}$, we get:

$$\frac{\partial \mu}{\partial y} = \frac{\left(\Delta u_z(\cdot,z) - E[u_{zz}(\cdot,z)]\right)\frac{\Delta u(\cdot,z)}{E[u_z(\cdot,z)]}}{2\frac{\partial p(\mu,b)}{\partial \mu}\Delta u_z(\cdot,z) - E[u_{zz}(\cdot,z)]\left(\frac{\partial p(\mu,b)}{\partial \mu}\right)^2 + E[u_z(\cdot,z)]\frac{\partial^2 p(\mu,b)}{\partial \mu^2}} \quad (85)$$

$$\frac{\partial \mu}{\partial b} = \frac{-\left(\Delta u_z(\cdot,z) - E[u_{zz}(\cdot,z)]\frac{\Delta u(\cdot,z)}{E[u_z(\cdot,z)]}\right)\frac{\partial c(b)}{\partial b}}{2\frac{\partial p(\mu,b)}{\partial \mu}\Delta u_z(\cdot,z) - E[u_{zz}(\cdot,z)]\left(\frac{\partial p(\mu,b)}{\partial \mu}\right)^2 + E[u_z(\cdot,z)]\frac{\partial^2 p(\mu,b)}{\partial \mu^2}} \quad (86)$$

Under the assumption that perceived quality is a normal good, the numerators in (85) and (86) are positive. On the other hand, (85) and (86) have a common denominator, the negative of the expression in equation (83). Therefore, $\frac{\partial \mu}{\partial y}$ and $\frac{\partial \mu}{\partial b}$ are positive if and only if the SOC is met.

From the FOC we obtain

$$\begin{aligned}\frac{\partial k(\mu)}{\partial \mu} &= \frac{\triangle u(\cdot, y - p(\mu,b))}{E[u_z(\cdot, y - p(\mu,b))]} \\ &= \frac{\triangle u(\cdot, y - k(\mu) - c(b))}{E[u_z(\cdot, y - k(\mu) - c(b))]}.\end{aligned}$$

Therefore, any combination of y and b such that $y - c(b)$ is constant results in the same choice of μ. Hence, the equilibrium assignment function takes the form:

$$\begin{aligned}\mu^*(b,y) &= \phi(y - c(b)) \\ \text{with } \phi' &> 0.\end{aligned}$$

This equilibrium assignment function is uniquely determined from the market clearing condition:

$$\int_{b_{\min}}^{b_{\max}} F\left(b, \phi^{-1}(\mu) + c(b)\right) \mathrm{d}b = G(\mu),$$

where $b_{\min}$ and $b_{\max}$ are obtained from the conditions

$$\phi^{-1}(\mu) + c(b) \in Y \quad \text{and} \quad b \in B.$$

Finally, the price function is uniquely determined from the solution to the differential equation:

$$\frac{\partial k(\mu)}{\partial \mu} = \frac{\triangle u\left(\cdot, \phi^{-1}(\mu) - k(\mu)\right)}{E\left[u_z\left(\cdot, \phi^{-1}(\mu) - k(\mu)\right)\right]},$$

with a boundary condition $p(\mu_{\min}, \underline{b}) = \underline{p} < \underline{y}$. This solution satisfies that $\phi^{-1}(\mu) = y - c(b) > k(\mu)$, and therefore $k(\mu)$ is strictly increasing.