

A Dynamic Model of Certification and Reputation

MIHAELA VAN DER SCHAAR, University of California, Los Angeles
SIMPSON ZHANG, University of California, Los Angeles

Markets typically have many ways of learning about quality, with two of the most important being reputational forces and certification, and these types of learning often interact with and influence each other. This paper is the first to consider markets where learning occurs through these different sources simultaneously, which allows us to demonstrate the rich interplay and dynamics that can arise. Our work offers four main insights: (1) Without certification, market learning through reputation alone can get “stuck” at inefficient levels and high quality agents may get forced out of the market. (2) Certification “frees” the reputation of agents, allowing good agents to keep working even after an unfortunate string of bad signals. (3) Certification can be both beneficial and harmful, and so the social planner must choose the certification scheme carefully. In particular, the market will tend to demand more certification than socially optimal because the market does not bear the certification costs. (4) Certification and reputational learning can act as complementary forces so that faster information revelation will increase the social welfare generated by certification.

Categories Descriptors: **J.4 [Social and Behavioral Sciences]**

General Terms: Certification, Reputation, Social Welfare, Design

Additional Key Words and Phrases: Economics of information, reputation and trust systems, market design, mechanism design

ACM Reference Format:

Mihaela van der Schaar and Simpson Zhang, 2014. A Dynamic Model of Certification and Reputation. *ACM Trans. Embedd. Comput. Syst.*
DOI:<http://dx.doi.org/10.1145/0000000.0000000>

URL Link:

<http://medianetlab.ee.ucla.edu/papers/DynamicCertificationReputation.pdf>

1. INTRODUCTION

A market typically has several avenues from which to learn about the quality of an agent, such as the agent’s prior work history, or reports from trusted outside sources like a certification board. An agent’s work results are often readily available and provide a steady, but noisy, flow of information to the market, continuously updating the agent’s reputation. Certification boards also exist in many industries such as healthcare and accounting, and certification verifies that a minimum level of quality has been met. Given that these different sources of information often coexist, important questions are raised regarding how they can, and should, interact in a marketplace. As an example, physicians can become board certified in order to demonstrate their proficiency in a medical specialty, and patients can also learn about physicians through the results of past procedures. In such a market, should board certification be undertaken by a physician immediately, or after some time to allow the market to learn from the doctor’s results? And how stringently should the medical board set its certification requirements if it wishes to maximize social welfare? Although a substantial economics literature has explored environments in which markets learn quality through each source of information separately, the combined effects of the two have not yet been well analyzed. We believe this paper is the first to explore markets with access to both sources.

Our main results show that certification remains a critical component of a marketplace even when learning can be done through other channels. The key observation is that certification provides a very different kind of information to the market, a type of information that is beneficial both due to its informativeness as well as its dependability. Without certification, even high quality agents may produce an

unfortunate string of bad results and so get kicked out of the market. Once out of the market, agents can no longer work and thus cannot produce more information that would change the market's opinion. Certification, however, remedies this problem by acting as a safety net, letting these socially beneficial agents continue to work. The type and amount of information revealed by the agent's work history has a strong effect on the social welfare provided by certification, and we show that contrary to the case without certification, faster information revelation can be socially beneficial. Thus certification and the agent's work history can act as complements.

Our model features reputational dynamics related to those in Holmström (1999), with an agent of uncertain quality working for a market of homogenous short lived principals. The agent knows its own quality, but the market does not and must infer quality through the agent's work history, which generates a stochastic process that is publically observable. The model is very general and imposes few assumptions on the stochastic process itself, or the agent's initial quality distribution. In contrast, most papers in the reputation literature make strong assumptions such as binary quality types with Poisson or Brownian motion signals. The agent's work history provides a steady stream of information to the market, continuously updating the agent's reputation. However, the agent can also pay a cost to certify at any moment in time, which sends additional information to the market. Certification verifies that true quality lies above some standard \underline{q} , but does not reveal the exact quality level. This type of "imperfect" certification is one of the most common in practice, with examples such as board certification for doctors, pass/fail exams for accountants, and security seals for websites.

The need for certification even when other sources of information are available is most apparent in markets with frictions that can cause information revelation to become "stuck" at inefficient levels. In this paper, such a friction manifests in the form of a fixed price for the agent. This type of pricing is applicable to many real life situations where wages are sticky. Large companies often offer a set starting salary to all workers independently of their reputation and choose simply not to hire agents that have bad reputations. Negotiations with unions may also fix prices for workers. Moreover, in some markets the government will mandate a fixed price. For instance, doctors who see Medicare patients are paid according to a fixed fee for service schedule. Patients can choose which doctor to see and will not go to ones that have bad reputations. Finally, in settings without monetary transfers the prices are often a fixed cost. For example people must decide whether to watch television shows based on the opinions of others and at the cost of a fixed amount of time.

We first consider a benchmark model without certification and learning only through the agent's work history. Even in this simpler setting, we find surprising results. The key driver of these results is that with fixed prices, high quality agents will be forced out of the market after a string of bad signals instead of being able to work forever. Thus learning may be halted inefficiently and having access to the agent's work history does not guarantee a socially-efficient outcome. In fact, we show that for many types of reputational mechanisms, it does not provide any social benefit at all. This result is reminiscent of work on "bad reputation" such as Ely and Välimäki (2003) and Ely et al (2008), where the ineptness of reputation results from the myopia of the short-lived principals. In our model, this myopia results in an informational friction. Principals are not willing to hire once the agent's expected quality falls to the price level, but there is still positive social value from hiring at this point due to experimentation. Thus principals are kicking agents out of the market inefficiently due to their short-sightedness.

With the benchmark case in hand, we then introduce certification into the model. The decision to become certified or not at any given moment is a strategic one, which depends on the agent's previous work history and the market's beliefs about the agent's strategy. Many different beliefs are possible and there can be many equilibria, a common feature of most signaling games. We characterize the specific equilibrium beliefs that maximize principal, agent, and social welfare, and we show a relative ordering on these beliefs: agents prefer less certification than socially optimal and principals prefer more certification than socially optimal.

After characterizing these three different types of equilibrium, we then perform comparative statics to find the optimal certification standard and price level for each type of equilibrium. We find that certification increases welfare if the standards are set appropriately and certification costs are low. In addition, the information produced by the agent's work history and certification can act in a complementary fashion, with faster information revelation increasing the welfare generated by certification. But welfare could also decrease if the standards are set poorly. One case is when certification costs are very high, because then it may be better for there not to be certification at all. Another possibility is that even when costs are low, the equilibrium market beliefs expect too much certification. Such a case can occur in the principal optimal equilibrium, since principals do not internalize certification costs. Because agents are forced to certify at inefficient reputation levels, overall welfare could actually be reduced.

2. LITERATURE REVIEW

2.1 Reputation Papers

Our work is closely related to several papers in the reputation literature, such as Holmström (1999), Bar-Isaac (2003), and Ely and Välimäki (2003). All of these papers have some form of reputational mechanism that follows from work history, but ours is the first to allow the agent to certify and send information through this channel. Holmström (1999) presents the classic "signal jamming" model where an agent of unknown quality can exert effort to bias the market's perception. Contingent contracts are not possible and without reputational incentives the agent would exert no effort. Holmström finds that reputation can provide work incentives in the short run while the agent's quality is unknown, but not in the long run once quality does become known. Crucially, there is no exit in this model and so the agent's quality will become known perfectly over time. We show that when there is an exit point, perhaps due to market frictions resulting in fixed prices, then agents may be forced out before their true quality is revealed. In this case certification is necessary as a form of insurance, so that high quality agents can stay in the market and continue selling.

A reputational paper that does focus on exit is Bar-Isaac (2003). This model arrives at the striking conclusion that high quality agents never exit the market, because staying in the market is a signal of quality that increases reputation. Even an agent who receives a string of bad signals can demonstrate resolve by refusing to quit, which boosts the market's perceptions. Importantly, Bar-Isaac assumes that the price varies at every moment in time to equal the agent's expected quality, so good agents can internalize the future benefits of reputation and are thus willing to sustain a period of negative payoffs. We refrain from a similar pricing assumption for a manifold of reasons. First, such an assumption requires the agent to have market power (in fact the agent is extracting full surplus from principals), which can be unrealistic in many types of markets where agents have unknown qualities. Without market power agents may be unable to internalize the benefits of high reputation, and thus unwilling to

endure negative flow payoffs. Second, good agents may be forced to sustain numerous periods of negative payoffs in equilibrium, which may not be feasible if they have liquidity constraints. Finally, such an assumption requires prices to be extremely fluid, when in actuality market customs and menu costs can make changing prices quite difficult. When prices are fixed, we show an opposite conclusion results: high quality agents are forced to exit the market with positive probability and the learning may not be socially efficient.

While most papers in the literature show that reputation can increase social welfare, the papers by Ely and Välimäki (2003) and Ely et al (2008) are noteworthy in showing cases where reputation is potentially harmful. Ely and Välimäki (2003) considers a car repair framework, with both honest and dishonest mechanics. Honest mechanics wish to recommend the best repair for a car, be it a cheap tune-up or an expensive engine repair, whereas dishonest mechanics always want to recommend engine repairs. The mechanic's reputation represents the probability of being honest, which will change over time depending on how many tune-ups or engine repairs are recommended. The paper shows that the reputational mechanism may actually destroy social surplus: honest mechanics have an incentive to recommend tune-ups even when engine repairs are truly needed in order to boost reputation. But such an action hurts the consumer, who is short-lived and so does not internalize the benefit of higher reputation for the mechanic. In equilibrium, the market may break down as consumers are not willing to go to either type of mechanic. Ely et al (2008) extends this model to more general frameworks and shows that similar results will hold in models that feature "temptation" actions for the long-lived player. These "temptations" are actions that boost reputation but are socially inefficient, like recommending tune-ups in place of engine repairs. The assumption of short-lived consumers is critical in both papers, as Ely and Välimäki (2003) shows that a long-lived consumer could devise a mechanism that benefits from the information process. Such a result mirrors that of our model, which also considers a market of short-lived players. Because of myopia, the market is unable to internalize the benefits of experimentation and so learning may stop early at an inefficient level. But our model does not require moral hazard in the way of "temptations", as the certification choice does not hurt the market. Still, we show that the market myopia by itself is sufficient to create a "bad reputation" effect when prices are fixed and agents are forced out of the market inefficiently.

2.2 Certification Papers

Our work represents a novel approach to the certification literature, where certification has never been combined with a similar reputational mechanism in a dynamic setting. Much of the focus of the theoretical certification literature is on the certifier's actions and revolves around studying the decisions made by a strategic certifier who can control the type of information that it releases about agents or the payments that it charges (e.g. Lizzeri (1999), Stahl and Strausz (2010), Farhi et al (2013)). Some papers also allow the certifier to collude with agents and assign false ratings as in Strausz (2004). In contrast, our paper focuses more on the agent's decision process and how its reputation affects its certification decision and the principals' beliefs. This allows us to analyze the strategic aspects of certification for the agent, and show how even the possibility of certification will affect an agent's reputation. We do not explicitly model the certifier, instead delegating it to the role of the mechanism designer and analyzing the comparative statics of our model with regard to the certification standard. This implicitly assumes that the certifier is not strategic, is always accurate in its judgments, and is not allowed to cheat. Thus our work is best

suited to markets where the certifier is a government agency, which would be trustworthy and sets certification standards in a benevolent fashion.

Although our work on combining agent reputation and certification is new from a theoretical perspective, there have been interesting empirical papers published on this subject. For instance, Xiao (2010) tests the value of certification in the childcare industry. In childcare, certification is voluntary and usually chosen by only a small fraction of firms. The paper shows that the social value generated by certification is positive, albeit small. Further, the costs of certification are not negligible, such as administrative and personnel costs. Thus certification on the whole is a negative for the industry, as the benefits are outweighed by the costs. In our model, we show that this is certainly a possibility when certification is present; in particular when market beliefs are cautious and require certification often. In such a situation, the informational value of certification is outweighed by the inherent costs, resulting in a decrease of surplus.

Jin and Leslie (2003) and Jin and Leslie (2009) analyze another type of certification through looking at the impact of government issued grade cards on restaurant cleanliness. The authors find that certification has a strong effect on the level of restaurant cleanliness even when the restaurant also has reputational concerns. In addition, restaurant revenue is strongly influenced by these grades, with high scoring restaurants earning more revenue and low scoring restaurants earning less. This results mirror the results in our paper, which show that certification will improve the market's perceptions and that certification can increase social welfare even when reputational forces are present.

3. THE MODEL

We consider an infinite horizon continuous time model with a single long lived agent and a marketplace of principals¹. The agent has a fixed quality q that is determined at the start of the game according to a commonly known continuous distribution $q \sim f_0(q)$. Denote the support of this distribution by D , which can be any (non-null) set in \mathbb{R} . The agent is privately informed about its own quality, but the market must learn this quality over time. Let μ_0 denote the mean of the initial quality distribution. We assume the agent has a reservation value of $c < \mu_0$, which represents the agent's disutility of work and outside options. At each moment in time the agent meets a different short lived principal, who can make an offer to the agent at a fixed price p . We will assume throughout the paper that the price satisfies $p \geq c$ so that agents are always willing to accept an offer.

If the agent accepts the offer, it will work for one infinitesimal instant of time and return some output with quality q . The observations (or reports) of q are noisy, and the evolution of these observations follows a càdlàg stochastic process $R(t; q)$. For every possible value of q , a different stochastic process can be generated. We define the agent's *work history* as the history of all previous observations by the market, $\mathcal{H}_t = \{R(t')\}_{t'=0}^t$. We also call the set $\mathcal{L} \equiv \{R(t; q)\}_{q \in D}$ an *information process* for the market. The set \mathcal{L} is common knowledge, so the market knows what kind of stochastic process to expect from each type of agent. Thus, given any history \mathcal{H}_t , the market can update its beliefs about the agent via Bayes' rule. We define the agent's time t *work history reputation* as the market's belief about quality given only information from the work history, $f_t(q|\mathcal{H}_t)$.

¹ Alternatively we can think of the agent as representing a firm and the principals as representing consumers.

We make some restrictions on the information process \mathcal{L} to ensure that the agent's reputation will evolve sufficiently continuously over time. Let $\mu_t \equiv E[f_t(q|\mathcal{H}_t)]$ be the time t expected mean of the agent's quality, and let $\mu_t^q \equiv E[f_t(q|\mathcal{H}_t, q < \underline{q})]$ be the time t expected mean of the agent's quality given that it lies below some cutoff \underline{q} . We say that the information process is *admissible* if almost surely μ_t and $\{\mu_t^q\}_{q \in \mathbb{R}}$ are right-continuous and upper semi-continuous in time. That is, at any point of discontinuity these means can jump upwards but not downwards. In this paper we will only focus on admissible information processes. Section A of the online appendix details examples of admissible information processes, which can range from Brownian motion to Poisson signals.

We further define a *blind process* to be any information process such that $R(t; q) = R(t; q') \forall q, q' \in D$. Under a blind process, the information provided to the market is completely uninformative of the agent's true quality, and so no learning occurs. At all times t , the agent's work history reputation is thus the same as the initial prior, $f_t(q|\mathcal{H}_t) = f_0(q)$.

Finally, note that the stochastic process produced by the agent will only run as long as the agent is still working. If at a time t the principal does not hire the agent, no output gets produced and the stochastic process is stopped at that value of $R(t)$. Since the agent is not working, no further information gets sent. If the agent does get hired again in the future, then the stochastic process can once again proceed at that time. We will show that this can only happen if the agent passes certification; that is unless the agent certifies he will be kicked out of the market forever.

At every moment of time t , the agent can choose to certify if it has not done so already by paying a certification cost k and getting certified if it has a quality level of at least \underline{q} , which is a fixed exogenous standard. We assume that $\underline{q} \geq p$, which will imply that certified agents never get forced out of the market. We also assume that $k \leq \frac{p-c}{\rho}$ so that the cost is low enough for certification to be chosen (this inequality implies that the net present value of staying in the market forever is greater than the certification cost). We denote the time t certification status by $\theta_t \in \{[0, t] \cup \phi\}$, where a number represents the time at which that agent became certified and ϕ means that the agent has not yet certified. Once the agent becomes certified, all future principals will know for certain that the agent quality is at least \underline{q} . Since this information is permanent and public knowledge to all future principals, an agent only needs to be certified once. We assume principals do not know if an agent attempts to certify and fails, so an agent that cannot pass certification has no incentive to certify.

There are thus two sources of information in our market: the first being the history of observations from work history, $\mathcal{H}_t = \{R(t')\}_{t'=0}^t$, and the second being the certification status of the agent θ_t . Upon observing these sources of information, the market uses Bayes' rule in updating its beliefs about agent quality. Together, these two signals combined with the prior quality distribution will result in a posterior belief distribution of agent quality $f_t(q|\theta_t, \mathcal{H}_t)$, which we call the *reputation* of the agent. This is different from the work history reputation that was defined previously because it also takes into account the agent's certification status. Note that the reputation can be calculated from the work history reputation by using Bayes' rule together with θ_t on

the work history reputation $f_t(q|\mathcal{H}_t)$. We will use $F_t(q|\theta_t, \mathcal{H}_t)$ to denote the cdf of the reputation².

We assume principals have identical linear utilities given by $U(p, q) = q - p$. Since principals are short-lived, they care only about maximizing their current utility values. Given the above linear utility function, principals will hire if and only if $p < \mu_t \equiv E[q|\theta_t, \mathcal{H}_t]$, where the expectation is taken with respect to the agent's time t reputation. For technical reasons, we assume principals who are indifferent will hire certified agents, but not uncertified agents. Note that once the agent's expected quality falls below p the agent will not be able to work again unless it certifies, because principals would not be willing to hire at any quality level less than p , so no outputs will be produced.

If the agent works at time t , it receives a flow payoff of $\pi_t = p - c$. Otherwise, the agent receives a payoff of $\pi_t = 0$. The agent also has a discount rate of ρ . Thus at time t the agent decides whether to certify by maximizing its expected discounted value, $\int_{t'=t}^{\infty} e^{-\rho t'} E[\pi_{t'}|\mathcal{H}_t, q, \underline{q}] dt'$, taking into account the cost of certification and given the state variables: the history of outputs \mathcal{H}_t , the agent's true quality q , and the certification standard \underline{q} . The agent's certification strategy is an optimal stopping time that is measurable with respect to the filtration generated by $R(t; q)$. We denote this stopping time by $\tau(q)$, and we note that $\tau = \infty$ for all $q < \underline{q}$ since these agents cannot pass certification. The market must also have beliefs about the agent's certification strategy, and we represent these beliefs using $\tilde{\tau}(q)$. In equilibrium, these beliefs must be correct so that $\tilde{\tau}(q) = \tau(q)$.

4. BENCHMARK CASE : NO CERTIFICATION

In order to analyze the welfare benefits of certification, we will first derive some welfare results for a benchmark setting with no certification. This can be considered a specialized case of the above model with $\underline{q} = \infty$. A principal will hire at time t if and only if $p < E[q|\theta_t, \mathcal{H}_t]$, with the expectation taken with respect to the agent's reputation, $f_t(q|\theta_t, \mathcal{H}_t)$. Since the agent will never certify, its reputation is always equal to the work history reputation, $f_t(q|\mathcal{H}_t)$. To analyze this model we thus need to determine how the agent's reputation evolves as a function of its work history. Working generates a stochastic process $R(t; q)$, and without certification the market will update its beliefs only through this stochastic process itself. Through Bayes' rule the agent's expected quality level, $\mu_t \equiv E_t[q|\mathcal{H}_t]$, will be continuously updated as the market receives information. The reputational dynamics in the benchmark are as follows: the agent starts out with a quality $\mu_0 > p$ (for $\mu_0 \leq p$ no hiring ever takes place) and continues working as long as its expected quality stays higher than p , but the first time that its expected quality drops to or below p it will be forced to stop. At this point, no principals are willing to hire the agent, and thus the stochastic process is stopped forever causing no future principals to hire either.

Given this characterization of the dynamics, once the agent stops working it can never start again. Thus, the probability that the agent is still working at time t , denoted by $P(S_t|q, p, \mathcal{L})$, is the same as the probability that the hitting time of μ_t against the price p or any lower value is greater than t . Since we only consider admissible information processes, at the first time that μ_t falls below p it must be that $\mu_t = p$. This is due to the fact that the expected mean can never jump downwards. Let

² We sometimes suppress the notation and write $f_t(q)$ and $F_t(q)$ for the reputation and cdf of the reputation respectively.

$t^* = \inf\{t | \mu_t = p\}$ denote this (stochastic) hitting time. So $P(S_t | q, p, \mathcal{L})$, the probability the agent is still working at time t , is equal to the probability that $t^* > t$. Now we can write out the *ex ante* expected social surplus (sum of principal and agent surplus) for a price p as:

$$W(p, \mathcal{L}) = \int_{-\infty}^{\infty} \left[\int_0^{\infty} e^{-\rho t} (q - c) P(S_t | q, p, \mathcal{L}) dt \right] f_0(q) dq$$

For an agent with quality q , the social welfare at every moment will be given by $q - c$ if the agent is hired and 0 otherwise. Thus, hiring is socially optimal if and only if $q \geq c$. There are two sources of inefficiency in our model: perceived quality may be above p even though true quality is below c so that principals are hiring when they should not (bad agents are working), or perceived quality is below p even though true quality is above c so that principals are not hiring when they should (good agents are not working).

Although in general the working probability may not have an analytic expression, and so the social surplus cannot be directly computed, we can nonetheless still prove results about welfare. We show a perhaps surprising result: for all information processes that are admissible, the *ex ante* social welfare will be bounded above by the blind process. That is, the information provided through the reputational mechanism cannot actually improve social welfare.

THEOREM 1 (BLIND BOUNDEDNESS):

Let \mathcal{L}^b denote any blind process. For any admissible information process \mathcal{L} , $W(p, \mathcal{L}) \leq W(p, \mathcal{L}^b)$ for all $p \in \mathbb{R}$.

PROOF: See online appendix. ■

The intuition for this theorem is that since the market is kicking out agents when their expected quality hits p , the social welfare impact must be weakly negative for any $p \geq c$ because beliefs are correct, and so the expected quality of an agent that leaves *must* be higher than the social cost of that agent. If, at the time the agent is getting kicked out, we could instead choose to let that agent stay in and work forever, we would wish to do so. Then since under the blind process no agent ever gets kicked out, the blind process must be weakly better than any other admissible information process. So all admissible information processes give a social welfare that is bounded above by the blind process.

In addition, from the proof we can see that the social welfare for any information process is the same if $p = c$, but can be strictly less if $p > c$. This implies that for *any* admissible information process, the optimal price is equal to the agent's reservation value. When price is exactly equal to the reservation value, we are indifferent between letting agents with $\mu_t = p$ stay and work forever, or leave immediately. But when $p > c$ we strictly prefer to have agents that are being kicked out stay in because they are still generating positive flow payoffs in expectation. Thus we want to set the price as low as possible in order to make the exiting process more efficient. Note that if we could subsidize agents to work below their reservation values, we would because an agent with expected quality equal to price still generate positive value to society through the informational benefits of working.

As a reason for the results in this section, it is important to note that the principals are *short-lived*, which creates a source of inefficiency in the marketplace because each principal does not internalize the social benefits of experimentation. Once the agent's expected quality level reaches the price, it is still socially beneficial for the agent to

work, because working sends information to the market and this information is valuable. A long-lived principal for instance would wish to continue hiring the agent because there is a chance the agent's reputation could improve and thus a positive value from experimentation. This result is related to results in papers such as (Ely and Välimäki 2003) and (Ely et al 2008) which show other cases where a reputational mechanism may not be beneficial when the market has short-lived players who do not internalize the benefits of experimentation.

5. GENERAL MODEL

5.1 Belief Updating with Certification

Now let us consider the general model with both work history and certification. If certification is possible, agents will have a decision to make regarding the exact time to undergo certification. For an agent whose true quality is above \underline{q} and has not undergone certification before time t , its expected discounted value if it chooses certification at time t will be $\int_t^\infty e^{-\rho(\tau-t)} E[\pi_t | \theta_t = t, \mathcal{H}_t] d\tau - k$. We can analyze π_t by solving the principals' problem. Principals update the agent's reputation with Bayes' rule, and they will hire the agent if the expected quality level is above the price given the certification status and history of outputs, or $p < E[q_t | \theta_t, \mathcal{H}_t]$.

Certification combined with the work history \mathcal{H}_t will result in a posterior that is a truncated distribution. Suppose for instance that in equilibrium all agents with quality at least \underline{q} are expected to certify at time 0. Given this fact, agents that become certified at time 0 will be believed to have a quality higher than \underline{q} . If the prior distribution is distributed $f_0(q)$, then the posterior distribution given certification will be $f_0(q)$ truncated to be over the interval $[\underline{q}, \infty)$. Since all agents in this quality region choose to certify, the relative density inside the support is unchanged. Likewise, the reputation for an agent that does not certify will also be updated. The market will believe that the agent has a quality less than \underline{q} , and so the posterior becomes $f_0(q)$ truncated to be over the region $(-\infty, \underline{q})$. An agent that certifies will continue to work forever (since we have assumed $\underline{q} \geq p$), and an agent that does not certify will stop working if the truncated mean drops below the price.

The key with certification, and the reason that it can improve welfare, is that high quality agents who certify will be able to keep selling even if they send an unlucky string of bad signals. Without certification such agents would instead be inefficiently kicked out of the market. In this way, certification is "freeing" the reputation of good agents which could otherwise get "stuck". Another way to look at it is from the perspective of the agents that are getting kicked out because they do not certify. Since their posterior gets truncated downwards, the expected mean of such agents can fall below the price level. Thus unlike the benchmark case where agents got kicked out when $\mu_t = p$, here we are kicking out agents with an expected mean less than the price, and so kicking out these agents is socially efficient. This aspect of certification is welfare improving, but since certification also has a cost, welfare may decrease as well if certification is not implemented correctly.

5.2 Agent's Certification Strategy

We will start solving the model by considering what the agent's optimal certification stopping time strategy $\tau(q)$ should be. In general this will also depend on what the market's beliefs $\tilde{\tau}(q)$ are, but we will show that for any beliefs the agent's strategy will have a common structure.

Throughout this section it is important to keep in mind two assumptions that we are making. For convenience, we assume that the certification standard $\underline{q} \geq p$, which implies that the agent will never stop working if it certifies. Using this assumption reduces the set of equilibria allowing for greater tractability, and implies that once certified the principals no longer need to doubt the agent's quality. With this assumption, the payoff after certifying is deterministic and given by $\frac{p-c}{\rho} - k$. The second important assumption is a restriction on the price, specifically that $p > k\rho + c$, which implies that $\frac{p-c}{\rho} - k > 0$. This condition means that the price is high enough to cover the cost of certification. Thus an agent that is being forced out of the market would choose to certify and continue working instead of exiting forever. Without this assumption, the only possible equilibrium would involve no certification.

We show in the following theorem that under our assumptions the only possible types of equilibrium are *type independent equilibrium*, where after any work history all agents with quality above \underline{q} will choose to certify, or none will. Thus in equilibrium all agents will use the same stopping time strategy τ^* .

THEOREM 2 (TYPE INDEPENDENCE):

In any equilibrium, all agents with quality $q \geq \underline{q}$ have the same certification stopping time strategy τ^ .*

PROOF: See online appendix. ■

Given our type independence result, we will maintain the following notation in this paper: for any time t , let μ_t^{NC} represent the mean of the non-certifying agents and μ_t^C represent the mean of certifying agents, if the market were to believe that all agents with quality $q \geq \underline{q}$ would choose to certify at that t . Note that there are some histories in which both no types of agents certifying or all types of agents certifying are supportable in equilibrium. The proof of Theorem 2 implies that as long as $\mu_t \geq p$ no agents certifying is supportable in equilibrium, and as long as $\mu_t^{NC} \leq p$, all agents certifying is supportable in equilibrium. The first condition indicates that if no agents are believed by the market to certify, all agents can keep working without certifying. The second condition indicates that if all agents with qualities in $[\underline{q}, \infty)$ are believed to certify, then agents that don't certify will be forced to stop working. Since it is possible that $\mu_t \leq p \leq \mu_t^{NC}$, we could have either type of behavior in equilibrium after some histories.

5.3 Equilibrium Characterization

In general, as is common in signaling games, many different equilibria are possible. In order to find out which specific equilibrium maximizes principal, agent, or social surplus, we need to find the specific beliefs and actions that will result in the highest payoffs to each party. Theorem 2 implies that all equilibrium can be summarized by a single stopping time strategy for all types of agents. We denote the agent optimal, socially optimal, and principal optimal equilibrium stopping times by τ_a^* , τ_s^* , and τ_p^* respectively. The next section characterizes the optimal equilibria for these three different cases. Theorem 3 summarizes our main findings: we see that principals prefer earlier certification than is socially optimal and that agents prefer later certification that is socially optimal.

THEOREM 3:

The optimal stopping times for the principal, agent, and social welfare cases must satisfy

$$\tau_p^* \leq \tau_s^* \leq \tau_a^*$$

PROOF: The result follows from a comparison of the stopping times in Propositions 1, 2, and 3 below. ■

5.3.1 Agent optimal equilibrium

We start with characterizing the best equilibrium from the agent's point of view. Since the agent is long-lived, the notion of optimality we will use is the one that maximizes the agent's expected value at time 0, given by $\int_0^\infty e^{-pt}(p - c)dt - E_{\tau_a^*}[e^{-p\tau_a^*}k]$ if $q \geq \underline{q}$ or $E_{\tau_a^*}[\int_0^\infty e^{-pt}(p - c)dt]$ if $q < \underline{q}$. In general, the agent's expected discounted value will depend on its own quality level, but we show that the specific equilibrium that maximizes the agent's value does not.

PROPOSITION 1:

The equilibrium that maximizes the agent's payoff is for $\tau_a^ = t^* = \inf\{t | \mu_t \leq p\}$.*

PROOF: See online appendix. ■

Note that t^* is the absolute latest that the agent can certify in equilibrium because at this point the agent must stop working no matter what the market certification beliefs are. This proposition shows that the best equilibrium from the agent's perspective will always delay certification for as long as possible, with certification occurring just as they are getting kicked out of the market. Since agents bear the full cost of certification and do not directly benefit in terms of flow payoffs, they are much less willing to certify than principals would want them to. In this equilibrium the beliefs expect certification at the latest possible time, and so we will call them *optimistic* beliefs.

If $\mu_0 \leq p$ then agents cannot work initially and must certify to get hired. In this case certification will be a *de facto license*, with non-certifying agents never being able to work. Certification thus acts as an entrance key into the market itself. But as long as $\mu_0 > p$, in any agent optimal equilibrium certification will not be a de facto license. Instead the equilibrium will feature *delayed certification*, with agents working and sending information initially before needing to certify if their reputation drops too low. In this case the particular information process will make a big difference on the timing of the agent's certification decision, as well as the overall social welfare that is generated by certification.

5.3.2 Social welfare optimal equilibrium

Next, we characterize the equilibrium beliefs that maximize the total social welfare for the market. We will use the same notion of *ex ante* social welfare that was defined in section 4, meaning that we want to maximize the working probability of good agents while minimizing the working probability of bad agents. The formula is given by:

$$W = \int_{\underline{q}}^\infty \left[E \left[\int_0^\infty e^{-pt}(q - c)dt - e^{-p\tau_s^*}k | q \right] \right] f_0(q) dq$$

$$+ \int_{-\infty}^{\underline{q}} \left[E \left[\int_0^{\tau_s^*} e^{-\rho t} (q - c) dt \mid q \right] \right] f_0(q) dq$$

Early certification is beneficial because it verifies that an agent is good more quickly. However, since certification is also costly, we may want agents to delay certification if their reputation is sufficiently high, since that saves the discounted flow cost of certification, as well as keeps the option value of certifying for the future. The following proposition characterizes a *necessary condition* for certification to be socially optimal: the agent should certify only if its reputation is lower than the following cutoff.

PROPOSITION 2:

If at the socially optimal stopping time τ_s^ , the expected quality of the agent is strictly higher than the price, $\mu_{\tau_s^*} > p$, then the following condition must also be satisfied:*

$$\mu_{\tau_s^*}^{NC} \leq c + \rho k - \frac{\rho k}{F_{\tau^*}(q)} \quad (1)$$

PROOF: See online appendix. ■

We note that this bound can also be a sufficient condition for some information processes. With a blind process (1) gives the exact value of the truncated mean below which it becomes optimal to have the agent certify. Certification is socially beneficial because it gets rid of “bad” agents with qualities less than c , but it is socially harmful because it can get rid of “good” agents with qualities above c but below \underline{q} and it also carries a cost of k . Certification at any time t must balance these two aspects together, and (1) gives a condition on how low the mean of non-certifying agents (i.e. the agents that would be kicked out by certification) must be for certification to be beneficial.

However it is possible for (2) to only be satisfied at some $\mu_t < p$. In this case we would have $\tau_s^* = t^* = \inf\{t \mid \mu_t \leq p\}$, because when the expected quality falls below the price, the unique equilibrium is for an agent to certify. In this case, the agent optimal stopping time and the socially optimal stopping time would coincide. In general however, the socially optimal stopping time can be strictly later than the agent optimal stopping time due to the beneficial information that certification provides.

5.3.3 Principal optimal equilibrium

Finally, we analyze the equilibrium that maximizes the utility of the principals. The right notion of the principal utility is a bit tricky since each principal is short lived and only cares about itself. The equilibrium that gives the time t principal the highest payoff will in general be different than the equilibrium that gives the time t' principal the highest payoff. We resolve this tension in the following way: principals that arrive earlier should have a greater say in the equilibrium than principals who like later, because presumably the later principals may even know of the agent at earlier times. Thus we define a *principal optimal equilibrium* to be the equilibrium in which at any time t the agent remains uncertified if and only if all principals at all times before t would not have received a higher payoff if the agent had been required to certify. Such an equilibrium always exists, and can be found by considering, at every time t and reputation $f_t(q)$, whether the principal that is active at that time t wishes for the agent to certify *at that moment* if it has not done so already. The principal optimal equilibrium stopping time is then given by the first time t at which that time t principal prefers certification. This notion of principal optimality is not to be confused

with the net present value of principal surplus, since the myopic time t principal does not consider the welfare of future principals.

PROPOSITION 3:

The principal optimal equilibrium stopping time is

$$\tau_p^* = \inf\{t | \mu_t^{NC} \leq p\} \quad (2)$$

PROOF: See online appendix. ■

This Proposition gives a cutoff in terms of μ_t^{NC} , with principals wanting the agent to certify the first time its truncated expected quality falls below the price. Certification is good for principals because it gets rid of “bad” agents with qualities less than p , but it may also hurt principals because it can get rid of “good” agents with qualities between p and q . The cutoff $\mu_t^{NC} = p$ is the quality such that these two effects exactly balance out. We see from (2) that the best equilibria for principals are those where the certification strategy beliefs τ^* expect agents to certify at the highest possible truncated quality that is still an equilibrium, since $\mu_t^{NC} \leq p$ is necessary for certification. Since this is the highest possible supportable value, we call such types of beliefs *cautious beliefs*. Because principals do not bear the cost of certification, they want agents to certify at the fastest possible time that is supportable in an equilibrium. We will denote the principal optimal equilibrium stopping time by τ_p^* .

Note that if $q = p$ then $\mu_t^{NC} \leq p$ for all t , and so principals will want certification by all agents immediately. In such a case, certification acts as a de facto license because any agent that does not certify will be immediately excluded from the market and unable to work at all. With a de facto license, the actual information process being used by the market is irrelevant since agents certify before any learning takes place, and they then stay in the market forever regardless of the signals that are sent.

However, for $q > p$ it could be possible that $\mu_0^{NC} > p$, and so principals may not want agents to certify immediately. In this case, if the agent’s initial reputation is high enough, then the principals will want the agent to work for a while and only certify if its reputation drops too low. Thus the equilibrium will be a delayed certification equilibrium. In this type of equilibrium, the information process itself matters because it affects how quickly the agent’s reputation changes.

6. OPTIMAL CERTIFICATION STANDARDS

6.1 Certification Standards for Socially Optimal Beliefs

Now that we have characterized what can happen in the various types of equilibrium, we will perform some comparative statics in order to analyze the socially optimal certification standards and prices. Doing so will tell us how certification should be implemented depending on what type of equilibrium is expected. In the next few sections we will compute how the certification standards should be set given the various types of equilibrium we named above. In doing so we maintain our assumption that the certification standard must be set higher than the price, $q \geq p$. We start with the socially optimal case because it is the most straightforward. We assume that for any choice of p, q , the market will hold the socially optimal beliefs that were discussed in the previous section. Since these beliefs are an equilibrium, the agent will also play accordingly. With this assumption on the resulting game equilibrium, we find the values of p, q that give the highest social welfare.

Although finding the socially optimal stopping time strategy is challenging in general, characterizing how the price and standard should be set is much simpler. The reason is that, unlike in the principal and agent optimal cases, the beliefs will always work in our favor and we do not need to rig the system to counteract possible “bad beliefs” that expect certification at inefficient reputation levels. Thus we can set the standard and price as low as possible (but still above c) to allow as many good firms to certify as necessary. The next Theorem classifies exactly how the price and standard should be set given socially optimal beliefs.

THEOREM 4:

Let W^ represent the greatest ex ante social welfare that is achievable in any equilibrium when $p = \underline{q} = c + \rho k$. Then if the equilibrium is always a socially optimal equilibrium for any choice of p, \underline{q} , the following holds:*

1. *For $W^* \geq \frac{\mu_0 - c}{\rho}$ it is optimal to set $p = c + k\rho$ and $\underline{q} = p$.*
2. *For $W^* \leq \frac{\mu_0 - c}{\rho}$ it is optimal to set $p = c$ and $\underline{q} = \infty$.*

PROOF: See online appendix. ■

6.2 Socially Optimal Certification Standards for Cautious Beliefs

We now characterize the social welfare maximizing certification standards and prices for the *cautious beliefs* principal optimal beliefs defined in section 5.3.3. Unlike with the socially optimal beliefs, these beliefs are not aimed at maximizing social welfare and so we may need to manipulate the standards and prices in order to increase welfare. We analyze the comparative statics of our model for the parameters \underline{q} and p , assuming that given any choice of the parameters a principal optimal equilibrium will occur. We present a perhaps surprising result: any socially optimal certification standard in a principal optimal equilibrium must act as a de facto license. Thus in these equilibrium, the principals never get to learn from an agent’s work history at all before certification occurs.

The next theorem tells us exactly what the optimal certification standards and prices are as a function of the certification cost. When the certification cost is sufficiently high, certification should not be allowed, and so we end up with our benchmark model. When the certification cost is low, certification should be allowed and implemented as a de facto license where agents certify immediately. In no situation is it optimal to allowed delayed certification, where agents do not certify immediately, but only after working and sending some signals.

THEOREM 5:

If the equilibrium is always a principal optimal equilibrium for any choice of p, \underline{q} , the following holds:

1. *For $k < \frac{\mu_0^c - c}{\rho} - \frac{\mu_0 - c}{\rho(1 - F_0(c + \rho k))}$ it is optimal to have $p = c + k\rho$ and $\underline{q} = p$.*
2. *For $k > \frac{\mu_0^c - c}{\rho} - \frac{\mu_0 - c}{\rho(1 - F_0(c + \rho k))}$ it is optimal to have $p = c$ and $\underline{q} = \infty$.*

PROOF: See online appendix. ■

The reason that delayed certification is never optimal, is if it were implemented then cautious beliefs would require certification by agents very early and often, and in

fact so early and often that certification destroys social welfare. To see this intuitively, note that under cautious beliefs the mean of a non-certifying agent will be exactly equal to the price level. Thus we have the same result as in the benchmark, where we are always kicking out agents who have expected qualities above c . And since $p > c$ in order to allow for certification, the kicking out of the agents will strictly reduce social welfare. Even worse, high quality agents are also forced to pay the cost of certification, destroying additional social value. Since the myopic principals would prefer the agent to certify at inefficient reputation levels, the value of certification gets completely undercut.

Because certification should act as a de facto license under these types of beliefs, learning through the work history and certification are interacting trivially. Agents that don't certify can never work, and agents that do certify never exit, so it doesn't matter what type of information the work history is sending. For any information process, the ex ante social welfare will be the same. When certification costs are low enough, de facto licenses will provide a higher social welfare than the no certification benchmark and thus they should be implemented. In fact, for $k = 0$ we can achieve the first best outcome through setting $\underline{q} = p = c$, which results in all socially beneficial agents working forever, and all other agents exiting immediately. And as $k \rightarrow 0$, we can asymptotically achieve first best by setting the price and certification standard lower and lower. However, if certification costs are too high, then too much social welfare is lost through certification itself, and it is thus better to just have no certification at all.

6.3 Socially Optimal Certification Standards for Optimistic Beliefs

Now we will analyze the agent optimal equilibrium that features the optimistic beliefs mentioned above. Since optimistic beliefs never result in de facto licenses if $\mu_0 > p$, the social welfare generated by these beliefs can be heavily dependent on the specific information process that the market has access to. In order for certification to occur, the price needs to be high enough to compensate for the cost of certification, so we need $p \geq c + k\rho$ or the result is no certification. And given that agents with qualities higher than $c + k\rho$ contribute positively by certifying instead of exiting, the certification standard should be set equal to p . But the inequality $p \geq c + k\rho$ need not be binding, because it may be optimal for the price to be raised so that certification occurs sooner. For instance, if the certification cost k is very low, then it is optimal to have p higher so that we can get certification earlier.

Also note that a de facto license can be implemented if the price is set at $p \geq \mu_0$, since then principals would never buy from non-certifying agents. The optimal way to implement a de facto license is to set $p = \underline{q} = \max(c + \rho k, \mu_0)$, for similar reasons as in the principal optimal case. But a de facto license would result in all agents with qualities less than μ_0 but greater than c getting forced out of the market. Thus it may be better not to have a de facto license, but delayed certification instead.

In general then, for k sufficiently low it will be optimal to have $\underline{q} = p \in [c + \rho k, \mu_0]$, where the exact value depends on the information process itself. This results in either delayed certification if $p < \mu_0$ or a de facto license if $p \geq \mu_0$. If k is too high, then again we should set $p = c$ and $\underline{q} = \infty$ to ensure that no certification occurs. The exact cutoff value for k , as well as the exact values of the price and certification standard, will depend on the information process because that will influence the social welfare of any delayed certification scheme.

Given that delayed certification can be desirable in an agent optimal equilibrium, one may think that the social welfare under optimistic beliefs will always be higher than the social welfare under cautious beliefs, where delayed certification could never be optimal. However, optimistic beliefs also have a downside – de facto licenses are now much harder to implement because the price needs to be set at a very high level to do so. Under cautious beliefs, we could asymptotically get first best as $k \rightarrow 0$ by implementing a de facto license. But that is not the case with optimistic beliefs, as the next theorem shows. With optimistic beliefs, a de facto license cannot achieve first best, and delayed certification cannot achieve first best either.³ For any value of k , we define $W_a(k, p, \underline{q}, \mathcal{L})$ as the *ex ante* social welfare in the agent optimal equilibrium.

THEOREM 6:

As $k \rightarrow 0$, in any agent optimal equilibrium the value of the ex ante social welfare is strictly bounded away from the first best value, $\limsup_{k \rightarrow 0} W_a(k, p, \underline{q}, \mathcal{L}) < \int_c^\infty \frac{q-c}{\rho} f_0(q)$.

PROOF: See online appendix. ■

This theorem shows that given any information process, the social welfare is bounded away from the optimal as costs become low. However, we note that if the information process itself become very informative, then the optimal social welfare can be asymptotically achieved. Define a *fully revealing* information process to be such that almost surely $\tau_a^* < \infty$ for all agents with $q < c$. This means that bad agents will for sure be forced to certify at some point, and thus will always get kicked out of the market in the long run. Such a property holds as long as the information process allows bad agents to be distinguished from good agents over time. For fully revealing information processes, faster information revelation can increase social welfare. Formally, we say that an information process \mathcal{L}' is *faster* than another process \mathcal{L} if for some $n > 1$, $R'(t; q) = R(nt; q)$ for all t, q . The parameter n tells us how much faster one process is at sending information than another. The reason that faster information is helpful is because low quality agents will get kicked out of the market very quickly while high quality agents can still stay in⁴. As the following proposition shows, first best can be asymptotically achieved if the certification costs become very low and the information arrives at a very fast speed. Thus information revelation and certification can act in a complimentary fashion, with more information increasing the welfare provided by certification.

PROPOSITION 4:

Suppose that \mathcal{L} is a fully revealing information process, and let $\{\mathcal{L}'_n\}$ be a sequence of information processes that are faster than \mathcal{L} by the factor n . Then we have

$$\lim_{n \rightarrow \infty} \limsup_{k \rightarrow 0} W_a(k, p, \underline{q}, \mathcal{L}'_n) = \int_c^\infty \frac{q-c}{\rho} f_0(q)$$

PROOF: See online appendix. ■

³ When $k = 0$, agents do not mind certifying at an expected quality above the price, and so first best can be achieved in an agent optimal equilibrium.

⁴ Note that without certification faster information cannot increase welfare as shown by Theorem 1.

On the other hand, if certification costs are high and a delayed certification scheme were attempted, faster information revelation could actually lower social welfare. When costs are high, the price and certification standard need to be high as well to get agents to certify. But setting high certification standards will make certification kick out agents less efficiently than before because exiting agents will have higher quality levels. A faster information process compounds this problem by getting to the kick out point faster, thus lowering social welfare. In this case it may be better not to implement certification.

7. CONCLUSION

This paper analyzed how two separate avenues of information, certification and work history, can interact to affect learning about the quality of an agent. We showed that the information provided by the work history alone cannot raise social surplus, making certification necessary even when the market has access to other informative signals. With certification, all equilibria will feature type independent certification strategies, and the various equilibria that maximize principal, agent, and social welfare can be ordered according to the agent's certification stopping time strategies. Principals will prefer the agent to certify at earliest time, and thus having the agent send information before certifying is inefficient. With cautious beliefs, only a de facto license type of certification can be socially beneficial. On the other hand, in the agent optimal case the agent will delay certification for as long as possible. This is harmful when certification costs are low and it is socially beneficial to have the agent certify quickly. However, in this case faster information revelation can increase social welfare, creating a complementarity between the reputational forces and certification.

There are several possible extensions for future research that could have interesting implications. One important case is to allow for variable prices that depend partially or even wholly on the agent's reputation. Such a change complicate the exit decision for the agent, because it may wish to keep working even if the price falls below its own reservation value. In extreme cases such as Bar-Isaac (2003), higher quality agents may never exit, which means that there is no loss in efficiency in the long run. Thus, certification would only be useful to flush bad agents out of the market sooner, and so certification would be less valuable than in the current model.

Another interesting extension is to allow the certification standards to be dynamic. For instance, a certifier may wish to not set a standard at all in the beginning, and only introduce a standard if the agent's reputation falls too low. If this is possible, then delayed certification could now be socially beneficial in a principal optimal equilibrium. Allowing a certifier to not set a standard in the beginning means that later on the certifier can choose a lower standard than with delayed certification, where the standard had to be high enough such that agents did not need to certify at the start. A lower standard means that more agents are able to certify, and thus welfare can be higher than with delayed certification. The problem would then entail finding the optimal time for the certifier to switch standards, and what standards to set. Such flexibility would of course allow the certifier to achieve better results than in our model.

ELECTRONIC APPENDIX

The electronic appendix for this article can be accessed in the ACM Digital Library.

ACKNOWLEDGMENTS

The authors are deeply indebted to Moritz Meyer-ter-Vehn, Simon Board, Oren Louidor, and Bill Zame for their guidance. This paper also benefitted from valuable discussions with Jeff Ely, Crysanthos Dellarocas, Sanjeev Goyal, Johannes Hörner, David Parkes, and Yangbo Song.

REFERENCES

- Atkeson, A., Hellwig, C., & Ordonez, G. (2012). Optimal regulation in the presence of reputation concerns (No. w17898). *National Bureau of Economic Research*.
- Bar-Isaac, H. (2003). Reputation and Survival: Learning in a Dynamic Outputting Model. *The Review of Economic Studies*, 231-251.
- Bar-Isaac, H., & Tadelis, S. (2008). Seller Reputation. *Foundations and Trends in Microeconomics*.
- Board, S., & Meyer-ter-Vehn, M. (2010). Reputation For Quality. *Mimeo*.
- Chung, K.-S., & Eso, P. (2007). Outputting with Career Concerns. *Mimeo*.
- Bonatti, A. (2011). Menu Pricing and Learning. *American Economic Journal: Microeconomics*, 3(3), 124-163.
- Bolton, P., Freixas, X., & Shapiro, J. (2009). The credit ratings game (No. w14712). National Bureau of Economic Research.
- Ely, J., Fudenberg, D., & Levine, D. K. (2008). When is reputation bad?. *Games and Economic Behavior*, 63(2), 498-526.
- Ely, J. C., & Välimäki, J. (2003). Bad reputation. *The Quarterly Journal of Economics*, 118(3), 785-814.
- Gill, D., & Sgroi, D. (2012). The optimal choice of pre-launch reviewer. *Journal of Economic Theory*, 147(3), 1247-1260.
- Farhi, E., Lerner, J. & Tirole, J. (2013). Fear of Rejection? Tiered Certification and Transparency. *mimeo*
- Holmström, B. (1999). Managerial incentive problems: A dynamic perspective. *The Review of Economic Studies*, 66(1), 169-182.
- Jin, G. Z., & Leslie, P. (2003). The effect of information on product quality: Evidence from restaurant hygiene grade cards. *The Quarterly Journal of Economics*, 118(2), 409-451.
- Jin, G., & Leslie, P. (2009). Reputational Incentives for Restaurant Hygiene. *American Economic Journal: Microeconomics*, 1(1), 237-267.
- Keller, G. (2011). Brownian Motion and Normally Distributed Beliefs.
- Lizzeri, A. (1999). Information revelation and certification intermediaries. *The RAND Journal of Economics*, 214-231.
- Leland, H. E. (1979). Quacks, lemons, and licensing: A theory of minimum quality standards. *The Journal of Political Economy*, 87(6), 1328-1346.
- Luehrman, T. A. (1998). Strategy as a portfolio of real options. *Harvard business review*, 76, 89-101.
- Mathis, J., McAndrews, J., & Rochet, J. C. (2009). Rating the raters: are reputation concerns powerful enough to discipline rating agencies?. *Journal of Monetary Economics*, 56(5), 657-674.
- Morrison, A., & White, L. (2004). Financial liberalisation and capital regulation in open economies.
- Mysliveček, J. (2008). Comparing Certification and Self-regulation. *CERGE-EI Working Paper Series*, (361).
- Ottaviani, M., & Sørensen, P. (2006). Reputational cheap talk. *RAND Journal of Economics*, 37(1), 155-175.
- Panzar, J. C., & Savage, I. (2011). Does a Minimum Quality Standard Always Reduce the Price of High Quality Products?. *The BE Journal of Economic Analysis & Policy*, 11(1).
- Reiss, D. J. (2005). Subprime Standardization: How Rating Agencies Allow Predatory Lending to Flourish in the Secondary Mortgage Market. *bepress Legal Series*, 764.
- Salminen, P. (1988). On the first hitting time and the last exit time for a Brownian motion to/from a moving boundary. *Advances in applied probability*, 411-426.
- Siegmund, D. (1986). Boundary Crossing Probabilities and Statistical Applications. *The Annals of Statistics*, 40(2), 361-404.
- Shapiro, C. (1983). Premiums for high quality products as returns to reputations. *The Quarterly Journal of Economics*, 98(4), 659-679.
- Stolper, A. (2009). Regulation of credit rating agencies. *Journal of Banking & Finance*, 33(7), 1266-1273.
- Strausz, R. (2005). Honest certification and the threat of capture. *International Journal of Industrial Organization*, 23(1), 45-62.
- Stahl, K., & Strausz, R. (2011). Who should pay for certification?.
- Wang, L., & Pötzelberger, K. (1997). Boundary crossing probability for Brownian motion and general boundaries. *Journal of Applied Probability*, 54-65.
- White, L. J. (2010). Markets: The credit rating agencies. *The Journal of Economic Perspectives*, 24(2), 211-226.
- Xiao, M. (2010). Is quality accreditation effective? Evidence from the childcare market. *International Journal of Industrial Organization*, 28(6), 708-721.

Online Appendix to: A Dynamic Model of Certification and Reputation

MIHAELA VAN DER SCHAAR, University of California, Los Angeles
SIMPSON ZHANG, University of California, Los Angeles

A. DISCUSSION OF ADMISSIBLE INFORMATION PROCESSES:

This section will discuss the significance of our admissibility requirement and give examples of information processes that are admissible. Note that a restriction on \mathcal{L} entails a restriction on the types of stochastic signals $R(t; q)$ that can be allowed. One way to satisfy our continuity requirement is for the market to only be able to infer a little bit of information from $R(t)$ at any point in time. For instance suppose that for any q , $R(t; q)$ is the diffusion $dR(t; q) = qdt + \sigma_r dZ(t)$, with drift q and variance σ_r^2 (and precision τ_r), and with $Z(t)$ a standard Brownian motion. This type of process would result if the market observes each unit of output with some normal noise, and we note that this is the continuous time non-moral hazard version of Holmström (1999). In this case, the agent's reputation would evolve continuously in time, and so our admissibility requirement is satisfied. Even more general diffusion processes are admissible, such as if $dR(t; q) = f(q, \mathcal{H}_t)dt + g(\mathcal{H}_t)dZ(t)$ so that the drift and variance do not have independent or stationary increments. However, stochastic processes such as $dR(t; q) = qdt + q\sigma_r dZ(t)$ are not admissible because the variance would be learned immediately through Bayes' rule, causing the means to jump downwards with positive probability.

Stochastic processes with jumps can also be considered. For instance, suppose we have a Poisson process with an arrival rate $\lambda(q)$, where λ is increasing in q . This is a good news Poisson process, where an arrival indicates a positive event like being mentioned in an article or winning an award. Without a signal arrival, the means drift slowly and continuously downwards, and at a signal arrival, the expected mean and all truncated means would jump upwards. Thus our admissibility condition is satisfied. More generally the arrival rate could depend on the history as well, $\lambda(q, \mathcal{H}_t)$, as long as it was still increasing in q at all t . And if the stochastic process was a combination of diffusions and good news Poisson processes, the result could still hold. However, if the stochastic process was a bad news Poisson process, where $\lambda(q)$ was decreasing in q , then admissibility would be violated because at a signal arrival the means would jump downwards. Thus jumps stochastic processes are admissible as long as they indicate good news, but not if they indicate bad news.

B. PROOF OF THEOREM 1:

First note that if the price is less than c , the agent would refuse to accept any offers, and so the social welfare is equal to 0 regardless of the information process. Thus we only need to consider prices $p \geq c$.

Note that under any blind process, the agent's expected quality is never updated, and so since $\mu_o > p$ the agent will never stop working. The *ex ante* social surplus can thus be calculated as $\frac{\mu_o - c}{\rho}$. For a general information process the market continues to

hire any agent until its expected quality drops below p , and for an admissible process this happens at the first time that $\mu_t = p$. We can write out the *ex ante* expected social welfare for any information process as

$$\begin{aligned} W(p, \mathcal{L}) &= \int_{-\infty}^{\infty} \left[\int_0^{\infty} e^{-\rho t} (q - c) dt \right] f_0(q) dq - E \left[\int_{t^*}^{\infty} e^{-\rho t} (q - c) dt \right] \\ &= \int_{-\infty}^{\infty} \left[\int_0^{\infty} e^{-\rho t} (q - c) dt \right] f_0(q) dq - E_{t^*} \left[\int_{t^*}^{\infty} E_q [e^{-\rho t} (q - c) | t \geq t^*] dt \right] \\ &= \frac{\mu_0 - c}{\rho} - \frac{p - c}{\rho} E_{t^*} [e^{-\rho t^*}] \leq \frac{\mu_0 - c}{\rho} \end{aligned}$$

■

C. PROOF OF THEOREM 2:

Fix any arbitrary market strategy beliefs at any time t' , and let the expected quality for non-certifying agents be given by $\mu_{t'}^{NC}$. There are two cases: $\mu_{t'}^{NC} \geq p$ and $\mu_{t'}^{NC} < p$, and we show that either case must lead to all types of agents certifying or no types of agents certifying.

If $\mu_{t'}^{NC} > p$, then agents can still work even after not certifying. The payoff of an agent that chooses to certify is given by $\frac{p-c}{\rho} - k$. Now consider the alternate strategy of waiting until the time $\hat{t} \equiv \inf\{t | \mu_t^{NC} \leq p\}$ and then certifying. This strategy gives a payoff of

$$\int_0^{\infty} e^{-\rho t} (p - c) dt - (e^{-\rho \hat{t}}) * \left(\frac{p - c}{\rho} - k \right)$$

The alternate strategy gives a payoff higher than certifying immediately by $(1 - e^{-\rho \hat{t}}) * k$. So certifying at time t' is not optimal, and with these beliefs all types of agents would choose not to certify. Now suppose that the market were to believe instead that no agents certify. In this case $\mu_{t'}^{NC} > p$ as well and by the previous argument no agents would want to certify. Therefore it is an equilibrium for agents not to certify if $\mu_{t'} > p$.

Next consider the case where $\mu_{t'}^{NC} \leq p$. An agent that does not certify will not be able to work at time t' , and so receives a maximum payoff of

$$(1 - \rho dt) * \left(\frac{p - c}{\rho} - k \right)$$

Given that no outputs are sent, the mean at a later time cannot be greater than p unless the agent were to certify. If the agent instead chooses to certify at time t' , it would get a payoff of

$$\frac{p - c}{\rho} - k > (1 - \rho dt) * \left(\frac{p - c}{\rho} - k \right)$$

Thus certifying in the current time step would increase the payoff by $\rho dt \left(\frac{p-c}{\rho} - k \right)$. Therefore all types of agents would choose to certify given these market beliefs. Now suppose that the market believed that all agents with qualities in $[q, \infty)$ choose certification. Given these market beliefs, $\mu_{t'}^{NC} \leq p$ and agents that don't certify will not be able to work, so agents will find it optimal to certify. Therefore it is an equilibrium for all agents with qualities in $[q, \infty)$ to choose certification given these market beliefs.

■

D. PROOF OF PROPOSITION 1:

Consider any equilibrium that requires the agent to certify at a time t' where $\mu_{t'} > p$. Now let's compare the payoffs to the agent against the equilibrium where the agent only certifies when $\mu_t \leq p$ for all t . In the second equilibrium, at time t' the agent would instead delay certification until the time $\hat{t} \equiv \inf\{t | \mu_t \leq p\}$. Since the payoff from certification is the same regardless, the agent would be able to improve its payoff by the amount $(1 - e^{-p\hat{t}}) * k$. Thus no equilibrium that requires that agent to certify at a $\mu_t > p$ can be optimal. ■

E. PROOF OF PROPOSITION 2:

Suppose for the sake of contradiction that (1) does not hold at the certification time τ^* . Then consider the social welfare generated if the agent could keep working up until the time $\tau^* + dt$ no matter what signals are sent, and the agent then certifies at $\tau^* + dt$ regardless of the signals, staying in the market only if it gets certified. We call this the *no stop* scheme, and we calculate the difference in the social welfare generated by this scheme and by certifying at time τ^* . Note that all the flow payoffs after $\tau^* + dt$ are the same, so we only have to consider the difference in the flow payoffs up to that point. With the no stop scheme, we have the flow payoff at any time t given by

$$E_t[q - c] = \int_{-\infty}^{\infty} (q - c) f_{\tau^*}(q) dq$$

With certification this would become

$$E_{\tau^*}[q - c] = \left(1 - F_{\tau^*}(\underline{q})\right) \int_{\underline{q}}^{\infty} (q - c) f_t(q | q > \underline{q}) dq = \int_{\underline{q}}^{\infty} (q - c) f_{\tau^*}(q) dq$$

Thus the difference in the two expectations is given by

$$\int_{-\infty}^{\underline{q}} (q - c) f_{\tau^*}(q) dq$$

Certification is preferred only if the difference in expectations plus the flow cost of certification is less than zero:

$$\int_{-\infty}^c (q - c) f_{\tau^*}(q) dq + \int_c^{\underline{q}} (q - c) f_{\tau^*}(q) dq + \rho k \left(1 - F_{\tau^*}(\underline{q})\right) \leq 0$$

Or equivalently $(\mu_t^{NC} - c - \rho k) F_t(\underline{q}) \leq -\rho k$.

This results in equation (1) in the Proposition. Thus we have proven the result if almost surely the agent can keep working for some positive length of time dt . However, this may not be true in general, because there may be a positive probability that the mean drops below the price for any finite dt . In this case, the agent stops working before time $\tau^* + dt$ so the above expression cannot hold in equilibrium.

However, we will consider instead the following revised strategy that can be an equilibrium, and we show that it still provides a higher welfare than stopping at time τ^* if (1) does not hold at τ^* . Suppose instead that the agent keeps working until the first $t > \tau^*$ such that (1) holds, or until the first $t > \tau^*$ such that $\mu_t \leq p$, or until $\tau^* + dt$, whichever of the three comes first, and then certifies at the respective time. For now we will assume that the agent can always certify at time $\tau^* + dt$, which is the case if $\underline{q} = p$ or almost surely $\mu_{\tau^*}^{NC} \leq p$. We will consider the other case later. For the above hitting times t , let $\Delta t_i = t_i - \tau^*$ for case i respectively. If the first case occurs, the welfare will be less than under certification at time τ^* , and this loss is given by $\rho \Delta t_1 \left((\mu_{t_1}^{NC} - c - \rho k) F_{t_1}(\underline{q}) + \rho k \right)$. If case 2 occurs, then the welfare will be strictly

greater than with certification at τ^* . The amount of increase is $\rho\Delta t_2 \left((\mu_{t_2}^{NC} - c - \rho k)F_{t_2}(\underline{q}) + \rho k \right)$. If case 3 occurs, the welfare will be the same as in the no stop scheme and greater than certifying at τ^* by $\rho dt \left((\mu_{\tau^*+dt}^{NC} - c - \rho k)F_{\tau^*+dt}(\underline{q}) + \rho k \right)$. So the total difference in welfare is given by

$$\Pr(1) \rho E \left[\Delta t_1 \left((\mu_{t_1}^{NC} - c - \rho k)F_{t_1}(\underline{q}) + \rho k \right) \right] + \Pr(2) \rho E \left[\Delta t_2 \left((\mu_{t_2}^{NC} - c - \rho k)F_{t_2}(\underline{q}) + \rho k \right) \right] + \Pr(3) \rho E \left[dt \left((\mu_{\tau^*+dt}^{NC} - c - \rho k)F_{\tau^*+dt}(\underline{q}) + \rho k \right) \right]$$

The only case where this scheme does not provide welfare greater than certifying at τ^* is case 1. Now the term $(\mu_t^{NC} - c - \rho k)F_t(\underline{q}) + \rho k$ is a martingale, so at any time t that case 1 occurs, cases 2 and 3 must also be occurring over that time interval with a greater relative weight. And any time that case 2 occurs, cases 2 and 3 must have a greater relative weight than case 1 over that time interval. After case 1 or case 2 occurs, the flow payoffs are the same in the two cases. This means that the term $\Pr(1) \rho E \left[\Delta t_1 \left((\mu_t^{NC} - c - \rho k)F_t(\underline{q}) + \rho k \right) \right]$ must be less in magnitude than the other two expectations. And after case 1 occurs, the flow payoffs are the same in the two cases. So the overall expected welfare is higher than certifying at time τ^* .

If the agent cannot certify at time $\tau^* + dt$, then we note that $\mu_{\tau^*+dt} > p$ and so equation (1) will not be satisfied. In this case, we note that the welfare from not having the agent certify must be higher than the welfare from having the agent only keep working if it does certify, which is what happens in the *no stop* scheme. This is shown by iterating the above argument at the time $\tau^* + dt$ using the same procedure. Through every step of the iteration, the flow payoffs will be higher from not certifying. ■

F. PROOF OF PROPOSITION 3:

The short run time t principal's utility is given by $\mu_t - p = E_t[q - p]$, so the principal prefers certification if and only if this expectation after certification is higher than the expectation before certification. Without certification, we have

$$E_t[q - p] = \int_{-\infty}^{\infty} (q - p)f_t(q) dq$$

With certification this would become

$$E_t[q - p] = \left(1 - F_t^-(\underline{q})\right) \int_{\underline{q}}^{\infty} (q - p)f_t(q|q \geq \underline{q}) dq = \int_{\underline{q}}^{\infty} (q - p)f_t(q) dq$$

Thus the difference in the two expectations is given by

$$\int_{-\infty}^{\underline{q}} (q - p)f_t(q) dq$$

Certification is preferred if and only if the above term is less than zero.

$$\int_{-\infty}^p (q - p)f_t(q) dq + \int_p^{\underline{q}} (q - p)f_t(q) dq \leq 0$$

This means that the benefit of removing bad agents (qualities below p) from the market outweighs the costs of removing the good agents (qualities above p). Or equivalently:

$$(\mu_t^{NC} - p)F_t(\underline{q}) \leq 0$$

This results in equation (2) in the Proposition. Since at each time t , that time t principal wishes for the agent to certify if and only if this equation holds, the resulting

equilibrium will feature the agent certifying at the first moment that this equation holds.

■

G. PROOF OF THEOREM 4:

First note that if $p < c + \rho k$ there can be no certification in equilibrium. The optimal way to implement no certification is to set $p = c$, $\underline{q} = \infty$ from Theorem 1. This results in a social welfare of $\frac{\mu_0 - c}{\rho}$ for any admissible information process. Next, suppose that we wish to allow certification in equilibrium. Thus we need to set $p \geq c + \rho k$. We show that for any $p > c + \rho k$ and any $\underline{q} \geq p$, we can achieve at least as high of a welfare by setting $\underline{q} = p = c + \rho k$. Given the first set of parameters, denote the socially optimal certification stopping time by τ_s^* . But under the second set of parameters, τ_s^* can also be implemented because $\mu_t^{NC} \leq p$ and $\mu_t > p$ must hold at τ_s^* . In addition, once implemented the social welfare provided by certification will be at least as high, because all agents with qualities $q > c + \rho k$ contribute positively to social welfare by certifying instead of exiting. Thus the social welfare with $\underline{q} = p = c + \rho k$ must be at least as high as with any other standard.

■

H. PROOF OF THEOREM 5:

This proof will proceed in several steps. Note that certification can be broken up into three possible types: immediate certification at $t = 0$, delayed certification that takes place at some $t > 0$, and no certification for all times. Which type of certification results will depend on the specific values of \underline{q} and p . We prove that under cautious principal beliefs, delayed certification is never optimal. Then, we characterize the social welfare generated by immediate and no certification. We prove that immediate certification and no certification can both be optimal depending on how high the certification cost is.

First we show that implementing delayed certification is never socially optimal. If the price is set lower than $c + \rho * k$ then certification can never occur because agents would prefer to exit than certify. Thus for any type of certification to be implemented, we must have $p \geq c + \rho k$. Next, note that under cautious principal beliefs, the certification standard \underline{q} must be set high enough such that $\mu_0^{NC} \geq p$ because otherwise agents would be expected, and thus forced, to certify immediately in a principal optimal equilibrium by Proposition 3. In particular this requires that $\underline{q} > p$.

Now we analyze the social welfare generated by any delayed certification scheme, and we show that the welfare is strictly less than under no certification. Let $t^{**} = \inf(t | \mu_t^{NC} \leq p)$. In a principal optimal equilibrium, t^{**} is the time at which certification would occur. Note that admissibility implies that at t^{**} , the truncated expected mean $\mu_{t^{**}}^{NC} = p$. For any p and \underline{q} that satisfy the above conditions, we can compute the resulting social welfare as:

$$W(p) = \int_{-\infty}^{\infty} \left[\int_0^{\infty} e^{-\rho t} (q - c) dt \right] f_0(q) dq - E_{t^{**}} \left[k e^{-\rho t^{**}} \left(1 - F_{t^{**}}(\underline{q}) \right) + F_{t^{**}}(\underline{q}) \int_{t^{**}}^{\infty} e^{-\rho t} (\mu_{t^{**}}^{NC} - c) dt \right]$$

$$= \frac{\mu_0 - c}{\rho} - \left(k E_{t^{**}} \left[e^{-\rho t^{**}} \left(1 - F_{t^{**}}(\underline{q}) \right) \right] + \frac{p - c}{\rho} E_{t^{**}} \left[e^{-\rho t^{**}} F_{t^{**}}(\underline{q}) \right] \right) \leq \frac{\mu_0 - c}{\rho}$$

This proves that the social welfare of any delayed certification scheme is bounded above by setting $p = c$ and $\underline{q} = \infty$, which results in no certification. From the blind boundedness theorem, we know that the payoff of such a scheme is exactly $\frac{\mu_0 - c}{\rho}$.

Now, fixing a \underline{q} , the welfare provided by a de facto license is given by the expression $\int_{\underline{q}}^{\infty} \left(\frac{q - c}{\rho} - k \right) f_0(q) dq$. Given a k , the optimal certification standard is $\underline{q} = c + \rho k$. The reason is that any agent that certifies will give a social welfare of $\frac{q - c}{\rho} - k$, and the above quality is where this value is equal to zero. Any agent with a quality higher than this amount contributes positively to welfare by certifying. Since we require that $p \leq \underline{q}$ and we need $p \geq c + \rho k$ for certification to occur, this implies that we need to set $p = c + \rho k$ in order to implement immediate certification. Thus the highest *ex ante* surplus generated by any immediate certification scheme is $\int_{c + \rho k}^{\infty} \left(\frac{q - c}{\rho} - k \right) f_0(q) dq = \left(\frac{\mu_0 - c}{\rho} - k \right) (1 - F_0(c + \rho k))$.

Thus to see whether immediate certification is better, or whether no certification is better, we need to see which of the two surpluses is higher. This depends on the value of k , and specifically the cutoff value will be given by $k^* = \frac{\mu_0 - c}{\rho} - \frac{\mu_0 - c}{\rho(1 - F_0(c + \rho k))}$.

■

I. PROOF OF THEOREM 6:

First note that if either no certification or a de facto license is implemented, the social welfare will always be bounded away from the social optimal. With no certification, the welfare always equals the benchmark welfare for any k , and with a de facto license, the price and standard have to be set at least at μ_0 . The social welfare of the standard is thus equal to $\int_{\mu_0}^{\infty} \left(\frac{q - c}{\rho} - k \right) f_0(q) dq$, which is bounded away from the first best welfare, $\int_c^{\infty} \left(\frac{q - c}{\rho} \right) f_0(q) dq$, for any value of k .

Thus in order to get asymptotic efficiency as $k \rightarrow 0$, we need to do it through a delayed certification scheme. We now show that the social welfare of any delayed certification scheme will also be bounded away from first best. First note that if the standard $\underline{q} \rightarrow c$, then first best cannot be achieved. The reason is that before certification occurs, we are losing welfare from letting bad agents work, and after certification occurs we also lose welfare since not all good agents are working. Then assume that $p, \underline{q} \rightarrow c$. Fix any path of the expected mean for the agent. Let $t_c^* = \inf\{t | \mu_t \leq p; p = c\}$ be the stopping time of this path in the limit as the price approaches c . In order to achieve first best as $k \rightarrow 0$, we must have $t_c^* = 0$ or else bad firms will be selling for some stretch of time. But for any admissible information process this is impossible since $\mu_0 > p$. Thus t_c^* is strictly above 0, and so delayed certification cannot achieve first best.

■

J. PROOF OF PROPOSITION 4:

From the proof of Theorem 6, we see that immediate certification and no certification cannot asymptotically achieve first best, because the informativeness of the reputational mechanism does not affect social welfare in either case. So we must show that the social welfare of a delayed certification scheme approaches first best. We wish to show that as the informativeness becomes very high, $t_c^*(q) \rightarrow 0 \forall q < c$ because doing so means that all agents who have socially inefficient qualities will be kicked out extremely quickly, and the good agents will be able to stay in forever (perhaps paying the certification cost that asymptotically approaches 0). Since the process is fully revealing, almost surely $t_c^*(q) < \infty$ for agents with quality $q < c$. Then as n gets large, agents will be kicked out at time $\frac{t_c^*(q)}{n}$ instead, which approaches 0 for all finite $t_c^*(q)$. Thus $t_c^*(q) \rightarrow 0 \forall q < c$ and so delayed certification asymptotically achieves the first best outcome.

■