SELF-CONFIDENCE AND SURVIVAL

Heski Bar-Isaac*
Department of Economics and Financial Markets Group
London School of Economics and Political Science

Contents:

	The Suntory Centre
	Suntory and Toyota International Centres for
	Economics and Related Disciplines
	London School of Economics and Political Science
Discussion Paper	Houghton Street
No. TE/01/428	London WC2A 2AE
October 2001	Tel.: 020-7955 6698

---

* I am very grateful to Sudipto Bhattacharya, Antoine Faure-Grimaud, Drew Fudenberg, Christian Hellwig, Michihiro Kandori, Godfrey Keller, Gilat Levy, Margaret Meyer, Marco Ottaviani, Andrea Prat, Michele Piccione, Steve Tadelis, Juuso Valimaki, and participants at a number of seminars and conferences, as well as many other insightful colleagues at the LSE's Economics Department and at the Financial Markets Group for their seemingly limitless patience, useful comments and advice. In particular, I thank George Malaith who inspired the extension in Section 6 and Margaret Bray and Leonardo Felli for countless and often critical advice, suggestions and support. Financial support from the Economic and Social Research Council (award no. R42200034040) is also gratefully acknowledged. All errors are my own.

Abstract

We consider the impact of history on the survival of a monopolist selling single units in discrete time periods, whose quality is learned slowly. If the seller learns her own quality at the same rate as customers, a sufficiently bad run of luck could induce her to stop selling. When she knows her quality, a good seller never stops selling. Furthermore, a seller with positive, though imperfect, information sells for the same number of periods whether her information is private or public. We further consider the robustness of the central result when the seller's opportunities for strategic behaviour are limited.

Keywords: Reputation, signalling, learning, one-armed bandit, monopolist, private information, public information.

JEL Nos.:C61, C73, D42, D83, L15.

© by the author. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission, provided that full credit, including © notice, is given to the source.

Contact address: Heski Bar-Isaac, Department of Economics, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, email: h.bar-isaac@lse.ac.uk/

1 Introduction

Reputation, reputation, reputation . . . Reputation is an idle and most false imposition: oft got without merit and lost without deserving

Othello, Act II, Scene 3

In the scene from which the above quote is taken, Cassio complains that he has lost his reputation—the "immortal part" of himself—and is comforted by Iago. Iago essentially tells Cassio that reputation does not have any meaning: good people can have bad reputations and bad people can have good reputations. However, in a typical example of Shakespearean irony, by the end of the play Iago is publicly revealed as the villain that he is and Cassio, who in subsequent acts of the play tries to repair his reputation, ends up as governor of Cyprus. This paper can be read as contrasting the stylised view of Iago, that good people can be unlucky and that bad people can be lucky, with Shakespeare's implicit view, that in the long run the truth will out. Specifically, we find support for the latter view under certain circumstances; though Iago's viewpoint might prevail in the short-term, self-confidence signals quality and may partially restore an inappropriately poor reputation.

We consider an optimal stopping problem in which, in every period, a monopolist provider chooses whether to sell a single unit of a service or good, incurring a fixed production cost each time that she does, or to stop trading. Initially the quality of the service is unknown to potential buyers but each purchase yields additional information about the true quality. We consider the problem, under different assumptions concerning what the seller knows about her quality, specifically we begin by examining feasible outcomes if she knows no more than buyers, and if she knows her quality perfectly.

In the case where the seller has no additional information to buyers we show that in this strategic setting, an analogous result to Rothschild (1974) emerges, that is that even though the seller may have a good product, she might stop selling. This then offers support to "Iago's view" as characterised above—that a good product can end up with a bad reputation.

However, introducing private information leads to a contrary result—when the seller knows the quality of her product with certainty, then a good seller (one with a good product) sells in every period—lending some support to "Shakespeare's view". The mechanism, which we term "self-confidence" supporting this result has some intuitive appeal. A seller with a poor reputation (that is one whose product buyers believes is quite likely to be bad) would be more willing to sell, even at a low price, if she were good than if she were bad. This is because a good seller would know that her reputation would most likely improve following future sales, and so would be relatively more likely to accept a loss in the current period. As a bad seller would be relatively more likely to cease trading if required to sell at a loss, it follows that simply being prepared to sell at a "reasonable" price and exhibiting confidence that her reputation will improve, can act as a signal of quality to buyers in the current period. This consideration changes the notion of what price might be reasonable in the current period, since now buyers' beliefs depend not only on previous experience of the product but also on the observation that the seller is willing to continue selling despite losses. In particular, this consideration leads to a higher reasonable price, albeit one at which the seller would still expect a short-term, or myopic, loss.[1]

[1]The experience of bankrupts in Silicon Valley might provide anecdotal evidence of this mechanism in oper-

Thus self-confidence, in the sense of accepting short-term losses for the sake of future expected gains, can signal information and limit the losses that the seller need incur in any period. Indeed, in the model these losses are limited to such an extent that a good seller will always wish to make a sale. Moreover, this intuition and the qualitative result do not rely on the discount factor—whenever a seller puts any weight on the future, a good seller who expects a better future than a bad one prefers to continue selling. Note, however, that this mechanism is tempered to the extent that a bad seller could mimic the well-founded self-confidence of a good one; formally this implies semi-separation so that when the reputation is low enough, a good seller would always sell and a bad seller would stop selling with some probability and continue selling with the complementary probability. Indeed, this consideration suggests that no fully separating equilibrium is possible.[2]

We then proceed to test the appeal and robustness of the signalling equilibrium characterised. First, in an extension of the model, where the seller's information is not perfect but she has more information than buyers, we show that self-confidence leads to an outcome whereby a seller with positive information continues selling for just as long as she would have done if her information had been public rather than private; thus a good private signal and a good public signal are equivalent for the survival of a seller. Second, although the introduction of asymmetric information leads to multiple equilibria, which we discuss at some length, we argue that the self-confidence equilibrium we identify is the most appealing one in a well-defined sense—specifically, natural restrictions on beliefs lead to uniqueness of this equilibrium as a limiting case in a richer environment. We explore varying the seller's opportunity to behave strategically by supposing that she can only decide to continue or quit trading with some probability, so that an observation that she continues trading need not reflect a deliberate decision, for example it may be that she makes a mistake. In this scenario we derive qualitatively similar equilibria and show that as the seller's probability of having an opportunity to make a strategic decision tends to 1, the equilibria (robust to natural restrictions on buyers' beliefs) tend to the one discussed in the paragraphs above and generated in the case where the seller always has the opportunity to make a decision to continue or quit.

1.1 Related Literature

Our point of departure is a seminal paper on experimentation, Rothschild (1974), in which an experimenter tries out different arms of a two-armed bandit machine, where each arm pays out a reward with some probability that is learned over time through trial. In each period, the experimenter must weigh up both the short-term benefit and the information gained, which will have a long-term value, from using one arm rather than the other. Rothschild's startling conclusion is that nothing guarantees that in the long-term the experimenter will end up choosing

ation. Apparently, in this environment, previous bankruptcies or failed ventures are not necessarily perceived as an insurmountable obstacle to raising new finance. According to the mechanism described above, this might be explained in part as lenders might perceive that only 'high quality' entrepreneurs would seek to start another new enterprise, whereas 'low quality' entrepreneurs might have given up or been discouraged by previous bankruptcy and the consequent higher costs of funds. The more common explanation is that lenders consider that entrepreneurs who have suffered previous bankruptcies might have gained knowledge or skills through their experiences and presumably learned from previous mistakes. These two explanations do not exclude each other and both mechanisms might operate.

[2]See Tadelis (1999) who proves a no-fully-separating-equilibrium result in a similar framework.

the right arm (that is the arm that is more likely to pay out the reward) and so inefficiency may arise, even in the long run. In a simplified framework with only one arm which pays an uncertain reward, we depart from Rothschild's approach in making the arm—the seller in our model—a much more active participant; we do this in two ways. First we allow the arm to make strategic decisions—the decision of whether or not to continue trading—and so must apply game theoretic techniques to analyse the problem; and secondly, we introduce asymmetric information.

The literature on experimentation and multi-armed bandits has developed in a number of interesting directions; however, of most relevance to this paper is work which has taken a similar approach to this paper in allowing the arms that are being experimented on a strategic role, in particular Bergemann and Valimaki (1996) and Felli and Harris (1996). In these papers, the strategic role is to vary the prices they charge for each trial. Although, considerably simplifying the framework in a couple of respects (in this paper we consider a one-and-a-half armed bandit problem rather than a two-armed problem and effectively suppress any strategic role for the experimenter), the innovation in this paper is the introduction of asymmetric information. Here, we suppose that the seller either knows her own type perfectly or else at least has better information than buyers, rather than supposing that learning and knowledge is symmetric among all participants.

Ottaviani (1999), like this paper, considers a situation where a monopolist effectively controls the learning process. In common with most learning models though in contrast with this paper, there is an assumption that there is no agent with perfect knowledge; however, Ottaviani considers privately informed buyers and focuses on social learning and the potential for informational cascades where a monopolist optimally sets prices over time and the buyers learn about quality by observing each other's decisions to buy. In that environment, the monopolist initially prefers prices that allow more transmission of information from current to future buyers, but eventually either quits the market or captures it entirely. In contrast, we suppose that buyers learn about quality by observing *ex-post* outcomes rather than purchases and focus on the monopolist's decision to trade rather than her pricing decision. Another related paper is Judd and Riordan (1994), which considers price signalling in the second period of a two-period model of a new product monopoly and allows for private information both for buyers and sellers; in contrast, our focus is on the introduction of a different signalling mechanism—willingness to trade at a loss—which can not arise in their model where there is an implicit assumption that there is always a non-trivial efficient level of trade.

Private information in this paper and the adverse selection that arises in this framework brings this paper close to a wide literature on reputation and signalling. In particular, the self-confidence mechanism used to generate the result that a good seller trades in every period is similar to one contained in the model of Milgrom and Roberts (1986). In a two-period model with heterogeneous consumers and free entry, they show that uninformative advertising can play a role in signalling quality in the market for an experience good. Uninformative and costly advertising in their model plays a similar role to the willingness to incur losses in this paper.

Tadelis (1999) presents a model with a similar framework to the one in this paper, in that the only difference between a good seller and a bad seller is the probability with which they produce a successful outcome. Tadelis, however, develops a model in a competitive environment and focuses on shifts in the ownership of firms and the value of the firm's name. In this paper, we consider a monopolist seller and consider the influence of luck on her survival.

1.2 Plan of the paper

In the following section, we introduce the framework for the models we go on to consider. In Section 3, we suppose that the seller has no private information regarding her own type and show that the outcome is equivalent to that described in Rothschild's paper. In Section 4, we suppose that the seller knows her own type perfectly; in this setting we present the central result of the paper—that there is an equilibrium where a good seller always trades. A similar outcome is obtained in the finite horizon case and efficiency is also discussed; in particular, lowering expected profits at all reputation levels can allow arbitrarily close approximation to full efficiency.

The intuition of self-confidence underlying these results—essentially that a seller who knows she is good values continuation more highly than a seller who knows she is bad—is extended in two different directions. First in Section 5, we suppose that the seller has an informative private signal regarding her own type and demonstrate that a seller who had *privately* received a high signal would survive (keep trading) for as long as a seller who had *publicly* received a high signal. Secondly, in Section 6, we vary the degree to which the seller has opportunities to behave strategically by supposing that in each period she can choose to cease trading only with some probability. In this case, we characterise equilibria qualitatively similar to those derived elsewhere. Moreover, we show that this class of equilibria is uniquely robust to a natural restriction on buyers' beliefs and argue that as the probability that the seller can behave strategically tends to 1, these equilibria tend to a unique equilibrium—the one discussed in Section 4.

Finally, Section 7 concludes and discusses further possible extensions.

2 The basic framework

Consider a monopolist seller of a good or service. The seller may be either a good quality seller or a bad quality one.[3] Both good and bad quality types incur a cost c in producing the good and can produce either one unit or no units at each discrete time period. Each time that a seller produces a good, it may turn out successful or unsuccessful with some probability.[4] This probability is independent and identical across time and depends only on the type of the seller.[5] In fact, the two types of seller differ only in that each time that a good seller produces a good,

[3] If quality is drawn from a continuum rather than restricted to two types, there is little substantive difference in the conclusions. At some reputation levels, rather than the semi-separation of this paper in which a seller with whom trade would be inefficient ceases trading with some non-zero and non-unitary probability, there is a cut-off level so that a seller whose innate true quality is below that threshold ceases trading for sure, and a seller whose quality is above that threshold continues trading for sure.

[4] The restriction that the seller can produce only a single unit in each period is an important one, which is fairly common in the experimentation literature. (For a paper on experimentation where quantity does play an important role see Khilstrom et al (1984)). This assumption might seem reasonable if we think of the monopolist seller as an expert selling her labour; here c would be interpreted as a reservation wage earned in alternative employment.

[5] It is reasonable to think of the seller as producing a customised good, such as legal advice or auto-repair. Although, the assumption that the probability of success in any period is independent and identically distributed, is a strong one, which considerably simplifies the analysis, it is not critical for the results. It would be sufficient, for example, that for any particular history a good type would be more likely to produce a success than a bad type, that the likelihood of producing a success is non-decreasing in the number of previous successes and that the stochastic processes generating successes for good and bad types are common knowledge.

it will be successful with probability g and unsuccessful with probability $1-g$, whereas for a bad type the corresponding probabilities are b and $1-b$, where $1>g>b>0$. The seller is risk neutral, and seeks to maximise the discounted present value of profits, where the per-period discount factor is $\beta \in (0,1)$.[6]

Buyers are homogeneous and risk neutral, and at the beginning of period t, they share a common belief that with probability λ_t the seller is good and with probability $1-\lambda_t$ she is bad—this belief is termed the seller's reputation. Buyers assign a value of 1 to a success and 0 to a failure. Thus the *ex-ante* value of buying from a good seller is g and from a bad seller is b. It is assumed that $g>c>b$, or equivalently that *ex-ante* trade with a good seller is socially efficient and trade with a bad seller is inefficient. In each period, given their beliefs and the price of the good, buyers decide whether or not to buy the good. We suppose that buyers buy only once and that in each period there are at least two potential buyers.

In each period the seller first decides whether or not she will seek to trade. If she decided to trade, the seller auctions a single unit of the good. Then this framework in which there is more than one potential buyer, all buyers are homogeneous and buyers buy only once, ensures that the seller extracts the full buyer valuation of her service.[7] Further note, that we therefore exclude any role for price signalling, we discuss the issue in the conclusion of the paper in Section 7 and deal with it explicitly in Appendix 2.

The success or failure of the good can only be determined after it has been bought and consumed. We suppose that this realisation is publicly observable. Although this realisation is publicly observable, it is not verifiable so that contracts contingent on outcomes cannot be written and so the price of the good in any period is independent of the outcome in that period, though it varies with the history of the seller's previous successes and failures. After this realisation has been observed, the seller's reputation is revised and the period ends.

Supposing that the seller begins with some initial reputation λ_0, we can summarise the action as the repetition of the following stages:

1. the seller begins the period with an initial reputation λ_t;

2. she decides whether or not to trade in this period;

3. on the basis of this decision beliefs are revised to the interim belief μ_t;

4. she then produces a single unit at a cost c and sells it, as discussed above, at a price $\mu_t g + (1-\mu_t)b$;

5. the buyer then consumes the good and its realisation as a success or failure is publicly observed; and,

[6]Note that for any discount factor $\beta \in (0,1)$ all the qualitative results derived below hold, though the discount factor would effect the *level* of the critical reputation levels discussed in the propositions below. However for $\beta \in \{0,1\}$, the qualitative results need not hold; in particular if $\beta = 0$, the game is effectively a one-shot game and dynamic reputational concerns are irrelevant, and if $\beta = 1$, there are potentially infinite future rewards.

[7]This assumption is analytically convenient; however it would be sufficient if the price of the good or service depended only on the seller's reputation, and was increasing in the seller's reputation, such that for high enough reputation the price would be higher than c and for low enough reputation it would be below c.

6. on the basis of this realisation buyers revise their beliefs so that the seller's end of period reputation is λ_{t+1} (if the seller did not trade in the period then $\lambda_{t+1} = \mu_t$).

Throughout we suppose that the information structure is common knowledge; for example in Section 4, we suppose not only that the monopolist knows her own type, but that buyers know that she knows her own type (though they do not know what this type) is, the monopolist knows that they know and so on.

3 Uninformed seller

In this section, we suppose that the seller has no private information regarding her type. Instead, her belief regarding her own type is identical to her public reputation; so the seller and potential buyers share the same beliefs and update these beliefs identically. In deciding whether to stop selling or not, the seller takes into account both the profit that she expects to earn in that period and the value of the information generated in a further trial of her ability.

Following a success, or failure, the seller's reputation λ is updated according to Bayes' rule, so that the reputation after a success and failure are given respectively by the following equations:

$$\lambda^s = \frac{\lambda g}{\lambda g + (1-\lambda)b} \tag{1}$$

$$\lambda^f = \frac{\lambda(1-g)}{\lambda(1-g) + (1-\lambda)(1-b)} \tag{2}$$

Symmetric information implies that the only factor that has an influence on a seller's value of trading is her current reputation. Using the equations above and recalling that at any time, the seller can choose not to trade, it follows that an uninformed seller's expected value of selling, given a reputation λ, can be defined by the following equation:

$$\begin{aligned} V^u(\lambda) \quad = \quad & \lambda g + (1-\lambda)b - c + \beta[(\lambda g + (1-\lambda)b)\max\{0, V^u(\lambda^s)\} \\ & +(1-\lambda g - (1-\lambda)b)\max\{0, V^u(\lambda^f)\}] \end{aligned} \tag{3}$$

In this expression, $\lambda g + (1-\lambda)b - c$ represents the profit in the current period. Future expected profits, whether successful—$\max\{0, V^u(\lambda^s)\}$—or not—$\max\{0, V^u(\lambda^f)\}$—are discounted at the discount rate β, and $\lambda g + (1-\lambda)b$ represents the seller's belief that the good will prove successful. Finally note that since the seller can choose not to sell in any period, the value of having the reputation λ is $\max\{0, V^u(\lambda)\}$.

In this environment, if there is a very low belief that the seller is good then the seller would sell at a loss in the current period and would not expect much valuable information to be generated by further trial—if her prior belief is close to 0 then the posterior belief derived as above according to Bayes' rule would not be much different. Equivalently, at a sufficiently low reputation, both the current sale price and the option to continue selling will be so low as not to compensate the seller for the cost of production. This intuition is formalised in the proposition below.

Proposition 1 *There exists a critical reputation level $\lambda^u \in (0,1)$, such that if the reputation at the beginning of a period, λ, is greater than or equal to this critical level, that is $\lambda \geq \lambda^u$, then there is trade and if $\lambda < \lambda^u$, trade ceases.*

Details of the proof appear in the appendix; however, the key elements of the proof involve the use of standard recursive techniques to show that $V^u(\lambda)$ is a unique, well-defined, continuous and increasing function. It can be readily verified that $V^u(0) < 0$ and $V^u(1) > 0$, and so by the continuity and monotonicity of $V^u(\lambda)$, there is a unique λ^u, such that $V^u(\lambda^u) = 0$. Since $V^u(\lambda)$ is monotonically increasing in λ, it follows that for all $\lambda \geq \lambda^u$, $V^u(\lambda) \geq 0$ so that the seller will seek to trade, whereas for $\lambda < \lambda^u$, $V^u(\lambda) < 0$ and the seller would prefer to stop trading. Note that once the seller has stopped trading in one period, the decision that she faces in the next period is identical, since no new information has been generated, so once the seller decides to stop trading, she will never resume.

An important implication of the result is that even a good seller might stop trading in finite time, this is because even a good type could suffer a run of such bad luck that it would drive her reputation down below the critical level. Further note that λ^u can be thought of as an absorbing barrier for reputation—if ever reputation falls below this level, the seller ceases trading.

As a final observation in this section, the problem described in this section is equivalent to a special case of the two-armed bandit problem of Rothschild, with one safe arm (always paying c) and one risky arm (which sometimes pays out 1 and sometimes 0, with a probability that can be learned through experience). Therefore the level of experimentation (the number of periods that a seller trades) in this problem would be identical to that in the one risky, one safe one-and-a-half-arm bandit problem, though note that here it is the arm rather than the experimenter who is strategic.

4 Perfectly informed seller

This section begins by presenting and proving the central result of the paper—that when the seller is perfectly informed about her type, a good seller never stops selling. Underlying this result is the mechanism that in equilibrium a seller's self-confidence when she sells at a loss acts as a signal of her quality. Throughout this section and indeed in the sections below, the introduction of private information allows a multiplicity of equilibria in contrast to the unique outcome that results in the case where the seller is uninformed. We discuss this multiplicity at some length below. Nevertheless, at the end of Section 4.1 and in Section 6 below, we argue that the equilibria we characterise in this section are appealing and robust equilibria.

4.1 The perpetual survival of a good seller in the infinite horizon model

Intuitively, underlying the result in this section and indeed most of the other results below is the observation that good sellers have more to lose from ceasing trade than bad sellers. In particular, a good seller knows that by continuing to trade, her reputation is likely to be enhanced. Potential buyers, who know that the seller knows her type perfectly, are aware of this and so take account of the fact that the seller is willing to trade in assessing the likelihood that she is good.

More formally, that a good seller prefers to continue trading relative to a bad type can be seen clearly from a comparison of the Bellman equations for the value of a sale for a good type and a bad type respectively.

$$V^g(\lambda) = \mu g + (1-\mu)b - c + \beta[g \max\{0, V^g(\mu^s)\} + (1-g)\max\{0, V^g(\mu^f)\}] \tag{4}$$

$$V^b(\lambda) = \mu g + (1-\mu)b - c + \beta[b \max\{0, V^b(\mu^s)\} + (1-b)\max\{0, V^b(\mu^f)\}] \tag{5}$$

Note that for a particular reputation, λ, the short-term profit is the same for a good or bad seller; it is simply $\mu g + (1-\mu)b - c$. It differs from the short-term profit for an uninformed seller (see Equation (3)), which is $\lambda g + (1-\lambda)b - c$, since in the case where the seller is informed her decision to trade is informative and so the buyers' belief is conditioned on the fact that the seller has traded. The interim reputation is revised to μ, which will be determined by equilibrium strategies and beliefs and may differ from λ.

The only difference between Equations (4) and (5) is that in the value function for a good seller a future success is more likely than in the value function for a bad seller; it occurs with probability g rather than b. Intuition suggests that for both good and bad sellers, the value of a sale should be increasing in current reputation ($V(\lambda)$ should be increasing in λ) so that, since $\mu^s > \mu^f$, good sellers value continuation more highly than bad sellers do, at almost all levels of reputation.[8] In the equilibrium described in Proposition 2 below, this intuition is borne out and these properties hold. Therefore, a bad seller would be more willing to drop out of the market than a good one. Indeed we show that a good seller never drops out for any discount factor $\beta \in (0, 1)$. This is a consequence of the buyers' knowledge that a good seller is relatively more willing to incur a loss to continue trading, which induces a lower bound for the reputation of a seller conditional on the fact that she has not ceased trading or, equivalently, a limit to how low the sale price can fall.

For a given prior belief λ_t, if buyers believe that both good types and bad types would trade for sure, then the interim belief is equal to this prior reputation, $\mu_t = \lambda_t$. If however, it is believed that bad types would drop out with some positive probability and good types would always stay in the market then $\mu_t > \lambda_t$. Equivalently, conditional on her remaining in the market, the seller is more likely to be good. In particular, if it is believed that bad types always drop out at a particular reputation level then $\mu_t = 1$.

If in equilibrium, a bad type always stayed in the market, regardless of her current reputation, then at some reputation levels the sum of the continuation value and the short-term profit would be negative, since buyers' beliefs could drop arbitrarily low. This could not be an equilibrium strategy. At another extreme, if at some reputation levels, a bad seller would drop out for sure, then buyers would assume that any seller seeking to trade must be a good type and so would be willing to pay g at all periods in the future, a bad seller would then clearly prefer to stay in and earn $(g - c)/(1 - \beta)$, again this can not be an equilibrium.[9] Thus it must be the case that, in equilibrium, for some reputation levels a bad seller would drop out with some probability and

[8] The value functions for the good and bad sellers will take the same value if the reputation is either 0 or 1, since in these cases future beliefs will stay at the initial level whatever happens.

[9] Note that, it is at this point that there will be a difference in Sections 5 and 6, where in equilibrium there can be full separation of types (though note that in Section 5 a type is defined by the signal received rather than innate ability). In essence this derives from the fact that in those environments, even if there is a full separation

continue trading with the complementary probability. Such mixed strategies on the part of a bad seller allows some continuity in the beliefs of buyers which is crucial for the existence of an equilibrium of the form described below. In particular, in equilibrium a bad seller employs a mixed strategy of continuing to trade that ensures that her value of continuing is zero. We state the result formally in Proposition 2 below.

First, consider the value function for a bad type if buyers believed that sellers never stopped trading regardless of type. This is implicitly defined by the following Bellman equation:

$$\begin{aligned} W^b(\lambda) &= \lambda g + (1-\lambda)b - c + \beta b \max\{0, W^b(\frac{\lambda g}{\lambda g + (1-\lambda)b})\} \\ &\quad + \beta(1-b)\max\{0, W^b(\frac{\lambda(1-g)}{\lambda(1-g)+(1-\lambda)(1-b)})\} \end{aligned} \tag{6}$$

Note that in Equation (6), since buyers believe that all types always sells, the price in the short-run is $\lambda g+(1-\lambda)b$. It can be shown, that there is a unique solution to this Bellman equation using the standard recursive techniques used to prove Proposition 1. Moreover, the solution $W^b(\lambda)$ is a well-defined, bounded, continuous and strictly increasing function. Furthermore, there exists a unique λ' such that $W^b(\lambda') = 0$. Note, that λ' depends on β, in particular λ' is decreasing in β. However, for any β, the following result holds.

Proposition 2 *Suppose that the seller's prior reputation is $\lambda \in (0,1)$. Then there is an equilibrium, in which if $\lambda \geq \lambda'$ trade occurs for sure—that is both a good and bad seller would trade with probability* 1*. If $\lambda < \lambda'$, then a good seller trades for sure and a bad seller continues trading with probability $d(\lambda)$ and ceases trading with a probability $(1-d(\lambda))$, where $d(\lambda) = \frac{\lambda(1-\lambda')}{(1-\lambda)\lambda'}$.*

Essentially both good and bad types sell for sure at high reputation levels, that is above λ', but if her reputation falls then a bad type would cease trading with a mixed strategy (that is she trades with probability $d(\lambda)$) such that she would be indifferent between continuing to trade or ceasing, and at this interim reputation a good seller strictly prefers to sell, which in equilibrium she does. This result is illustrated in Figure 1, which represents a sample path for the reputation of a seller (contingent on her willingness to sell). The reputation level λ' acts as a lower bound for the reputation contingent on selling. Thus for a good seller λ' acts as a reflecting barrier and for a bad seller it is a partially absorbing barrier. It is further worth noting that $\lambda' < \frac{c-b}{g-b}$, where this latter value is the reputation level that allows the seller to break even in the one period game.[10]

(Figure 1 page 43)

The outcome characterised in this proposition is not the unique equilibrium outcome in this environment. Other equilibria can be generated though they rely on specific (and arbitrary)

of types, there is still "noise" that allows Bayesian updating to continue and this leads to a region in the space of reputations where there is full separation between different types of sellers. See Sections 5 and 6 for further details.

[10] In a one-shot game, a seller would behave identically whether good or bad, and would sell whenever the profit—the difference between the sale price and the cost—was positive, that is when $\lambda g + (1-\lambda)b - c > 0$.

off-equilibrium beliefs. For example, if buyers believed that a good seller would always cease trading whatever the current reputation, then there would never be any trade. This kind of equilibrium is robust to the Cho and Kreps (1987) Intuitive Criterion, for example, since in the environment of this paper the only difference between the two types of seller is with respect to the frequency with which each type generates success, and at the *ex-ante* stage when the sale price is determined, the profit realised depends only on buyers' beliefs, so that arbitrary beliefs would have exactly the same impact on a good seller as on a bad seller. However, equilibria other than the one described in Proposition 2 above, do rely on specific off-equilibrium beliefs. In the equilibrium outlined in Proposition 2, the only possible off-equilibrium actions are not selling when reputation is high (since the only decision is whether or not to continue trading and continuation is always an action that might be chosen in equilibrium) but such behaviour does not require a response from buyers. Any other behaviour is equilibrium behaviour and so, is not reliant on specific arbitrary off-equilibrium beliefs. Moreover, as shown in Section 6 the equilibrium outlined in Proposition 2 is the unique outcome as a limit within a class of equilibria robust to a natural restriction on beliefs—that the buyers' belief contingent on observing that the seller is willing to trade should be non-decreasing in the beginning of period belief—in the more general environment considered in that section.

4.2 The finite horizon problem

In this section, we show that an analogous result to Proposition 2, which held in the infinite horizon problem, applies in the finite horizon model (where there are only a fixed finite number of opportunities to trade). The result for the finite case is somewhat altered in that in the last period even a good seller may not trade, though in all previous periods a good seller would trade with certainty. A good seller's self-confidence underlies this result—the value of remaining in the market and having further opportunities to trade is worth more to a good seller than to a bad one and so a good seller is more willing to incur a short-term loss to remain in the market (except in the last period where there is no continuation value). Again, this is recognised by buyers whose beliefs are contingent on the observation that the seller has not yet ceased trading, and this sets a floor to the losses that a seller need incur in order to continue trading, in such a way that a good seller always trades (except possibly in the last period). Note, however that the level of this effective floor to reputation varies with the number of periods remaining. In particular, the greater the number of periods remaining, then for any given reputation level the higher the value of continuing for both types of sellers, and so the lower the effective floor to reputation, which is the level of reputation at which a bad seller is indifferent between trading and stopping to trade. These considerations are summarised in the proposition below; the proof of which is based on backward induction and appears in the Appendix. Again, although the outcome described is not a unique equilibrium, it does show that self-confidence is not reliant on an infinite horizon, but can apply in a finite horizon problem. Moreover, the equilibrium described is the unique limit of equilibria described in Section 6, which are uniquely robust to natural restrictions on beliefs.

Proposition 3 *Suppose that there are $N > 1$ trading opportunities and that the current reputation level is $\lambda \in (0, 1)$, then there exists an equil;ibrium with the following characteristics. A good seller sells with probability 1 whatever her reputation level. There exists a reputation level λ_N such that a bad seller sells with probability 1 if $\lambda \geq \lambda_N$. Otherwise a bad seller sells with* *probability less than one. Furthermore λ_N is decreasing in N. Finally if $N = 1$, then both a good seller and a bad seller would behave in the same way, specifically they would sell if and only if $\lambda g + (1 - \lambda)b - c \geq 0$.*

4.3 Efficiency

In this framework, an informed central planner, who knew the seller's type and wished to maximise the total payoffs to the seller and the buyers, would ensure that there was trade if and only if the seller were good. Although, all decision making is *ex-ante* efficient in the sense of Holmstrom and Myerson (1983), the lack of such a planner and the lack of information concerning the seller's type (in Section 3) or asymmetry of information (in Sections 4.1 and 4.2) can lead to two types of inefficiency, or "errors" in trading:

Type I: No trade with a good seller.

Type II: Trade with a bad seller.

Thus the outcome in Section 3, analogous to the outcome in the Rothschild framework, suffers from both types of inefficiency—the seller might stop selling even though she is good and a bad seller might sell. The equilibrium characterised in Proposition 2 has no Type I inefficiency—a good seller always trades—but suffers from Type II inefficiency—a bad seller would also trade with some probability.

If there are lower potential profits available, for example if there were taxes or for some other reason the seller could only appropriate a fraction of the buyer's valuation, then the value to having any reputation would be lower for both good and bad sellers. In particular, the lower value for a bad seller suggests that the reputation threshold at which a bad seller would start to stop trading would be higher, or equivalently that she would be (weakly) more likely to cease trading at all reputation levels. This would reduce the risk of the potential Type II error. The value for a good seller is almost always higher than the value for a bad seller, and so a good seller would never cease trading, through the self-confidence mechanism described above, so that as before there would be no Type I error. We formalise this intuition below in the specific case where the seller can only recover a proportion of the buyer's valuation.

Specifically, suppose that the seller keeps only a fraction γ of the price at which trade occurs. Further suppose that $\gamma \geq c/g$. (If this were not the case then even if it were common knowledge that the seller were good it would still not be worthwhile for her to trade as her profit would be $\gamma g - c < 0$). Then as in Proposition 2, there is a natural equilibrium with no Type I inefficiency and by choosing γ appropriately, the Type II inefficiency can be made arbitrarily small. Loosely, the lower is γ, the lower the potential gains from trading for a seller whether good or bad, therefore the higher the probability that a bad seller stops trading at any given prior reputation. As before the continuation value for a good seller is higher than that for a bad one, so a good seller is still willing to trade. Thus by reducing γ (equivalently increasing a per-unit tax), the Type II inefficiency can be reduced and there is still no Type I inefficiency. This is made explicit in Proposition 4 which is stated below and proved in the Appendix.

Proposition 4 *Suppose that for a given interim reputation level μ, the seller's revenue is $\gamma[\mu g + (1 - \mu)b]$, with $\gamma > c/g$, then there is an equilibrium with the following characteristics. A good seller always sells and there exists a $\lambda'(\gamma) \in (0, 1)$ such that a bad seller sells with certainty if the* *prior belief* $\lambda \geq \lambda'(\gamma)$ *and sells with probability* $d(\lambda, \gamma)$ *otherwise. Furthermore* $\lambda'(\gamma)$ *is decreasing in* γ *and* $d(\lambda, \gamma)$ *is increasing in* γ.

Note that $\lambda'(\gamma)$ is decreasing in γ and $d(\lambda, \gamma)$ is increasing in γ means that the lower the proportion of the consumer valuation that the seller can appropriate, the more likely that a bad seller would cease trading, and the greater the level of efficiency. In the limiting case where $\gamma = c/g$, only good sellers would trade; such behaviour is rational, since there would be neither profits nor losses in trading, so that continuing to trade is a rational strategy for a good seller, and never trading is a rational strategy for a bad one.

5 Imperfect private information

In this section we consider an intermediate situation in which the seller has some imperfect private information concerning her own type, which is superior to the information that buyers commonly hold.

The self-confidence mechanism identified in the previous section operates in this more general environment and leads to a more general result. A seller who has discouraging private information regarding her own type will have a lower continuation value than a seller with encouraging information, and therefore will be more likely to drop out in either a deterministic or probabalistic sense at low reputation levels. This would imply that the interim reputation level—the reputation contingent on the observation that the seller is self-confident, in the sense of being willing to continue trading—would rise to the encouraged seller's own assessment of her quality and would reach this level before a seller with encouraging private information was induced to cease trading. Thus for the survival of a seller with "good news" as to her type, there is an equilibrium in which it does not matter whether this "good news" is publicly known or her own private information.

We illustrate this first in an infinite horizon setting, in which the seller receives information regarding her type before any trading begins, refining the intuition discussed in the paragraph above. Then we suppose that the seller keeps receiving informative signals as to her type through time, but for the tractability of this problem, we analyse it assuming a finite horizon and a number of natural restrictions on the evolution of the buyers' and seller's beliefs.

5.1 An informative signal at the beginning of time: infinite horizon

We suppose that before any opportunity for trading arises, the seller privately receives either a high or a low signal. If she receives a high signal then she believes initially that she is good with probability $\overline{\lambda}_0$ and otherwise bad, whereas if she receives a low signal the corresponding probability is $\underline{\lambda}_0$, where $1 > \overline{\lambda}_0 > \underline{\lambda}_0 > 0$. Furthermore, we assume that this signal is private information and that initially the buyers have some belief, r_0, such that buyers assign a probability r_0 to the seller having received a high signal and a probability $(1 - r_0)$ to the seller having received a low signal, so that $\lambda_0 = r_0\overline{\lambda}_0 + (1 - r_0)\underline{\lambda}_0$. Thus the information set of the seller is superior to the buyers' information set.

In this environment, there are a number of different beliefs to bear in mind, first as above we denote the buyers' belief that the seller is good λ_t. In contrast to Section 4, the seller is not perfectly informed as to her own type and so revises her belief in the light of the success and failures that she realises. Thus the seller's own belief that she is good depends on her history and also on the signal that she received. We denote the belief at time t of a seller with a good signal by $\overline{\lambda}_t$ and the belief of a seller with a bad signal by $\underline{\lambda}_t$. Note that these beliefs are interrelated and depend on a commonly observed history. Specifically, we assume that $\overline{\lambda}_0$, and $\underline{\lambda}_0$ are common knowledge, so that for any history of successes and failures h_t, all agents can calculate $\overline{\lambda}_t(h_t)$ and $\underline{\lambda}_t(h_t)$, we further assume that r_0 is common knowledge so that in equilibrium $r_t(h_t)$ and so also $\lambda_t(h_t)$ would be common knowledge.

In this environment, the seller does not have perfect information as to her type and revises her belief that she is good in the light of previous history. It follows that even a good seller who had a high signal might stop trading. She could have such a run of bad luck that buyers would be fairly sure that she was bad (so that she would have to make significant short-term losses in order to continue selling) and moreover she would herself be fairly certain that she was bad so that she would consider it unlikely that her reputation would improve as a result of further trading. However in this environment the self-confidence mechanism still applies. We discuss this below.

First, consider the benchmark case in which the seller's signal is public information. In this case, the value function for a seller with a high signal is identical to the value function of the uninformed seller defined in Equation (3), though the initial and subsequent beliefs depend on the signal that the seller received, so that for a given history h_t, the value of selling in this period is given by $V^u(\overline{\lambda}_t(h_t))$. Therefore, following Proposition 1, there is a critical reputation level λ^u, such that as soon as $\overline{\lambda}_t(h_t) < \lambda^u$ the seller with the high (public) signal stops selling. Note that since $\overline{\lambda}_t(h_t)$ can take any value in $(0, 1)$ with an appropriate history, the set of histories, $\overline{H}$, for which a seller with a high (public) signal stops selling is not empty—this intuition is discussed in the paragraph above.[11]

Now returning to the case where the seller's signal is *private* information (though it is common knowledge that she receives such a signal), there is an equilibrium in which a seller with a high signal behaves in exactly the same way that she would if that signal were public. That is, there is no distortion for the survival of a seller with a high signal between seeing that signal privately and publicly. Equivalently, and somewhat more formally, where the seller's signal is private information, there is an equilibrium in which the set of histories that induce a seller with a high signal to stop selling, $\overline{PH}$, is exactly the same set of histories $\overline{H}$ that would induce a seller with a high public signal to stop selling.

Proposition 5 *There is an equilibrium in which the set of histories $\overline{PH}$ which induce a seller with a high* private *signal to stop selling is identical to $\overline{H}$, the set of histories that would induce a seller with a high* public *signal to stop selling, that is $\overline{PH} = \overline{H}$.*

The intuition underlying this outcome still derives from the high-signal seller's self-confidence. A seller with a high signal always prefers to continue selling than a seller with a low signal. Therefore, the seller's decision to continue selling is informative and affects buyers' beliefs, so that these rise above the prior belief. However, the buyers' belief can be no higher than if the

[11] Note that similarly a seller with a bad (public) signal will stop selling as soon as $\underline{\lambda}_t(h_t) < \lambda^u$; however, since $\underline{\lambda}_0 < \overline{\lambda}_0$ and so $\underline{\lambda}_t(h_t) < \overline{\lambda}_t(h_t)$ for all h_t, this event would be sooner than for a seller with a good public signal. Formally, if $\underline{H}$ is the set of histories which lead a seller with a bad (public) signal to stop selling, then $\overline{H} \subset \underline{H}$.

buyers were certain that the seller had received a high signal, though it can reach this level. It follows that $\overline{PH} = \overline{H}$. A more formal proof appears in the appendix.

A few other aspects of the equilibrium described in Proposition 5 are noteworthy. First note that similar to the equilibria described in Section 4, there will be a range of histories for which sellers who had received low signals will employ mixed strategies and continue selling with some non-zero and non-certain probability. However, in contrast with those equilibria, a consequence of *imperfect* private information is that there will be a range of histories for which sellers who had received a high signal will trade with certainty and sellers who had received a low signal will cease trading with certainty. We state this more formally, as follows.

Remark 1 *There exists a non-empty range of reputation levels* $[\lambda^u, \lambda'(1)]$ *for which there is full separation of sellers in the sense that a high-signal seller with reputation* $\lambda \in [\lambda^u, \lambda'(1)]$ *would sell for sure and a low-signal seller with reputation* $\lambda \in [\lambda^u, \lambda'(1)]$ *would cease trading for sure.*

This follows since even if buyers believed that the seller had received a high signal, this would not imply that she is good (this contrasts with the environment of Section 4) and so this may not forestall short-term losses, and moreover a seller with a low signal would expect that her reputation is relatively likely to deteriorate in comparison with a seller with a high signal and so would still value continuation less than one with a high signal. A formal proof appears in the Appendix.

Furthermore, as discussed above, there will be a range of histories after which a seller with a high signal (and *a fortiori* a seller with a low signal) would cease trading. Finally, again the equilibrium generating the outcome in Proposition 5 is not the unique equilibrium in this environment, however in contrast to other equilibrium outcomes it is the only one not sensitive to off-equilibrium beliefs. It is an equilibrium regardless of what buyers would believe if they were to observe a seller continuing to sell following a history in $\overline{PH}$ (that is whether they would believe that this shows that the seller had a low signal or a high signal or would have any intermediate belief).[12]

Figure 2 page 43

Figure 2 illustrates the result characterised in Proposition 5. The result is that there is an equilibrium in which public perception λ would converge to the self-belief of a seller who had had a positive signal $\overline{\lambda}$, before such a seller would cease trading (time C in the figure). This follows because a seller who had had a negative signal would cease trading sooner. The low-signal seller might cease trading with some positive probability as at A in the figure, so that the public reputation would move closer to the belief of a seller who had had a good signal. Alternatively, the seller with bad news might cease trading with certainty, as at B, so that the public reputation (which is contingent on the seller continuing to trade) would be exactly equal to $\overline{\lambda}$.

5.2 Seller learning over time: finite horizon

Intuitively, we conjecture that the result need not be confined to the case where the seller receives private signals only at the beginning of time; rather, there is an equilibrium in which the seller

[12] See the discussion in the final paragraph of Section 4.1 for more on this issue.

would continue trading whenever she has private information that would ensure survival in that period if it were public. In order to keep the treatment simple, we state such a result in the finite horizon version of the model, below. The proof of the result appears in the appendix.

We suppose that the seller keeps getting better information as to her own type and we term this the seller's self-belief or self-perception p and summarise the beliefs of the public and the seller at the beginning of each period by (λ, p). We assume that the information available to the public is also available to the seller, so that p is always more informative than λ and that if the seller's information was observed by the public, they would believe that the seller was good with probability p. Consider the cumulative distribution function $F(\nu, q \mid \mu, p)$ which characterises the evolution of public reputation and self perception, given that just prior to consumption (that is at the interim stage) these are given by (μ, p). Suppose that $F(\nu, q \mid \mu, p)$ exhibits First Order Stochastic Dominance with respect to p and μ, that is for any μ, if $q > p$ then $F(\cdot, \cdot \mid \mu, q)$ first order stochastically dominates $F(\cdot, \cdot \mid \mu, p)$ and for any p, if $\mu > \nu$ then $F(\cdot, \cdot \mid \mu, p)$ first order stochastically dominates $F(\cdot, \cdot \mid \nu, p)$, moreover we assume that $\int f(\mu, q \mid \mu, p) dq = 1$ if and only if $\mu = 0$ or 1 and $\int f(\nu, p \mid \mu, p) d\nu = 1$ if and only if $p = 0$ or 1 and that $f(\nu, p \mid \mu, p)$ is continuous in μ. Finally $E(\nu, q \mid p, p) = p$ and $E(\nu, q \mid \mu, p) = (\cdot, p)$.

First consider the problem in the case where, the public has access to all of the seller's information so that $\lambda = \mu = p$ always and consider the recursively defined:

$U_1(p) = \max\{0, pg + (1-p)b - c\}$

$U_n(p) = \max\{0, pg + (1-p)b - c + \beta E[U_{n-1}(q) \mid p]\}$

Thus $U_n(p)$ denotes the value of having reputation p when there are n remaining trading opportunities. With this definition, we state the following proposition:

Proposition 6 *Suppose that there are $n > 1$ remaining trading opportunities and the current reputation and self-belief are given by (λ, p). Then there is an equilibrium which generates the value functions $V_m(\lambda, p)$ for $m \leq n$ which has the following property: If $U_m(p) > 0$ and $m > 1$, then $V_m(\lambda, p) > 0$ for all $\lambda > 0$.*

Thus if the seller's own self-perception p were sufficiently high such that if it were public knowledge then the seller would continue trading, then in the equilibrium described in the proposition, the seller would continue selling even though p is private information, whatever the seller's public reputation (except the case where the public reputation is that it is certain that the seller is bad).

The proof of the result proceeds by constructing strategies through backward induction. The logic of the proof can not be transferred to the infinite horizon framework. The difficulty in proving the case for the infinite horizon model and with Markov strategies is that changing strategies at one state affects the value at all other states (so if we try and modify an equilibrium where all drop out at one reputation level it affects the equilibrium strategies at all other reputation levels). The finite case is easier since here a state depends on both the reputations (public and self-belief) and the time, so changing what happens in one state will not affect anything that happens after, only before, so we can proceed as in the Appendix by backward induction.

6 Varying the degree to which the seller can behave strategically

In this section, we further consider the robustness of the self-confidence mechanism introduced in Section 4 and, in particular, the equilibrium characterised in Proposition 2. We introduce a more general environment and a class of equilibria robust to natural restrictions on buyers' beliefs and show that the equilibrium characterised in Proposition 2 is the unique limit of equilibria among that class of equilibria. Moreover, in the more general environment of this section the self-confidence mechanism continues to operate.

We suppose throughout this section that the seller is fully informed as to her type, as in Section 4, but we consider varying the degree to which the seller can act on her information—that is varying her opportunity to behave strategically. Specifically, suppose now that in each period the seller has the opportunity to cease trading not with certainty as before but only with probability α. Further suppose that potential buyers can not observe whether the seller continues trading out of choice or because she did not have the opportunity to cease. Thus α is essentially a measure of how strategic the seller can be. If $\alpha = 0$ then the seller has no decision to make—both a good and a bad seller must keep selling forever. If $\alpha = 1$, then the seller has the opportunity to cease trading in every period (this is the case discussed in Section 4) and so in a sense there is no limit to how revealing the observation that the seller continues trading can be.

Akin to trembling hand refinements, $1 - \alpha$ can be though of as the probability with which the seller "makes a mistake" and continues to trade when it is not rational to do so.[13] In this case, however, we do assume that in all other periods the seller is rational to the extent of taking into account her potential future fallibility.

For $\alpha = 0$, there is no decision making and the outcome is unique and $\alpha = 1$ is the fully strategic case considered in Section 4. Suppose that $\alpha \in (0, 1)$ then we begin by stating the following result, which characterises a seller's behaviour when given the opportunity to make the strategic decision to continue or cease trading.

Lemma 1 *In any equilibrium there exist reputation levels $0 < l < h < 1$, such that if the prior reputation is λ, then (i) if $\lambda < l$ then both bad types and good types would cease trading, and (ii) if $\lambda > h$ then if given the opportunity both types would continue trading.*

The intuition for this result is that beliefs cannot move "too fast". In previous sections, the interim reputation—the reputation contingent on the observation that the seller continues to trade—could, in principle take any value between 0 and 1. Here it is related to the prior reputation. If the seller's reputation is high enough initially, then even if the seller has failures from now on and buyers' beliefs are as pessimistic as possible—that is buyers believe that when given the opportunity only bad types would continue trading—then the current period reward and a sufficient number of future periods' reward would be high enough above cost to ensure that either type of seller prefers to continue trading. Similarly, since beliefs cannot move too fast,

[13]Supposing that similarly there is a probability that the seller would cease trading when it is not rational to do so would not change the qualitative results in this section, so long as we assume that once a seller stops trading she cannot resume trading in later periods.

if a seller has a very low reputation it takes a long time to recover even if things go as well as possible; since the seller would be incurring losses in this time, she might prefer to cease trading. We formalise this intuition in the Appendix, for one of the two cases (the other would have an analogous proof), specifically we assume that there is no such l and show that this leads to a contradiction.

In the sections below, we argue that natural restrictions on beliefs lead to particular classes of equilibria, which in the limit converge to those described in Proposition 2 for the infinite horizon case and Proposition 3 for the finite horizon version; though at $\alpha = 1$, these restrictions are not sufficient to ensure uniqueness.

6.1 Infinite horizon

We begin by making a couple of natural assumptions on buyers' beliefs and argue that these lead to a particular form of equilibrium. Specifically we make the following assumptions:

Assumption 1: *Buyers' interim beliefs are non-decreasing in their prior (beginning of period) beliefs, i.e. $\mu(\lambda)$ is non-decreasing in λ.*

We argue that under this assumption and for any $\alpha \in (0, 1)$ self-confidence leads to equilibria of a particular form:

Lemma 2 *In any equilibrium in which buyers' beliefs satisfy Assumption 1, there exist values $0 \leq k \leq l \leq m \leq h \leq 1$ such that the seller makes strategic decisions to continue or cease trading as follows:*

- *with prior reputations $\lambda < k$, both good and bad types would cease trading for sure;*
- *for $k \leq \lambda < l$, a good type continues trading with some probability and a bad type would cease trading for sure;*
- *for $l \leq \lambda < m$, a good type continues trading for sure and a bad type would cease trading for sure;*
- *for $m \leq \lambda < h$, a good type continues trading for sure and a bad type would continue trading with some probability that ensures that in equilibrium the interim reputation would be h;*
- *for $h \leq \lambda$, both types continue trading for sure.*

The result follows since the assumptions imply that the value functions for both good and bad types are increasing in the prior belief and that the value of any prior reputation is never lower for a good type then for a bad type; these two properties of the value functions of good and bad types ensure that any equilibrium must be of the hypothesised form.

Next we show that an equilibrium of this form does indeed exist, we do this by showing the existence of equilibria within a subclass of the above form, for which $k = l$.

Proposition 7 *For any $\alpha \in (0, 1)$, there exist $0 < l < m < 1$ such that the following strategies with associated beliefs form an equilibrium: when given the opportunity (a) with prior reputations $\lambda < l$, both good and bad types would cease trading for sure (b) for $l \leq \lambda < m$, a good type* *continues trading for sure and a bad type would cease trading for sure (c) for $m \leq \lambda < m/[m+(1-m)(1-\alpha)]$, a good type continues trading for sure and a bad type would continue trading with a probability that ensures that in equilibrium the interim reputation would be $m/[m+(1-m)(1-\alpha)]$ and (d) for $m/[m+(1-m)(1-\alpha)] \leq \lambda$, both types continue trading for sure.*

The formal proof of the result, the details of which appear in the Appendix, sets out Bellman equations for the values of a good and bad seller which depend on how the interim reputation is derived from the prior reputation (which is determined by the equilibrium beliefs). Standard recursive techniques are then invoked to characterise properties of these value functions. Finally, sufficient conditions to ensure equilibrium can be summarised as conditions on l and m, the proof concludes in showing that these conditions can be satisfied with $0 < l < m < 1$.[14]

Proposition 7 states that an equilibrium robust to Assumption 1 exists; however, many such equilibria might exist. Loosely, this multiplicity arises at relatively low levels of prior reputation within the range where a bad seller would not trade even if buyers beliefs were optimistic (that is they believed that only good sellers choose to trade when given the chance). Within this range, there are reputation levels where a good seller would be willing to trade if buyers were optimistic but not if they were more pessimistic (that is if they believed that neither type would trade). Thus either combination of beliefs and strategies would be mutually consistent and a multiplicity of equilibria might arise. The intuition is somewhat complicated in as much as changing the good seller's equilibrium strategies in a range of reputations would change all other strategies, whereas in the finite horizon version of the model, as discussed below, the effect can be isolated within a particular time period and so we can explicitly show the multiplicity of equilibria. However this intuition does suggest that a multiplicity of equilibria exist.

We argue, however, that as α tends to 1, the limiting case is unique, moreover the unique limiting equilibrium is the one characterised in Proposition 2. Note however, that in the case $\alpha = 1$, there are a multiplicity of equilibria of the form described, consistent with Assumptions 1 and 2. For example, in addition to the equilibrium outlined in Proposition 2, there is an equilibrium where buyers have the off-equilibrium belief that only a bad type would seek to sell at any reputation level, then $\mu(\lambda) = 0$ for all $\lambda \in [0, 1)$ and so is non-decreasing and the equilibrium is of the hypothesised form with $l = m = h = 1$.

Proposition 8 *As α tends to 1, the set of feasible equilibria of the hypothesised form (that is as described in Lemma 2) converge to a singleton. Moreover, this is the equilibrium characterised in Proposition 2.*

The proof rests on showing that as α tends to 0, the set of feasible equilibrium l tends to $\{0\}$. This is done by characterising an upper bound for this set (trivially the lower bound is 0) and showing that this upper bound tends to 0. The details of the proof appear in the Appendix.

This result, that the class of equilibria which satisfy two natural assumptions—that buyers' interim beliefs are non-decreasing in their prior beliefs and that buyers believe that the probability with which a seller of either type sells is non-decreasing in her reputation—converges to the

[14] Note that if the prior reputation is m and buyers believe that a good seller would continue trading when given the opportunity and a bad seller would stop, then the interim reputation would be $m(\alpha+1-\alpha)/[m(\alpha+1-\alpha)+(1-m)(1-\alpha)] = m/[m+(1-m)(1-\alpha)]$. Further note that m is the lowest prior reputation for which this interim reputation could be attained.

equilibrium characterised in Proposition 2 makes a strong case for the earlier focus of the paper on that result and the similar results discussed in earlier sections.

6.2 Finite horizon

We begin by arguing that the qualitative results shown for the infinite horizon, also apply in the finite case.

First we argue that the finite horizon analogues of Assumption 1, together with another natural assumption, restricts the set of feasible equilibria to equilibria of a similar particular form in the finite horizon problem. Specifically we make the following assumptions:

Assumption 2: *With n periods remaining, buyers' interim beliefs are non-decreasing in their prior beliefs, i.e.* $\mu_n(\lambda)$ *is non-decreasing in* λ.

Assumption 3: *Buyers believe that both types of seller behave identically in the final period.*

Lemma 3 *In the finite horizon model, in any equilibrium in which buyers beliefs satisfy Assumption 2 and 3, there exist values* $0 \leq k_n \leq l_n \leq m_n \leq h_n \leq 1$ *such that the seller makes strategic decisions to continue or cease trading as follows:*

- *with prior reputations* $\lambda < k_n$*, both good and bad types would cease trading for sure;*
- *for* $k_n \leq \lambda < l_n$*, a good type continues trading with some probability and a bad type would cease trading for sure;*
- *for* $l_n \leq \lambda < m_n$*, a good type continues trading for sure and a bad type would cease trading for sure;*
- *for* $m_n \leq \lambda < h_n$*, a good type continues trading for sure and a bad type would continue trading with some probability that ensures that in equilibrium the interim reputation would be h;*
- *for* $h_n \leq \lambda$*, both types continue trading for sure.*

In the Appendix we prove the result through backward induction, noting that Assumption 3 implies that in the final period good and bad sellers will choose to trade if and only if $\lambda > (c-b)/(g-b)$ and proceeding to show that this together with Assumptions 2 implies non-decreasing value functions such that the value of a good seller for any reputation level is strictly greater than the value to a bad one if there are more than one remaining trading opportunities, implying that any equilibria must be of the hypothesised form.

Next we show that an equilibrium of the hypothesised form can exist.

Proposition 9 *With n remaining trading opportunities, there exists an equilibrium defined by parameters* l_n*,* m_n *and* h_n*, and the following strategies and corresponding beliefs (a)* $\lambda \leq l_n$ *then both good and bad types would cease trading (b) if* $l_n \leq \lambda < m_n$ *then a good seller continues for sure and bad would cease for sure (c) if* $m_n \leq \lambda < h_n$ *then a good seller continues for sure and a bad would cease probabilistically so that the interim belief is* h_n *and (d) if* $h_n \leq \lambda$ *then both types continue for sure.*

The intuition underlying this result again rests on the self-confidence mechanism which suggests that a good seller would be more willing to continue trading at any reputation level than a bad seller; the proof proceeds by inductively constructing an equilibrium of this form and can be found in the Appendix.

For general α, there is a multiplicity of equilibria of the type characterised in Proposition 9. Note, however, that there is no multiplicity of equilibria of this type even with one period remaining; in this case, it can be readily verified that the unique equilibrium of this form has $l_1 = m_1 = h_1 = \frac{c-b}{g-b}$.

However, in the model with two or more remaining periods, multiplicity can arise. This follows since a good seller will choose to continue trading so long as the expected value from doing so is non-negative. This value depends only on the interim reputation and the same interim reputation could be achieved with a relatively low reputation and optimistic beliefs (that is with buyers believing that the good seller chooses to trade) or with a relatively higher reputation and pessimistic beliefs (that is with buyers believing that the good seller would cease trading if given the opportunity). Hence a multiplicity of equilibria can arise;[15] though we argue that as α tends to 1, the set of feasible equilibria converges to a singleton.

Proposition 10 *As α tends to 1, the set of feasible equilibria of the hypothesised form (that is as described in Lemma 3) converge to a singleton. Moreover, this is the equilibrium characterised in Proposition 3.*[16]

Thus we have shown that natural restrictions on beliefs lead to a particular class of equilibria, similar to that considered for the infinite horizon case in the subsection above, and that as α tends to 1, the set of feasible equilibria within this class tend to a singleton. Again, these results lend further support for the earlier focus of the paper on the equilibria characterised in Sections 4 and 5.

7 Conclusion and further work

Returning to the quote from *Othello* at the beginning of this paper, we have found some support for the view that the truth will eventually emerge in contrast to the alternative that luck plays the critical role. Specifically, the central result of this paper is that a perfectly informed seller who knows that her quality is good will always sell. This result is demonstrated in particular settings and the self-confidence mechanism driving the result—that a good seller has more to gain from continuing to trade because her reputation is relatively more likely to be enhanced rather than to deteriorate—might be expected to have some wider application. In particular, self-confidence has been shown to apply both in the finite and the infinite horizon cases, in the case where the seller has information which, though imperfect, is more informative than the buyers' knowledge and in the case where the seller has limited opportunities for strategic behaviour.

An aspect which has been consciously avoided in the discussion above, is the issue of signalling through mechanisms other than a willingness to trade. In particular, the pricing mechanism outlined in the model—an auction amongst many buyers who buy only once—does not allow a

[15] See the Appendix for a specific example.

[16] The proof is similar to the proof of Proposition 11 and so is omitted.

role for price to be used as a signal. Allowing price a signalling role, for example, by allowing the seller to make take it or leave it offers, leads to a greater multiplicity of equilibria. However, the equilibria described here would always be equilibria of such a model and as discussed in Appendix 2, which deals with price signalling, a restriction on equilibria shows that the signalling mechanism described here (signalling through a willingness to trade) and price signalling merely complement each other.

Relaxing the assumption that all outcomes are commonly observable might prove an interesting extension, either by supposing that the seller does not know whether the buyer perceived the good as successful at all, or by supposing that the seller has some noisy information on this. An approach towards considering this problem is to begin by returning first to the case where the seller has no private information.

Finally it is of considerable interest to develop an analogue of a multi-armed bandit framework in which a number of sellers are competing. As Bergemann and Valimaki (1996) and Felli and Harris (1996) do, it seems natural to start with two sellers.

A first step in this direction might be to consider one seller whose type is common knowledge—an incumbent—and a privately informed new entrant, whose quality may be better or worse than the incumbent's. The framework might further need to be altered so that there was a single buyer buying in every period, while the sellers compete to sell a single unit in each period. Intuitively, it seems reasonable to suppose that the mechanism underlying the results derived in this paper could apply in this case as well. In particular, in equilibrium, the price that the incumbent would charge in any period would be predictable. If a bad new entrant were prepared to match the price/anticipated quality offered by the incumbent, the new entrant would make a sale. In contrast, if a bad seller were not prepared to sell with certainty, then following the logic of the sections above, in equilibrium her strategy would be a mixed strategy, which left her indifferent between selling and not selling, though a good seller would always sell. The pricing strategies in this environment may not be trivial to calculate, and, in particular, price signalling would have to be considered. Multiplicity of equilibria would also arise in this context. Nevertheless, following the logic developed in this paper it seems reasonable to believe that there are equilibria in which a good entrant survives, and the approach of Section 6 and the uniqueness result which arose in that context might be applicable here too. Such thinking should also apply to a case with two sellers of privately-known quality, since in the long run the type of one of the sellers will be known arbitrarily accurately, and thus the better of the two sellers should survive.

Appendix

Proof of Proposition 1

Let $B[0,1]$ represent the set of bounded, continuous real-valued functions with domain $[0,1]$. We begin by defining the operator $S : B[0,1] \longrightarrow B[0,1]$ as follows:

$$\begin{aligned} S(f(\lambda)) &= \lambda g + (1-\lambda)b - c + \beta(\lambda g + (1-\lambda)b)\max\{0, f(\lambda^s)\} \\ &\quad +\beta(1-\lambda g - (1-\lambda)b)\max\{0, f(\lambda^f)\} \end{aligned} \tag{A1}$$

where λ^s and λ^f are as defined in Equations (1) and (2) above.

We then proceed as follows, first we prove that it is true that $S : B[0,1] \longrightarrow B[0,1]$, next we prove Blackwell's two sufficiency conditions for a contraction, which ensure that the contraction mapping theorem applies and that there exists a unique solution to the recursive Equation (3).[17] Finally we prove the that S takes increasing functions to increasing functions. That S preserves continuity and is monotonically increasing is sufficient for ensuring that $V^u(\lambda)$ is continuous and increasing. Putting these together we must show that:

1. S takes bounded, continuous real-valued functions with domain $[0,1]$ to bounded continuous real-valued functions with domain $[0,1]$.

2. S satisfies monotonicity, that is for any $f, h \in B[0,1]$ with $f(x) \geq h(x)$ for all $x \in [0,1]$ then $(Sf)(x) \geq (Sh)(x)$ for all $x \in [0,1]$.

3. S satisfies discounting, that is there exists some constant $\alpha \in (0,1)$ such that, $(Sf)(x) + \alpha a \geq (S(f+a))(x)$, for all constants a.

4. If f is such that $f(\lambda) > f(\mu)$ for all $\lambda \geq \mu$, then $(Sf)(\lambda) > (Sf)(\mu)$ for all $\lambda \geq \mu$.

We begin by proving the first point:

Let f be a bounded real-valued continuous function with domain $[0,1]$. That $S : B[0,1] \longrightarrow B[0,1]$ can be trivially verified. To show that (Sf) is continuous, given any $x \in [0,1]$ and $\varepsilon > 0$, there must be some $\delta > 0$ such that for all$\mid x - y \mid < \delta$, $\mid (Sf)(x) - (Sf)(y) \mid < \varepsilon$. Consider first any $z \in [0,1]$, then

$$\begin{aligned} \mid (Sf)(x) - (Sf)(z) \mid = \; & \mid xg + (1-x)b - c + \beta(xg + (1-x)b)\max\{0, f(x^s)\} \\ & +\beta(1 - xg - (1-x)b)\max\{0, f(x^f)\} \\ & -zg - (1-z)b + c - \beta(zg + (1-z)b)\max\{0, f(z^s)\} \\ & -\beta(1 - zg - (1-z)b)\max\{0, f(z^f)\} \mid \end{aligned} \tag{A2}$$

$$\begin{aligned} \mid (Sf)(x) - (Sf)(z) \mid = \; & \mid (x-z)(g-b) \\ & +\beta(zg + (1-z)b)(\max\{0, f(x^s)\} - \max\{0, f(z^s)\}) \\ & +\beta(1 - xg - (1-x)b)(\max\{0, f(x^f)\} - \max\{0, f(z^f)\}) \\ & +\beta(x-z)(g-b)(\max\{0, f(x^s)\} - \max\{0, f(z^f)\} \mid \end{aligned} \tag{A3}$$

[17] See, for example, Ch. 3.2 (pp 49-55) Stokey and Lucas (1989).

$$
\begin{aligned}
\mid (Sf)(x) - (Sf)(z) \mid = \quad & \mid x - z \mid (g-b)+ \\
& + \mid \max\{0, f(x^s)\} - \max\{0, f(z^s)\} \mid \\
& + \mid \max\{0, f(x^f)\} - \max\{0, f(z^f)\} \mid \\
& + \mid x - z \mid (g-b) \mid \max\{0, f(x^s)\} - \max\{0, f(z^f)\} \mid
\end{aligned} \tag{A4}
$$

Well by continuity of f, there exists some δ_1 such that for all $\mid x - z \mid < \delta_1$, $\mid \max\{0, f(x^s)\} - \max\{0, f(z^s)\} \mid < \frac{\varepsilon}{4}$ and $\mid \max\{0, f(x^f)\} - \max\{0, f(z^f)\} \mid < \frac{\varepsilon}{4}$, and by the boundedness of f, $\mid \max\{0, f(x^s)\} - \max\{0, f(z^f)\} \mid < M$ for some M.

So let $\delta = \min\{\delta_1, \frac{\varepsilon}{4}, \frac{\varepsilon}{4M}, 1\}$, then for any $\mid x - y \mid < \delta$

$$
\mid (Sf)(x) - (Sf)(y) \mid < \frac{\varepsilon}{4}(g-b) + \frac{\varepsilon}{4} + \frac{\varepsilon}{4} + \frac{\varepsilon}{4M}M < \varepsilon \tag{A5}
$$

Thus S preserves continuity.

Next we turn to monotonicity. Suppose that $f(x) \geq h(x)$ for all x, then it follows that for all x, $\max\{0, f(\frac{xg}{xg+(1-x)b})\} \geq \max\{0, h(\frac{xg}{xg+(1-x)b})\}$

and $\max\{0, f(\frac{x(1-g)}{x(1-g)+(1-x)(1-b)})\} \geq \max\{0, h(\frac{x(1-g)}{x(1-g)+(1-x)(1-b)})\}$. It follows that $(Sf)(x) \geq (Sh)(x)$.

Then we turn to the discounting condition.

$$
\begin{aligned}
(S(f+a))(x) = \quad & xg + (1-x)b - c \\
& + \beta(xg + (1-x)b) \max\{0, f(\tfrac{xg}{xg+(1-x)b}) + a\} \\
& + \beta(1 - xg - (1-x)b) \max\{0, f(\tfrac{x(1-g)}{x(1-g)+(1-x)(1-b)}) + a\}
\end{aligned} \tag{A6}
$$

So

$$
\begin{aligned}
(S(f+a))(x) \leqslant \quad & \beta a + xg + (1-x)b - c \\
& + \beta(xg + (1-x)b) \max\{0, f(\tfrac{xg}{xg+(1-x)b})\} \\
& + \beta(1 - xg - (1-x)b) \max\{0, f(\tfrac{x(1-g)}{x(1-g)+(1-x)(1-b)})\}
\end{aligned} \tag{A7}
$$

or equivalently $(S(f+a))(x) \leqslant (Sf)(x) + \beta a$ which is precisely the discounting condition since $0 < \beta < 1$.

Finally we must show that S takes strictly increasing functions to strictly increasing functions. Suppose that f is strictly increasing and that $\lambda > \mu$. Consider $(Sf)(\lambda) - (Sf)(\mu)$.

$$
\begin{aligned}
(Sf)(\lambda) - (Sf)(\mu) = \quad & (\lambda - \mu)(g-b) + \beta(\lambda g + (1-\lambda)b) \max\{0, f(\lambda^s)\} \\
& - \beta(\mu g + (1-\mu)b) \max\{0, f(\mu^s)\} \\
& + \beta(1 - \lambda g - (1-\lambda)b) \max\{0, f(\lambda^f)\} \\
& - \beta(1 - \mu g - (1-\mu)b) \max\{0, f(\mu^f)\}
\end{aligned} \tag{A8}
$$

$$(Sf)(\lambda) - (Sf)(\mu) = \begin{array}{c} (\lambda-\mu)(g-b) \\ +\beta(\mu g + (1-\mu)b)(\max\{0, f(\lambda^s)\} - \max\{0, f(\mu^s)\}) \\ +\beta(1-\lambda g - (1-\lambda)b)(\max\{0, f(\lambda^f)\} - \max\{0, f(\mu^f)\}) \\ +\beta(\lambda-\mu)(g-b)(\max\{0, f(\lambda^s)\} - \max\{0, f(\mu^f)\}) \end{array} \tag{A9}$$

Since $\lambda > \mu$, $\lambda^s > \mu^s$ and $\lambda^f > \mu^f$, and so since f is increasing $\max\{0, f(\lambda^s)\} \geq \max\{0, f(\mu^s)\}$ and $\max\{0, f(\lambda^f)\} \geq \max\{0, f(\mu^f)\}$. Furthermore $\lambda^s > \mu^f$, so $\max\{0, f(\lambda^s)\} \geq \max\{0, f(\mu^f)\}$. Hence

$$(Sf)(\lambda) - (Sf)(\mu) \geq (\lambda-\mu)(g-b) > 0 \tag{A10}$$

Thus (Sf) is a strictly increasing function.

Finally note that $V^u(0) = b - c < 0$ and $V^u(1) = \frac{g-c}{1-\beta} > 0$. ■

Proof of Proposition 2

First note that, the proof that $W^b(\lambda)$ is a well-defined, bounded, continuous and strictly increasing function and that there exists a unique λ', for which $W^b(\lambda') = 0$ is analogous to the proof of Proposition 1 and so details are omitted.

We seek to show that the outcome and strategies can be supported as a Perfect Bayesian Equilibrium in Markov stationary strategies. As discussed above, the sale price given interim belief μ will be given by $\mu g + (1-\mu)b$. Next note that all beliefs are consistent with the strategies used and are updated according to Bayes' rule. Suppose that $\lambda < \lambda'$ and the seller continues to trade then, by Bayes' law, the interim belief would be $\frac{\lambda}{\lambda+(1-\lambda)d(\lambda)} = \lambda'$. With the good types always trading and bad types trading according to the strategies given in the statement of the proposition, it follows that in Equations 4 and 5, $\mu = \lambda$ for all $\lambda \geq \lambda'$ and $\mu = \lambda'$ for all $\lambda < \lambda'$. It follows that the value for a bad seller of having a prior reputation λ is $V^b(\lambda) = W^b(\lambda)$ for all $\lambda \geq \lambda'$, $V^b(\lambda) = 0$ otherwise and that letting $V^g(\lambda)$ denote the value to a good seller of having a prior reputation λ, $V^g(\lambda) > V^b(\lambda)$ for all $\lambda \in (0,1)$, and in particular $V^g(\lambda) = V^g(\lambda') > 0$ for all $\lambda < \lambda'$.[18]

Therefore if the prior belief is $\lambda \in (0,1)$, if $\lambda \geq \lambda'$, both good and bad types would want to trade and if the reputation $\lambda < \lambda'$, given that buyers believe that these are the strategies being played, the interim reputation will be λ', so that a bad seller would be indifferent between entering or not and a good seller would want to enter. ■

Proof of Proposition 3

The case where $N = 1$ is trivial. With both good and bad selling according to these strategies, the price given entry is $\lambda g + (1-\lambda)b$, and so both good and bad do not want to deviate from these strategies when $\lambda g + (1-\lambda)b - c \geq 0$. When λ is such that $\lambda g + (1-\lambda)b - c < 0$, this equilibrium can be supported by the off-equilibrium beliefs for the buyers that if a seller does seek to trade, then the belief conditioned on this information is equal to the prior belief λ.

Next consider the case where $N = 2$, following a similar approach to the infinite horizon problem, we can write the value function for a bad seller if buyers believed that both types of sellers always sold in the current period. This is $W_2^b(\lambda)$ which satisfies the following equation.

[18] Formally we should define $V^g(\lambda)$ and $V^b(\lambda)$ using appropriate Bellman equations and show that these equations have unique well-defined solutions. Such details are analogous to the corresponding details in the Proof of Proposition 1 and are therefore omitted.

$$W_2^b(\lambda) = \lambda g + (1-\lambda)b - c + \beta[bV_1^b(\lambda^s) + (1-b)V_1^b(\lambda^f)] \tag{A11}$$

where $V_1^b(\lambda)$ represents the equilibrium value of having reputation λ with one period remaining (so that $V_1^b(\lambda) = \max\{0, \lambda g + (1-\lambda)b - c\}$.

We define λ_2 implicitly as the value that satisfies $W_2^b(\lambda_2) = 0$; notice that λ_2 will be uniquely defined and that $1 > \lambda_2 > 0$. Then, a good seller's strategy always to sell in this period and a bad seller's to sell with probability $d_2(\lambda) = \min\{1, \frac{\lambda(1-\lambda_2)}{(1-\lambda)\lambda_2}\}$ are equilibrium strategies.

Similarly, an inductive argument can show that, there is an equilibrium in which a good seller always sells and a bad seller sells with probability one if her current reputation is greater than or equal to an appropriately defined λ_N.

Finally to see that λ_N is decreasing in N, we prove this inductively, at the same time showing that $V_N^b(\lambda)$ is increasing in N in the range $\lambda > \lambda_N$ and non-decreasing otherwise. These properties can be shown easily for the case $N = 2$, and so we show only the inductive step. Again we define $W_N^b(\lambda)$:

$$W_{N+1}^b(\lambda) = \lambda g + (1-\lambda)b - c + \beta[bV_N^b(\lambda^s) + (1-b)V_N^b(\lambda^f)] \tag{A12}$$

by the assumed properties of $V_N^b(\lambda)$, it follows that

$$\begin{aligned} W_{N+1}^b(\lambda_N) &> \lambda_N g + (1-\lambda_N)b - c + \beta[bV_{N-1}^b(\lambda_N^s) + (1-b)V_N^b(\lambda_N^f)] \\ &\geq \lambda_N g + (1-\lambda_N)b - c + \beta[bV_{N-1}^b(\lambda_N^s) + (1-b)V_{N-1}^b(\lambda_N^f)] \\ &= V_N^b(\lambda_N) = 0 \end{aligned} \tag{A13}$$

Thus $W_{N+1}^b(\lambda_N) > 0$, and since, as can easily be verified, $W_{N+1}^b(\lambda)$ is increasing in λ, it follows that $\lambda_{N+1} < \lambda_N$. Then for $\lambda \in (\lambda_{N+1}, \lambda_N]$ $V_{N+1}^b(\lambda) > 0 = V_N^b(\lambda)$ and for $\lambda > \lambda_N$, again using the assumed inductive properties of $V_N^b(\lambda)$:

$$\begin{aligned} V_{N+1}^b(\lambda) &= \lambda g + (1-\lambda)b - c + \beta[bV_N^b(\lambda^S) + (1-b)V_N^b(\lambda^f)] \\ &> \lambda g + (1-\lambda)b - c + \beta[bV_{N-1}^b(\lambda^S) + (1-b)V_N^b(\lambda^f)] \\ &\geq \lambda g + (1-\lambda)b - c + \beta[bV_{N-1}^b(\lambda^S) + (1-b)V_{N-1}^b(\lambda^f)] \\ &= V_N^b(\lambda) \end{aligned} \tag{A14}$$

so for $\lambda \geq \lambda_{N+1}$, $V_{N+1}^b(\lambda) > V_N^b(\lambda)$, and it can similarly be shown that for $\lambda < \lambda_{N+1}$, $V_{N+1}^b(\lambda) \geq V_N^b(\lambda)$. This concludes the inductive step and the proof. ■

Proof of Proposition 4

The first part of the proposition—that for a given $\gamma > \frac{c}{g}$ there exists a $\lambda'(\gamma) \in (0,1)$ and an equilibrium in which a good seller always sells and in this equilibrium, a bad seller sells with certainty if the prior belief $\lambda \geq \lambda'(\gamma)$ and sells with probability $d(\lambda, \gamma)$ otherwise—follows in an identical fashion to the proof of Proposition 2. The details for this part are therefore omitted.

First to see that $\lambda'(\gamma)$ is decreasing in γ, consider the following Bellman operator S_γ on real-valued functions with domain $[0,1]$, defined as follows:

$$(S_\gamma f)(x) = \gamma[xg + (1-x)g] - c + \beta[b \max\{0, f(x^s)\} + (1-b)\max\{0, f(x^f)\}] \tag{A15}$$

Suppose that $\gamma_1 > \gamma_2$. Then $(S_{\gamma_1} f)(x) > (S_{\gamma_2} f)(x)$ for all $x \in [0,1]$. It follows that $W^b_{\gamma_1}(\lambda) > W^b_{\gamma_2}(\lambda)$ where these are defined in a consistent way to the above notation, that is where $W^b_\gamma(\lambda)$ is the value of a sale to a bad seller if buyers believed that sellers did not stop trading regardless of type and the seller appropriates a proportion γ of consumers' valuation. $\lambda'(\gamma)$ is the unique solution of $W^b_\gamma(\lambda) = 0$, and so since $W^b_{\gamma_1}(\lambda) > W^b_{\gamma_2}(\lambda)$ and these are increasing functions, $\lambda'(\gamma_1) < \lambda'(\gamma_2)$.[19] Hence $\lambda'(\gamma)$ is decreasing in γ.

Finally since $d(\lambda, \gamma) = \min\{1, \frac{\lambda(1-\lambda'(\gamma))}{(1-\lambda)\lambda'(\gamma)}\}$, it is decreasing in $\lambda'(\gamma)$, and $\lambda'(\gamma)$ is decreasing in γ, $d(\lambda, \gamma)$ is increasing in γ. ■

Proof of Proposition 5

We begin by writing down value of making a sale for a seller with a low signal, supposing that buyers believed that the strategies of sellers with high and low signals were identical, so that $r_0 = r_t \forall t$. In this context, the buyers' belief $\lambda_t(r_0, h_t)$ as to the seller's type, and the buyers' belief r_0 as to which signal the buyer received, are sufficient statistics for determining the seller's belief as to her own type $\underline{\lambda}_t(h_t)$, since for all histories h_t, j_t such that $\lambda_t(r_0, h_t) = \lambda_t(r_0, j_t)$ it will be the case that $\underline{\lambda}_t(h_t) = \underline{\lambda}_t(j_t)$. Therefore we can write the seller's belief as a function of λ_t and r_0 and the value of a sale for a seller with a low signal given that buyers believe that she is good with probability λ and that she received a low signal with probability r as follows.

$$\begin{aligned} W^b(r,\lambda) &= \lambda g + (1-\lambda)b - c + \beta(\underline{\lambda}(r,\lambda)g + (1-\underline{\lambda}(r,\lambda)b)\max\{0, W^b(r,\lambda^s)\} \\ &\quad + \beta(1-\underline{\lambda}_t(r,\lambda)g - (1-\underline{\lambda}_t(r,\lambda)b)\max\{0, W^b(r,\lambda^f)\} \end{aligned} \tag{A16}$$

In a similar fashion to the proof of Proposition 1, it can be shown that this equation has a unique well-defined solution which is continuous and increasing in λ. Moreover for each r, there is a unique $\lambda'(r) \in (0,1)$ such that $W^b(r, \lambda'(r)) = 0$. Furthermore, in a similar fashion to the proof of Proposition 4, it can be shown that $W^b(r,\lambda)$ is decreasing in r (the intuition here is clear, for a given λ, the higher the public belief that the seller had a high signal, the lower the self-belief $\underline{\lambda}(r,\lambda)$ of the seller). This in turn implies that $\lambda'(r)$ is increasing in r.

Let s_t be the buyers' interim belief that the seller received a high signal, and as before let μ_t denote the probability (as assessed by buyers) that the seller is good. Now consider the following strategies:

- a seller with a high signal, continues selling with certainty as long as $\overline{\lambda}_t \geq \lambda^u$ and otherwise stops selling for sure;
- a seller with a low signal, continues selling with certainty as long as $\underline{\lambda}_t \geq \lambda'(r_t)$, otherwise she continues selling with probability d such that $\mu_t(d, \lambda_t) = \lambda'(s_t(d))$. If $\mu_t(1, \lambda_t) < \lambda'(s_t(1))$, then she stops selling with probability one.

[19]The existence and uniqueness, continuity and monotonicity of $W^b_\gamma(\lambda)$ can be proved analogously to Proposition 1 and so details are omitted.

Note that on the equilibrium path, by Bayes' law $s_t = \frac{r_t}{r_t+(1-r_t)d}$. On the equilibrium path, beliefs are determined according to Bayes' law and off the equilibrium path, it does not matter what beliefs are assigned, in particular we can assume that if buyers observe a seller selling where $\mu_t(1,\lambda_t) < \lambda^u$, then the buyers believe that the seller is good.[20] It can be verified that these strategies form a Perfect Bayesian Equilibrium, and that this equilibrium satisfies the property outlined in the statement of the Proposition. ■

Proof of Remark 1

Using the notation of the Proof of Proposition 5, we argue that $\lambda'(1) > \lambda^u$. This follows by comparing $W^b(1,\lambda)$ and the value of selling for a seller with a high signal, who is believed to have received a high signal $V^u(\lambda)$. Since $\underline{\lambda}(1,\lambda) < \lambda = \overline{\lambda}(1,\lambda)$, it follows that $W^b(1,\lambda) < V^u(\lambda)$ for all λ and so $\lambda'(1) > \lambda^u$. Thus the range $[\lambda^u, \lambda'(1)]$ is non-empty. ■

Proof of Proposition 6

Note that $U_n(p)$ is an increasing function in p and that we can always define $p_n = \inf\{p : U_n(p) > 0\}$ which is the lowest reputation which will ensure the seller continues trading.

We show that there is an equilibrium which generates the following properties, in addition to the property described in the proposition.

For all (λ, p) in $(0,1)^2$:

1. *$V_m(p,p) \geq U_m(p)$;*
2. *$V_m(\lambda,p)$ is non-decreasing in λ;*
3. *$V_m(\lambda,p)$ is non-decreasing in p; and,*
4. *$V_m(\lambda,p)$ is continuous in λ.*

We generate appropriate equilibrium strategies inductively and show that they are indeed equilibrium strategies and that they do indeed deliver the desired properties.

Well in the last remaining period, we have, irrespective of own belief, sell if $\lambda g+(1-\lambda)b-c \geq 0$, then trivially this is an equilibrium and generates $V_1(\lambda,p) = \max\{0, \lambda g + (1-\lambda)b - c\}$. This is also non-decreasing and continuous in λ and in p, and $V_1(\lambda,p) = U_1(p,p)$.

Inductive step: Then we proceed inductively and consider the nth stage:

Given the reputation level λ, the nth period strategy for each self-belief type(the strategies at the remaining $n-1$ periods are defined inductively) are given as follows:

Suppose that $\lambda g + (1-\lambda)b - c + \beta E[V_{n-1}(\nu,q) \mid \lambda, p] \geq 0$ for all p then the strategy for a seller of any self-belief type p is to sell for sure.

Suppose that $\lambda g + (1-\lambda)b - c + \beta E[V_{n-1}(\nu,q) \mid \lambda, p] < 0$ for some p. Then let the strategies be defined by the following. Let $p_n(\lambda)$ and $d_n(\lambda)$ be implicitly defined by the following: suppose that when the public reputation is λ, all sellers with self-belief $p < p_n(\lambda)$ cease trading with probability 1 and all sellers with self-belief $p > p_n(\lambda)$ continue trading with probability 1, and those sellers with self-belief $p_n(\lambda)$ continue trading with probability $d_n(\lambda)$.[21]

[20] Then *a fortiori* if the off-equilibrium beliefs place less weight on the seller having received a high signal, neither kind of seller would still not want to continue selling.

[21] Note that in the case where self-belief is continuously distributed i.e. no mass points then we can ignore this and set $d_n(\lambda) = 1$ always.

Then such strategies would induce an interim reputation μ which is the belief contingent on the seller's willingness to continue trading, With μ thus defined (and note that μ would be lexicographic increasing in $p_n(\lambda)$ and $d_n(\lambda)$) we can implicitly define $p_n(\lambda)$ and $d_n(\lambda)$ by the equation:

$$\mu g + (1-\mu)b - c + \beta E[V_{n-1}(\nu, q) \mid \mu, p_n(\lambda)] = 0. \tag{A17}$$

Then the proposed strategies would indeed form an equilibrium.

It then only remains to show that (a) such a $p_n(\lambda)$ and $d_n(\lambda)$ exist and that (b) $p_n(\lambda) < p_n = \inf\{p : U_n(p) > 0\}$ and that (c) the other properties of the generated $V_n(\lambda, p)$ hold. We consider each in turn:

(a) Well as mentioned above μ is lexicographic increasing and continuous in $p_n(\lambda)$ and $d_n(\lambda)$.[22] Furthermore we have that $V_{n-1}(\nu, q)$ is non-decreasing in ν; non-decreasing in q and continuous in ν; this combined with the hypothesised FOSD properties imply that $\beta E[V_{n-1}(\nu, q) \mid \mu, p]$ is non-decreasing in μ, and non-decreasing in p; it is also continuous in μ since $f(\nu, q \mid \mu, p)$ is continuous in μ and $V_{n-1}(\nu, q)$ is continuous in ν. By the continuity of these functions it follows that $E[V_{n-1}(\nu, q) \mid \lambda, p]$ is continuous in λ and so $\mu g + (1-\mu)b - c + \beta E[V_{n-1}(\nu, q) \mid \mu, p_n]$ is continuous in μ; using this and the fact that $\lambda g + (1-\lambda)b - c + \beta E[V_{n-1}(\nu, q) \mid \lambda, p] < 0$ and $g - c + \beta E[V_{n-1}(\nu, q) \mid 1, p] > 0$ for all p gives us existence.

(b) We argue by contradiction suppose that this is not the case then it must be that $\mu > p_n$ and:

$$\mu g + (1-\mu)b - c + \beta E[V_{n-1}(\nu, q) \mid \mu, p_n] \leq 0 \tag{A18}$$

but it also follows that

$$\mu g + (1-\mu)b - c + \beta E[V_{n-1}(\nu, q) \mid \mu, p_n] > p_n g + (1-p_n)b - c + \beta E[V_{n-1}(\nu, q) \mid p_n, p_n] \geq 0 \tag{A19}$$

where the first inequality follows by the FOSD property and since $V_{n-1}(\nu, q)$ is non-decreasing in ν and $p_n < \mu$, and the second inequality follows by definition of p_n, so long as $E[V_{n-1}(\nu, q) \mid p_n, p_n] \geq E[U_{n-1}(q) \mid p_n]$. This last inequality follows since we know that $V_{n-1}(q, q) \geq U_{n-1}(q)$, that $E[\nu, q \mid p_n, p_n] = (p_n, p_n)$ and $E[q \mid p_n] = p_n$ and it is easy to verify that $V_{n-1}(\nu, q)$ is convex in ν. Specifically:

$V_1(\nu, q)$ is convex in ν

Suppose $V_{n-1}(\nu, q)$ is convex in ν

then claim $V_n(a\nu + (1-a)\lambda, q) \geq aV_n(\nu, q) + (1-a)V_n(\lambda, q)$

[22] More formally, let the strategy set $S(p_n(\lambda), d_n(\lambda))$ be the set of strategies whereby all seller's with self-belief $p < p_n(\lambda)$ cease trading with probability 1 and all sellers with self-belief $p > p_n(\lambda)$ continue trading with probability 1, and those sellers with self-belief $p_n(\lambda)$ continue trading with probability $d_n(\lambda)$. Then μ is a function of this strategy set, and if we consider the composition of μ as a function of S and S as a function of $p_n(\lambda)$ and $d_n(\lambda)$, then we can think of μ as a function of $p_n(\lambda)$ and $d_n(\lambda)$. Then $\mu(p, d)$ is increasing in p and in d; and if $p > q$ then $\mu(p, d) > \mu(q, e)$ for any d and e, and any value of $\mu \in [\lambda, 1)$ is attainable through appropriate choice of p and d.

Well this follows so long as $a\mu_n(\nu) + (1-a)\mu_n(\lambda)) \geq \mu_n(a\nu + (1-a)\lambda)$

Well since $\mu_n(\nu)$ is continuous and a constant for ν low and then linear increasing, it is convex.

(c) Property 2 follows directly and the remaining properties follow from the observation that for all λ's which induce at least one self-belief type to cease trading with some non-zero probability all give rise to the same interim belief μ. ■

Proof of Lemma 1

We argue by contradiction, so suppose that given any ε there is some $\lambda < \varepsilon$ such that either good or bad (or both) will continue trading if given the opportunity with some probability. However, since trivially an upper bound for the interim belief in the current period is $\frac{\lambda}{\lambda+(1-\lambda)(1-\alpha)}$ (this is the belief contingent on all good sellers continuing to trade and all the bad ones who have the opportunity ceasing), so an upper bound for the prior belief in the next period is $\frac{\lambda g}{\lambda g+(1-\lambda)(1-\alpha)b}$ and so an upper bound for the interim belief in the next period is

$$\frac{\frac{\lambda g}{\lambda g+(1-\lambda)(1-\alpha)b}}{\frac{\lambda g}{\lambda g+(1-\lambda)(1-\alpha)b} + \frac{(1-\lambda)(1-\alpha)b}{\lambda g+(1-\lambda)(1-\alpha)b}(1-\alpha)b} = \frac{\lambda g}{\lambda g + (1-\lambda)(1-\alpha)^2 b^2} \tag{A20}$$

and an upper bound for the prior belief in the period after is $\frac{\lambda g^2}{\lambda g^2+(1-\lambda)(1-\alpha)^3 b^3}$ and so on. Since these upper bounds on beliefs are increasing in time, so it follows that if it's worth staying in now it always will be with belief increasing in this way, so it must be that the upper bound for the total reward of making a sale is bounded by

$\frac{\lambda}{\lambda+(1-\lambda)(1-\alpha)} g + \frac{(1-\lambda)(1-\alpha)}{\lambda+(1-\lambda)(1-\alpha)} b - c +$
$\sum_{n=1}^{\infty} \beta^n \left(\frac{\lambda g^n}{\lambda g^n+(1-\lambda)(1-\alpha)^{n+1} b^{n+1}} g + \frac{(1-\lambda)(1-\alpha)^{n+1} b^{n+1}}{\lambda g^n+(1-\lambda)(1-\alpha)^{n+1} b^{n+1}} b - c\right) > 0$

but the left hand side of the inequality is strictly less than

$\frac{\lambda}{\lambda+(1-\lambda)(1-\alpha)} g + b - c + \sum_{n=1}^{\infty} \beta^n \left(\frac{\lambda g^n}{\lambda g^n+(1-\lambda)(1-\alpha)^{n+1} b^{n+1}} g + b - c\right)$
$< \frac{\lambda}{(1-\lambda)(1-\alpha)} g + b - c + \sum_{n=1}^{\infty} \beta^n \left(\frac{\lambda g^n}{(1-\lambda)(1-\alpha)^{n+1} b^{n+1}} g + b - c\right)$
$< \frac{\lambda}{(1-\lambda)(1-\alpha)} g + b - c + \frac{\lambda}{(1-\lambda)(1-\alpha)} g \frac{(1-\alpha)b}{(1-\alpha)b - \beta g} - (c-b)\frac{1}{1-\beta}$

for λ low enough this is negative which is the contradiction that we needed. ■

Proof of Lemma 2:

We define the following value functions which depend on the buyers' interim beliefs

$$\begin{aligned} V_\alpha^g(\lambda;\mu(.)) = \; & g\mu(\lambda) + (1-\mu(\lambda))b - c \\ & +\beta[g\alpha \max\{0, V_\alpha^g(\mu(\lambda)^s;\mu(.))\} + g(1-\alpha)V_\alpha^g(\mu(\lambda)^s;\mu(.)) \\ & +(1-g)\alpha \max\{0, V_\alpha^g(\mu(\lambda)^f;\mu(.))\} + (1-g)(1-\alpha)V_\alpha^g(\mu(\lambda)^f;\mu(.))] \end{aligned}$$

and

$$\begin{aligned} V_\alpha^b(\lambda;\mu(.)) = \; & g\mu(\lambda) + (1-\mu(\lambda))b - c \\ & +\beta[b\alpha \max\{0, V_\alpha^b(\mu(\lambda)^s;\mu(.))\} + b(1-\alpha)V_\alpha^b(\mu(\lambda)^s;\mu(.)) \\ & +(1-b)\alpha \max\{0, V_\alpha^b(\mu(\lambda)^f;\mu(.))\} + (1-b)(1-\alpha)V_\alpha^b(\mu(\lambda)^f;\mu(.))] \end{aligned}$$

Associated with these value functions, we can define Bellman operators G_μ and B_μ by

$$G_\mu(h)(\lambda) = \begin{array}{c} g\mu(\lambda)+(1-\mu(\lambda))b-c \\ +\beta[g\alpha\max\{0,h(\mu(\lambda)^s)\}+g(1-\alpha)h(\mu(\lambda)^s) \\ +(1-g)\alpha\max\{0,h(\mu(\lambda)^f)\}+(1-g)(1-\alpha)h(\mu(\lambda)^f)] \end{array}$$

and

$$B_\mu(h)(\lambda) = \begin{array}{c} g\mu(\lambda)+(1-\mu(\lambda))b-c \\ +\beta[b\alpha\max\{0,h(\mu(\lambda)^s)\}+b(1-\alpha)h(\mu(\lambda)^s) \\ +(1-b)\alpha\max\{0,h(\mu(\lambda)^f)\}+(1-b)(1-\alpha)h(\mu(\lambda)^f)] \end{array}$$

Then it can readily be verified that both operators take non-decreasing functions to non-decreasing functions, and that for any non-decreasing function h, $G_\mu(h) \geq B_\mu(h)$. It follows that $V^g(\lambda \mid \mu(.))$, and $V^b(\lambda \mid \mu(.))$, which following standard recursive techniques are uniquely defined, must be non-decreasing functions and $V^g(\lambda \mid \mu(.)) \geq V^b(\lambda \mid \mu(.))$.

These facts, together with Lemma 1 imply that at least in a bottom range the value of having a particular reputation is strictly greater than the value for a bad type. That the value for a good type for a good type is strictly greater than the value for a bad type in some range $(h,1)$ and is never lower than the value for a bad type, and given that with some non-zero probability given any starting belief λ, the reputation will some stage enter $(h,1)$, it follows that the value for a good type is strictly greater than the value of a bad type for all reputations in $(0,1)$.[23] This and the previously established result that the values of both good and bad types are non-decreasing in the prior reputations imply that any equilibrium must be of the hypothesised form. ■

Proof of Proposition 7:

We begin by introducing the following notation:

$$V^g_\alpha(\lambda;l,m) = \begin{array}{c} g\mu(l,m)+(1-\mu(l,m))b-c \\ +\beta[g\alpha\max\{0,V^g_\alpha(\mu(l,m)^s)\}+g(1-\alpha)V^g_\alpha(\mu(l,m)^s) \\ +(1-g)\alpha\max\{0,V^g_\alpha(\mu(l,m)^f)\}+(1-g)(1-\alpha)V^g_\alpha(\mu(l,m)^f)] \end{array} \tag{A21}$$

and

$$V^b_\alpha(\lambda;l,m) = \begin{array}{c} g\mu(l,m)+(1-\mu(l,m))b-c \\ +\beta[b\alpha\max\{0,V^b_\alpha(\mu(l,m)^s)\}+b(1-\alpha)V^b_\alpha(\mu(l,m)^s) \\ +(1-b)\alpha\max\{0,V^b_\alpha(\mu(l,m)^f)\}+(1-b)(1-\alpha)V^b_\alpha(\mu(l,m)^f)] \end{array} \tag{A22}$$

to represent the value to a good and bad type respectively of selling when the prior reputation is λ, and the interim reputation is given by:

$\mu = \lambda$ for $l < \lambda$

$\mu = \frac{\lambda}{\lambda+(1-\lambda)(1-\alpha)}$ for $l \leq \lambda < m$

[23] Suppose that it is not the case that there is some reputation such that the reputation will never enter $(h,1)$ well it must be that this initial reputation is strictly less than m. Let r be the highest such reputation, then it must be that $\mu(r)^s \leq r$. In particular it must be that $\mu(r) < r$ but by definition of $\mu(.)$ this cannot be.

$\mu = h$ for $m \leq \lambda < h$
$\mu = \lambda$ for $h \leq \lambda$
where $h = \frac{m}{m+(1-m)(1-\alpha)}$.

For an equilibrium to exist necessary and sufficient conditions are that there are values $0 \leq l \leq m \leq 1$, for which the following two conditions holds simultaneously:

(i) $V_{\alpha}^{g}(l; l, m) = 0$
(ii) $V_{\alpha}^{b}(m; l, m) = 0$

Note in particular, that if condition (ii) holds, it will also to follow that $V_{\alpha}^{b}(\lambda; l, m) = 0$ for all $\lambda \epsilon (m, h]$.

Therefore, since these value functions can easily shown to be uniquely defined, continuous and increasing in λ and that $V_{\alpha}^{g}(\lambda; l, m) > V_{\alpha}^{b}(\lambda; l, m)$ using the standard recursive techniques applied, for example, in the Proof of Proposition 1, it follows that the neither the good nor the bad seller would have an incentive to deviate from their hypothesised strategies and this is indeed an equilibrium. Moreover, note that standard techniques can be used to show that the value functions are continuous in l and m.

Now from Lemma 1, it must be the true that any solutions for (l, m) will be interior. It remains to show that such solutions exist.

We argue as follows, first note that trivially, $V_{\alpha}^{g}(0; l, m) < 0$ for all l and m and $V_{\alpha}^{g}(1; 1, 1) > 0$. Thus by the continuity of $V_{\alpha}^{g}(\lambda; l, m)$, the set $\{\lambda : V_{\alpha}^{g}(\lambda; \lambda, \lambda) = 0\}$ is non-empty, let $\bar{l} = \sup\{\lambda : V_{\alpha}^{g}(\lambda; \lambda, \lambda) = 0\}$. Then for any $m > \bar{l}$, since $V_{\alpha}^{g}(1; 1, 1) > 0$, it must be that $V_{\alpha}^{g}(m; m, m) > 0$. Since, $V_{\alpha}^{g}(0; 0, m) < 0$, $V_{\alpha}^{g}(l; l, m)$ is continuous in l and $V_{\alpha}^{g}(m; m, m) > 0$, it follows that the set $\{\lambda < m : V_{\alpha}^{g}(\lambda; \lambda, m) = 0\}$ is non-empty and we can implicitly define $l(m)$ by $l(m) = \inf\{\lambda < m : V_{\alpha}^{g}(\lambda; \lambda, m) = 0\}$ and $l(m)$ is continuous in m.

Then since $V_{\alpha}^{g}(\bar{l}; \bar{l}, \bar{l}) = 0$, and $\bar{l}$ is interior, it follows that $V_{\alpha}^{b}(\bar{l}; \bar{l}, \bar{l}) < 0$; in addition, $V_{\alpha}^{b}(1; l, 1) > 0$ for all l so by the continuity of $V_{\alpha}^{b}(\lambda; l, h)$ and of $l(m)$ it follows that there exists $m > \bar{l}$ such that $V_{\alpha}^{b}(m; l(m), m) = 0$.

Thus there exist m and $l(m)$ with $V_{\alpha}^{b}(m; l(m), m) = 0$ and $V_{\alpha}^{g}(l(m); l(m), m) = 0$—thus conditions (i) and (ii) are satisfied simultaneously, which is what is required for the existence of an equilibrium. ■

Proof of Proposition 8:

The bulk of the proof consists of showing that as α tends to 1, the set of feasible l tends to the singleton $\{0\}$, this in turn implies a singleton feasible m consistent with equilibrium. It then follows that as α tends to 1, the unique equilibrium outcome has $k = l = 0$ and $m = 0$, this is precisely the equilibrium characterised in Proposition 2.

First, by Lemma 1, above, it is clear that for any $\alpha \epsilon (0, 1)$ a lower bound for $l(\alpha)$ is 0. We claim that in the limit 0 is also an upper bound for $l(\alpha)$. We argue by contradiction, suppose the claim is false, then there must be some $\varepsilon > 0$ such that there is no $\overline{\alpha}$ with $l(\alpha) < \varepsilon$ for all $\alpha \in (\overline{\alpha}, 1)$. Now note that the value of having prior reputation λ is always bounded below by $(1-\alpha)\frac{b-c}{1-\beta(1-\alpha)} + \alpha.0$—this is the probability of not having the opportunity to stop in the current period multiplied by the value of having buyers convinced that the seller is bad from now to eternity. Thus it follows that an upper bound for $l(\alpha)$ would be a λ satisfying:

$$\frac{\lambda}{\lambda + (1-\alpha)(1-\lambda)} g + \frac{(1-\alpha)(1-\lambda)}{\lambda + (1-\alpha)(1-\lambda)} b - c - (1-\alpha)\frac{c-b}{1-\beta(1-\alpha)} = 0 \tag{A23}$$

but trivially as α tends to 1, it is clear that the λ solving this equation tends to 0. This is the required contradiction which completes the proof. ■

Proof of Lemma 3:

Assumption 3 implies that in the final period, good and bad sellers will choose to trade if and only if $\lambda > \frac{c-b}{g-b}$. This implies that in the penultimate period the value of *ending* the period with a reputation λ is strictly increasing in λ.[24] This, together with Assumption 2, implies that an equilibrium must be of the hypothesised form and such that for both good and bad types the value of having a reputation λ is non-decreasing in λ and is strictly greater for a good type than a bad type for $\lambda \in (0, 1)$.

Similarly, by induction, when there are n remaining periods, any equilibrium must be of the hypothesised form such that for both good and bad types the value of having a reputation λ is non-decreasing in λ and is strictly greater for a good type than a bad type for $\lambda \in (0, 1)$. ■

Proof of Proposition 9: In the last period, there is an equilibrium of this form where $l_1 = m_1 = h_1 = \frac{g-c}{c-b}$. Then we can proceed by constructing the equilibrium, proceeding by backward induction. We show the inductive step.

For λ low enough neither type wants to continue trading, this follows from Lemma 1 and earlier discussion. Then if beliefs are as optimistic as possible they will be that only a good type seller would continue trading if given the opportunity. In some range, even with such generous beliefs, a bad type won't want to continue trading but a good type will, we suppose that a good type will continue trading but a bad type will not. That such a range exists follows since the next period value is increasing in λ for both types and is not lower for good type than for bad type. Moreover it is negative for low enough reputations and positive for high enough reputations—these properties hold for the final period value function and the backward induction described here preserves these properties.

Then at some level with the most generous beliefs, as described above, a bad type would want to trade, and so, following a similar argument to the paragraph above, the equilibrium would require that a bad seller will employ mixed strategies, continuing to trade with a probability that ensures that the interim reputation is at a level where she is indifferent between continuing or ceasing, and the good type would continue trading. Then at some point even if buyers believed that bad types would continue trading with certainty, bad types would seek to trade—this point characterises the lower bound, h_n, of the highest range. ■

Proof that a multiplicity of equilibria of the form described in Proposition can arise in the 2-period case

Note that the value of ending the second period with reputation λ would be

$$\beta V_1(\lambda, \alpha) = \beta[\alpha \max\{0, \lambda g + (1-\lambda)b - c\} + (1-\alpha)(\lambda g + (1-\lambda)b - c)\} \tag{A24}$$

Therefore the values of trading for a good seller and bad seller respectively, when beginning the period with reputation λ are given by:

$$\mu(\lambda)g + (1-\mu(\lambda))b - c + g\beta V_1(\mu(\lambda)^s) + (1-g)\beta V_1(\mu(\lambda)^f) \tag{A25}$$

[24] Note that this follows, since even if the seller would choose to cease trading, she may not have the opportunity and may be compelled to continue trading.

$$\mu(\lambda)g + (1 - \mu(\lambda))b - c + b\beta V_1(\mu(\lambda)^s) + (1 - b)\beta V_1(\mu(\lambda)^f) \tag{A26}$$

It can be readily verified that there is a multiplicity of equilibria where $l_2 \in [\underline{l}, \max\{\bar{l}, m_2\}]$, and m_2 and h_2 satisfy the following:

$$\underline{l}g + (1 - \underline{l})b - c + g\beta V_1(\underline{l}^s) + (1 - g)\beta V_1(\underline{l}^f) = 0 \tag{A27}$$

$$\bar{l} = \frac{\underline{l}}{\underline{l} + (1 - \underline{l})(1 - \alpha)} \tag{A28}$$

$$h_2 g + (1 - h_2)b - c + b\beta V_1(h_2^s) + (1 - b)\beta V_1(h_2^f) = 0 \tag{A29}$$

$$h_2 = \frac{m_2}{m_2 + (1 - m_2)(1 - \alpha)}. \tag{A30}$$

Note that the first of these equations characterise the prior reputation, $\underline{l}$, at which if the buyers beliefs were pessimistic that is they believe that a good and bad type seller behave identically (both cease trading where possible) then a good seller would still be willing to trade. $\bar{l}$ is the reputation at which a good seller would be willing to trade if buyers were as optimistic as possible. In order to show that there is indeed a multiplicity of equilibria in this two-period case, it is sufficient to show that $\underline{l} < m_2$. It can be easily verified that for the parameter values, $b = 0$, $g = 0.7$, $c = 0.5$, $\beta = 0.9$, $\alpha = 0.5$ that $\underline{l} = 0.59$, $h_2 = 0.79$ and $m_2 = 0.66$ so that $\underline{l} < m_2$ and a multiplicity does indeed arise.

Appendix 2: Price as a signalling device

In Sections 3 and 4 above, the potential signalling role of price has been suppressed through the imposition of the assumption that competition among buyers and the sale procedure ensure that the price of the good is always the full buyers' valuation. Introducing this simplifying assumption allowed us to focus on the signalling role of self-confidence, in the specific sense in this paper of continuing to trade despite short-term losses. In this section, we relax this assumption and modify the model presented in Section 2 by supposing that if the seller chooses to trade, she makes a take-it-or-leave-it offer to one of the potential buyers.

In this context, in addition to deciding whether or not to continue trading, the seller also chooses a sale price (up to the buyers' valuation which is the most that a buyer would be prepared to pay). Thus both the decision to continue trading and the sale price can have signalling roles. In this environment, a plethora of equilibria are possible. This follows since the seller's action space, or the scope of her decision-making, increases significantly, and this can allow off-equilibrium beliefs a more critical role. Specific off-equilibrium beliefs can admit many equilibria. Two extreme situations are of particular interest and highlight this observation.

First, in a no-trade equilibrium, buyers believe that any seller who trades and offers a sale price of more than or equal to b is bad. The strategies of the good and bad types of seller are identical—they both cease trading.

In the efficient trade equilibrium, the buyers believe that any seller who trades and offers any price other than c is bad and that a seller who trades and offers a price c is good. The strategy of a good seller (who is indifferent between selling and not selling) is to sell at c in every period and the strategy of the bad type (who is similarly indifferent) is to cease trading.

These two equilibria point to two factors that lead to multiple equilibria and that are in contrast to the equilibria characterised in Section 4. In particular, in those equilibria there is always trade when trade is efficient and trade is at a price such that the seller extracts the full surplus from trading in that period. In general, neither of these two properties need hold for an equilibrium in the scenario where the seller makes take-it-or-leave-it offers. In particular, no trade can be supported, as demonstrated in the no-trade equilibrium, and if trade occurs, then the price need not be the one that gives the seller all gains from trade—this merely defines a maximum price at which trade could occur.[25]

It is worth noting that the equilibria described in Section 4 where price has no signalling role can also be sustained as equilibria in the case where price does have a signalling role. For example, consider the equilibrium described in Proposition 2. Trivially, this can be sustained as an equilibrium in the case where price has a signalling role by assuming that off-equilibrium beliefs are such that buyers believe that an offer at a price other than that specified in Proposition 2 indicates that the seller is bad. However, an appealing feature of this equilibrium (and a consequence of the fact that it is an equilibrium where price has no signalling role) is that the off-equilibrium beliefs need not be so severe. For example, the same equilibrium would be maintained if buyers faced with an off-equilibrium choice of price do not change their beliefs—that is they maintain the same belief as they held before observing the off-equilibrium price.

Appendix 2.1 The seller's preferred equilibrium

In the spirit of the earlier sections, in which the seller extracted the full consumer surplus in each period, in this section we seek to characterise the equilibrium that a good seller would choose if she were able to. Note that if the seller who had the opportunity to choose which equilibrium to play did not choose one preferred by a good type of seller, then she would reveal herself to be a bad seller. So the focus on the equilibrium outcome preferred by a good seller seems appropriate.

Two important observations are first, and as discussed above, that the outcomes characterised in Section 4 can be sustained as equilibrium outcomes where the seller makes take-it-or-leave-it offers to buyers, and secondly, that these outcomes may not be the ones preferred by a good seller. In particular, and as described below, it is the case that a good seller would strictly prefer an equilibrium which at some low levels of reputations, entailed a price below the full consumer valuation. Such a price would increase the probability that a bad seller ceases trading and so

[25] This is clear from considering the one-period game. For a given belief μ, the maximum price at which trade can occur is given by $\mu g + (1-\mu)b$. However, a price p below this level can be sustained in equilibrium if buyers believe that the strategies of both good and bad types are to sell at p and off-equilibrium beliefs are that any higher offer must be from a bad seller.

might enhance the seller's future prospects and result in a higher total expected value for the seller, this is a kind of costly reputation-building. This intuition is formalised below.

We consider a finite horizon problem, and suppose that in each period the seller can impose her preferred equilibrium, but is constrained in as much as the strategies that she proposes must be dynamically consistent (that is if she has N periods to go, then she can only impose her preferred equilibrium among the set of equilibria which entail preferred equilibria in all $n < N$ period subgames). The intuition behind this restriction is that if at every period the seller can propose a new equilibrium, for such proposals to be credible, they must be dynamically consistent. This restriction is close in spirit to the refinement that strategies in repeated or dynamic games should be renegotiation proof, and leads to an analogous definition.[26]

Definition 1 *In the 1 period game, a **consistent preferred equilibrium** is an equilibrium which yields greatest profits to the good seller, suppose that this value is $Y_1^g(\lambda)$. Then for any $N > 1$, suppose that $\{Q^N(\lambda)\}$ denotes the set of expected payoffs that are feasible in an equilibrium in which the continuation payoffs for a good adviser are given by $Y_{N-1}^g(.)$, then we define $Y_N^g(\lambda) = \max\{Q^N(\lambda)\}$. An equilibrium is in the N-period game is a **consistent preferred equilibrium** (CPE) if for every $n \leqslant N$, each n-period subgame with the prior buyers' belief λ, yields an expected value to a good seller of $Y_n^g(\lambda)$.*

We prove in the appendix that working with this definition and proceeding by backward induction yields a unique equilibrium outcome, which is similar to that of Proposition 3 in that a good seller always sells except possibly in the last period, and a bad seller either sells with certainty or else stops selling with some non-zero probability.

Proposition 11 *There exist constants p_n and $\widehat{\lambda}_n$ such that in the N-period game there is a unique CPE outcome, in which when there are n remaining periods and the seller's reputation at the beginning of the period is λ, then if $\lambda < \widehat{\lambda}_n$ a good seller sells with certainty, a bad seller sells with some probability and the sale price is p_n, and if $\lambda \geq \widehat{\lambda}_n$ then both a good and bad type would continue selling and the sale price is $\lambda g + (1-\lambda)b$.*

However, this consistent preferred equilibrium outcome differs from the one characterised in Proposition 3, in particular it is not always the case that the seller extracts the full consumer surplus, in particular it may be that $p_n < \widehat{\lambda}_n g + (1-\widehat{\lambda}_n)b$. This is because a good seller may prefer to build her reputation in the sense discussed above. An interesting consequence is that where $Y_n^b(\lambda)$ is the expected value in the CPE outcome for a bad type with n periods remaining and a reputation λ, then for some reputation levels, $Y_n^b(\lambda)$ may decrease in n, the number of remaining periods.[27] Though it may at first seem counter-intuitive that at some reputation levels a bad seller might prefer that there were fewer trading opportunities remaining, this result does have some intuitive appeal. First recall that a bad seller must mimic the behaviour of a good seller in order to continue trading and in order that she does not reveal herself to be bad. A good seller effectively has two potentially reasonable courses of action in any period. She can "milk" her reputation—that is charge as much as the buyer is prepared to pay—though notice that even

[26]For definitions of renegotiation-proofness, see for example Benoit and Krishna (1993), Fudenberg and Tirole (1991), Bernheim and Ray (1989) and Farrell and Maskin (1989).

[27]For example, at the parameter values $g = 0.8$, $b = 0.3$, $c = 0.7$ and $\beta = 0.9$, $Y_1^b(0.42) = 0.01$, but $Y_2^b(0.42) = 0$.

when milking her reputation, the seller's reputation will rise following a success. Alternatively, she can build up her reputation by charging a lower price but one at which in equilibrium a bad seller would be more likely to drop out. When there are more periods remaining, reputation-building might be more attractive for a good adviser than reputation-milking, whereas a bad seller would always prefer to be milking her reputation—any reputation building that a good seller would do is credible only in that it makes continuing to trade less attractive to a bad seller. Thus a bad seller might prefer that there are fewer remaining trading opportunities.

Appendix 2.2 Proof of Proposition 10

We make the following claims for the value for a good and bad seller of having a reputation λ with $n \leqslant N$ remaining trading opportunities in the CPE outcome:

- $Y_n^g(\lambda)$ is (strictly) increasing in n and (weakly) increasing and continuous in λ;
- $Y_n^b(\lambda)$ is (weakly) increasing and continuous in λ;
- There exist $\widehat{\lambda}_n$, p_n and $\underline{Y}_n^g$ such that:
  - for all $\lambda < \widehat{\lambda}_n$, $Y_n^g(\lambda) = \underline{Y}_n^g$, $Y_n^b(\lambda) = 0$ and the sale price in the current period (if sale occurs) is p_n.
  - for all $\lambda \geq \widehat{\lambda}_n$, then

$$Y_n^g(\lambda) = \lambda g + (1-\lambda)b - c + \beta g Y_{n-1}^g(\lambda^s) + \beta(1-g)Y_{n-1}^g(\lambda^f) \tag{7}$$

$$Y_n^b(\lambda) = \lambda g + (1-\lambda)b - c + \beta b Y_{n-1}^b(\lambda^s) + \beta(1-b)Y_{n-1}^b(\lambda^f) \tag{8}$$

Note that below we show that $\underline{Y}_1^g = 0$, so since we will show that $Y_n^g(\lambda)$ is strictly increasing in n and weakly increasing in λ, it follows that $Y_n^g(\lambda) > 0$ for $n > 1$, this can only be possible, under the equilibrium strategies derived below, if a good seller always trades in any period $n > 1$.

We prove these claims inductively and hence prove the proposition.

n=1

In the one period case, the problem is trivial. In the CPE the current period price is the maximum possible $\lambda g + (1-\lambda)b$ and there is trade so long as this realises non-negative profits, that is so long as $\lambda \geq \widehat{\lambda}_1 = \frac{c-b}{g-b}$, here $\underline{Y}_1^g = 0$, p_1 is arbitrary and $Y_0^g(\lambda) = Y_0^b(\lambda) = 0 \forall \lambda$.

Inductive step

Suppose that the claims hold for all $m-1 < N$. For now we assume $Y_m^g(\lambda) - Y_{m-1}^g(\lambda) > 0$, which as discussed above implies that a good seller always trades—we restate this claim and prove it below as Lemma 4. Then, for any prior reputation λ, there are two possible equilibrium outcomes in the period $m+1$: where a bad seller always wants to sell, or where a bad seller employs mixed strategies.[28] We consider each of these two possibilities separately.

[28] Note that it cannot be the case that in equilibrium a good seller wants to trade for sure and a bad seller wants to cease trading for sure, since in this situation a bad seller who did trade would be mistaken for a good seller. So she would obtain the same value from trading as a good seller (the current and future belief of buyers would be that the seller was good with certainty). However, since $Y_1^g(\lambda) \geq 0$ and by Lemma 8 $Y_n^g(\lambda) > Y_1^g(\lambda)$, the value for a good seller is always strictly positive and so the bad seller would want to deviate and obtain this positive value.

First consider the case that the bad type always enters, then the value for a good seller is given by:

$$p - c + \beta g Y^g_{m-1}(\lambda^s) + \beta(1-g)Y^g_{m-1}(\lambda^f) \tag{9}$$

in this expression the current period p is feasible if and only if $p \leqslant \lambda g + (1-\lambda)b$ and the continuation payoffs in the next period, by the definition of the CPE, are given by $Y^g_{m-1}(\lambda^s)$ and $Y^g_{m-1}(\lambda^f)$, and the future beliefs must be λ^s and λ^f for consistency with the hypothesised equilibrium strategies that both good and bad sellers always sell. Furthermore, for this to be an equilibrium, it must be true that[29]

$$p - c + \beta b Y^b_{m-1}(\lambda^s) + \beta(1-b)Y^b_{m-1}(\lambda^f) \geq 0 \tag{10}$$

This constraint ensures that it is rational for a bad seller to sell with probability one. Since $p - c + \beta g Y^g_{m-1}(\lambda^s) + \beta(1-g)Y^g_{m-1}(\lambda^f)$ is increasing in p, its maximal value is attained when p is at its maximal feasible value—that is when $p = \lambda g + (1-\lambda)b$—and so in any outcome in which the bad seller sells with probability one, this will be the sale price.

We define $\underline{\lambda}_m$ implicitly by the following equation[30]

$$\underline{\lambda}_m g + (1-\underline{\lambda}_m)b + \beta b Y^b_{m-1}(\underline{\lambda}^s_m) + \beta(1-b)Y^b_{m-1}(\underline{\lambda}^f_m) = 0 \tag{11}$$

Then since $Y^b_{m-1}(\lambda)$ is increasing in λ, for any $\lambda \geq \underline{\lambda}_m$, such an equilibrium outcome (in which a bad seller would sell with probability one) is feasible.

Now, suppose that in the equilibrium, a bad seller would employ a mixed strategy, then the value that a good seller can obtain is given by

$$p - c + \beta g Y^g_{m-1}(\mu^s) + \beta(1-g)Y^g_{m-1}(\mu^f) \tag{12}$$

where μ is the interim reputation, so that if the strategy for a bad seller is to trade with probability d, then $\mu = \frac{\lambda}{\lambda + (1-\lambda)d}$, so that any $\mu \in [\lambda, 1)$ is attainable and for a current period price p to be feasible, it must be that $p \leqslant \mu g + (1-\mu)b$. Furthermore, for employing mixed strategies to be rational for a bad seller, it must be the case that

$$p - c + \beta b Y^b_{m-1}(\mu^s) + \beta(1-b)Y^b_{m-1}(\mu^f) = 0 \tag{13}$$

[29]It must also be the case that
$p - c + \beta g Y^g_{n-1}(\lambda^s) + \beta(1-g)Y^g_{n-1}(\lambda^f) \geq 0$, this can be verified.

[30]Note we abuse notation and $\underline{\lambda}_m$ here is different from $\underline{\lambda}_t$ which appears in Section 5.

Consider the problem of maximising the expression in (12) subject to the constraint (13) and $p \leqslant \mu g + (1-\mu)b$. Substituting for p from (13), the problem can be recast as the maximisation of

$$\beta g Y^g_{m-1}(\mu^s) + \beta(1-g)Y^g_{m-1}(\mu^f) - \beta b Y^b_{m-1}(\mu^s) - \beta(1-b)Y^b_{m-1}(\mu^f) \tag{14}$$

subject to the relevant constraints. This will have a unique solution $\underline{Y}^g_m$ at μ_m, with an associated p_m. Furthermore, this value is attainable from any prior reputation $\lambda \leqslant \mu$ by an appropriate choice of d. Furthermore $\mu_m \geq \underline{\lambda}_m$, since by (13) and the fact that $p_m \leqslant \mu_m g + (1-\mu_m)b$ and the assumed property that $Y^b_{m-1}(\lambda)$ is increasing in λ, it follows that

$$\mu_m g + (1-\mu_m)b + \beta b Y^b_{m-1}(\mu^s_m) + \beta(1-b)Y^b_{m-1}(\mu^f_m) \geq 0 \tag{15}$$

Moreover since $\underline{\lambda}_m$ and $p = \underline{\lambda}_m g + (1-\underline{\lambda}_m)b$ satisfy the constraint (13), it follows that

$$\underline{Y}^g_m \geq \underline{\lambda}_m g + (1-\underline{\lambda}_m)b - c + \beta g Y^g_{m-1}(\underline{\lambda}^s_m) + \beta(1-g)Y^g_{m-1}(\underline{\lambda}^f_m) \tag{16}$$

From the assumed properties of $Y^g_{m-1}(\lambda)$, $\lambda g + (1-\lambda)b - c + \beta b Y^b_{m-1}(\lambda^s) + \beta(1-b)Y^b_{m-1}(\lambda^f)$ is continuous and increasing in λ and so there exists a $\widehat{\lambda}_m \geq \underline{\lambda}_m$ such that $\underline{Y}^g_m = \widehat{\lambda}_m g + (1-\widehat{\lambda}_m)b + \beta g Y^g_{m-1}(\widehat{\lambda}^s_m) + \beta(1-g)Y^g_{m-1}(\widehat{\lambda}^f_m)$.[31]

Finally it is a simple exercise to show that there is a CPE with $\widehat{\lambda}_m$, p_m and $\underline{Y}^g_m$ as defined above and with both good and bad types always selling at a price that extracts full consumer surplus for $\lambda \geq \widehat{\lambda}_m$, and otherwise at a sale price of p_m with good types always selling and bad types selling with a probability that induces the interim reputation μ_m and buyers beliefs such that any off-equilibrium action must have been taken by a bad type, and moreover that the appropriate hypothesised inductive properties do hold. ■

Lemma 4 $Y^g_m(\lambda) - Y^g_{m-1}(\lambda) > 0$

Proof The proof strategy is to proceed by induction and for each m, to construct an equilibrium in which strategies in any strict subgames form CPEs and which gives an expected value $Z^g_m(\lambda)$, such that $Z^g_m(\lambda) - Y^g_{m-1}(\lambda) > 0$. Then by definition $Y^g_m(\lambda) \geq Z^g_m(\lambda)$ and so the lemma would be proven. It remains to prove the existence of such an equilibrium with such a payoff.

n=2 Recall Section 4.2 and the two period version of Proposition 3, the equilibrium strategies in this proposition and the induced value which we here denote $Z^g_2(\lambda)$ satisfy the hypothesised properties.

Inductive step

Recall $\underline{\lambda}_m$ as defined in equation (11), we consider two cases separately and $\widehat{\lambda}_m$ as defined in the paragraph below equation (16). First suppose that $\underline{\lambda}_m \geq \widehat{\lambda}_{m-1}$, and then that $\underline{\lambda}_m < \widehat{\lambda}_{m-1}$.

Case I: $\underline{\lambda}_m \geq \widehat{\lambda}_{m-1}$ First if $\underline{\lambda}_m \geq \widehat{\lambda}_{m-1}$ then it is possible to construct an equilibrium which entails CPE in all strict subgames and yields $Z^g_m(\lambda) = \lambda g + (1-\lambda)b - c + \beta g Y^g_{m-1}(\lambda^s) + \beta(1-$

[31] It is easy to verify that $\widehat{\lambda}_m \leqslant 1$ since it is easy to verify that it can not be the case that $\underline{Y}^g_m > \frac{(g-c)(1-\beta^{m+1})}{1-\beta}$.

$g)Y^g_{m-1}(\lambda^f)$ for $\lambda \geq \underline{\lambda}_m$ and $Z^g_m(\lambda) = \underline{\lambda}_m g + (1-\underline{\lambda}_m)b - c + \beta g Y^g_{m-1}(\underline{\lambda}^s_m) + \beta(1-g)Y^g_{m-1}(\underline{\lambda}^f_m)$ otherwise.

Then for $\lambda \geq \underline{\lambda}_m$,

$$\begin{aligned} Z^g_m(\lambda) - Y^g_{m-1}(\lambda) \quad &= \lambda g + (1-\lambda)b - c + \beta g Y^g_{m-1}(\lambda^s) + \beta(1-g)Y^g_{m-1}(\lambda^f) \\ &\quad -[\lambda g + (1-\lambda)b - c + \beta g Y^g_{m-2}(\lambda^s) + \beta(1-g)Y^g_{m-2}(\lambda^f)] \\ &= \beta g[Y^g_{m-1}(\lambda^s) - Y^g_{m-2}(\lambda^s)] + \beta(1-g)[Y^g_{m-1}(\lambda^f) - Y^g_{m-2}(\lambda^f)] \\ &> 0 \end{aligned} \tag{17}$$

where the inequality follows from the hypothesised inductive property.

For $\underline{\lambda}_m > \lambda \geq \widehat{\lambda}_{m-1}$

$$\begin{aligned} Z^g_m(\lambda) - Y^g_{m-1}(\lambda) \quad &= \underline{\lambda}_m g + (1-\underline{\lambda}_m)b - c + \beta g Y^g_{m-1}(\underline{\lambda}^s_m) + \beta(1-g)Y^g_{m-1}(\underline{\lambda}^f_m) \\ &\quad -[\lambda g + (1-\lambda)b - c + \beta g Y^g_{m-2}(\lambda^s) + \beta(1-g)Y^g_{m-2}(\lambda^f)] \\ &> \lambda g + (1-\lambda)b - c + \beta g Y^g_{m-1}(\lambda^s) + \beta(1-g)Y^g_{m-1}(\lambda^f) \\ &\quad -[\lambda g + (1-\lambda)b - c + \beta g Y^g_{m-2}(\lambda^s) + \beta(1-g)Y^g_{m-2}(\lambda^f)] \\ &> 0 \end{aligned} \tag{18}$$

where the first inequality follows since $Y^g_{m-1}(\lambda)$ is increasing in λ, and the second inequality follows in analogous way to the inequality in (17).

Finally for $\widehat{\lambda}_{m-1} > \lambda$

$$\begin{aligned} Z^g_m(\lambda) - Y^g_{m-1}(\lambda) \quad &= \underline{\lambda}_m g + (1-\underline{\lambda}_m)b - c + \beta g Y^g_{m-1}(\underline{\lambda}^s_m) + \beta(1-g)Y^g_{m-1}(\underline{\lambda}^f_m) \\ &\quad -[\widehat{\lambda}_{m-1} g + (1-\widehat{\lambda}_{m-1})b - c + \beta g Y^g_{m-2}(\widehat{\lambda}^s_{m-1}) + \beta(1-g)Y^g_{m-2}(\widehat{\lambda}^f_{m-1})] \\ &> \widehat{\lambda}_{m-1} g + (1-\widehat{\lambda}_{m-1})b - c + \beta g Y^g_{m-1}(\widehat{\lambda}^s_{m-1}) + \beta(1-g)Y^g_{m-1}(\widehat{\lambda}^f_{m-1}) \\ &\quad -[\widehat{\lambda}_{m-1} g + (1-\widehat{\lambda}_{m-1})b - c + \beta g Y^g_{m-2}(\widehat{\lambda}^s_{m-1}) + \beta(1-g)Y^g_{m-2}(\widehat{\lambda}^f_{m-1})] \\ &> 0 \end{aligned} \tag{19}$$

Case II: $\underline{\lambda}_m < \widehat{\lambda}_{m-1}$ Now it is possible to construct an equilibrium which entails CPE in all strict subgames and yields $Z^g_m(\lambda) = \lambda g + (1-\lambda)b - c + \beta g Y^g_{m-1}(\lambda^s) + \beta(1-g)Y^g_{m-1}(\lambda^f)$ for $\lambda \geq \widehat{\lambda}_{m-1}$ and $Z^g_m(\lambda) = \widehat{\lambda}_{m-1} g + (1-\widehat{\lambda}_{m-1})b - c + \beta g Y^g_{m-1}(\widehat{\lambda}^s_{m-1}) + \beta(1-g)Y^g_{m-1}(\widehat{\lambda}^f_{m-1})$ otherwise.

Then for $\lambda \geq \widehat{\lambda}_{m-1}$

$$\begin{aligned} Z^g_m(\lambda) - Y^g_{m-1}(\lambda) \quad &= \quad \lambda g + (1-\lambda)b - c + \beta g Y^g_{m-1}(\lambda^s) + \beta(1-g)Y^g_{m-1}(\lambda^f) \\ &\quad\quad -[\lambda g + (1-\lambda)b - c + \beta g Y^g_{m-2}(\lambda^s) + \beta(1-g)Y^g_{m-2}(\lambda^f)] \\ &= \quad \beta g[Y^g_{m-1}(\lambda^s) - Y^g_{m-2}(\lambda^s)] + \beta(1-g)[Y^g_{m-1}(\lambda^f) - Y^g_{m-2}(\lambda^f)] \\ &> \quad 0 \end{aligned} \tag{20}$$

and for $\lambda < \widehat{\lambda}_{m-1}$

$$
\begin{aligned}
Z^g_m(\lambda) - Y^g_{m-1}(\lambda) \qquad & = 1 - \widehat{\lambda}_{m-1} g + (1 - \widehat{\lambda}_{m-1}) b - c \\
& +\beta g Y^g_{m-1}(\widehat{\lambda}^s_{m-1}) + \beta(1-g) Y^g_{m-1}(\widehat{\lambda}^f_{m-1}) \\
& -[\widehat{\lambda}_{m-1} g + (1 - \widehat{\lambda}_{m-1}) b - c \\
& \beta g Y^g_{m-2}(\widehat{\lambda}^s_{m-1}) + \beta(1-g) Y^g_{m-2}(\widehat{\lambda}^f_{m-1})] \\
& > 0
\end{aligned} \tag{21}
$$

This covers all possibilities and concludes the proof. ■

References

[1] AGHION, P., P. BOLTON, C. HARRIS and B. JULLIEN (1991): "Optimal Learning by Experimentation," *Review of Economic Studies*, **58**, 621-654.

[2] BENOIT, J.-P. and V. KRISHNA (1993): "Renegotiation in Finitely Repeated Games," *Econometrica*, **61**, 303-323.

[3] BERGEMANN, D. and J. VALIMAKI (1996): "Learning and Strategic Pricing," *Econometrica*, **64**, 1124-1149.

[4] BERNHEIM, B. D., and D. RAY (1989): "Collective Dynamic Consistency in Repeated Games," *Games and Economic Behavior*, **1**, 295-326.

[5] BOLTON, P. and C. HARRIS (1999): "Strategic Experimentation,"*Econometrica*, **69**, 349-374.

[6] CHO, I.-K. and D. M. KREPS (1987): "Signalling Games and Stable Equilibria," *Quarterly Journal of Economics,* **102**, 179-222.

[7] FARRELL, J. and E. MASKIN (1989): "Renegotiation in Repeated Games," *Games and Economic Behavior*, **1**, 326-360.

[8] FELLI, L. and C. HARRIS (1996): "Learning, Wage Dynamics, and Firm-Specific Human Capital," *Journal of Political Economy,* **104**, 838-868.

[9] FUDENBERG, D. and J. TIROLE (1991): *Game Theory*, MIT Press.

[10] HOLMSTROM, B. and R. B. MYERSON (1983): "Efficient and Durable Decision Rules with Incomplete Information," *Econometrica*, **51**, 1799-1820.

[11] JUDD, K. L. and M. H. RIORDAN (1994): "Price and Quality in a New Product Monopoly," *Review of Economic Studies,* **61**, 773-789.

[12] KELLER, G. and S. RADY (1999a): "Optimal Experimentation in a Changing Environment,"*Review of Economic Studies,* **66**, 475-507.

[13] KELLER, G. and S. RADY (1999b): "Market Experimentation in a Dynamic Differentiated-Goods Duopoly," *Research Paper No 1529*, Graduate School of Business, Stanford University.

[14] KHILSTROM, R. E., L.J. MIRMAN and A. POSTLEWAITE (1984): "Experimental consumption and the 'Rothschild Effect"' in eds. BOYER, M. and KHILSTROM, R., *Bayesian Models in Economic Theory*

[15] MILGROM, P. and J. ROBERTS (1986): "Price and Advertising Signals of Product Quality," *Journal of Political Economy*, **94**, 796-821.

[16] OTTAVIANI, M. (1999): "Monopoly Pricing with Social Learning, " *mimeo*, University College London

[17] ROTHSCHILD, M. (1974): “A Two-Armed Bandit Theory of Market Pricing,” *Journal of Economic Theory*, **9**, 185-202.

[18] STOKEY, N. L. and R. E. LUCAS with E. C. PRESCOTT (1989): *Recursive Methods in Economic Dynamics*, Harvard University Press.

[19] TADELIS, S. (1999): “What’s in a Name? Reputation as a Tradeable Asset,” *American Economic Review*, **89**, 548-563.

Figure 1

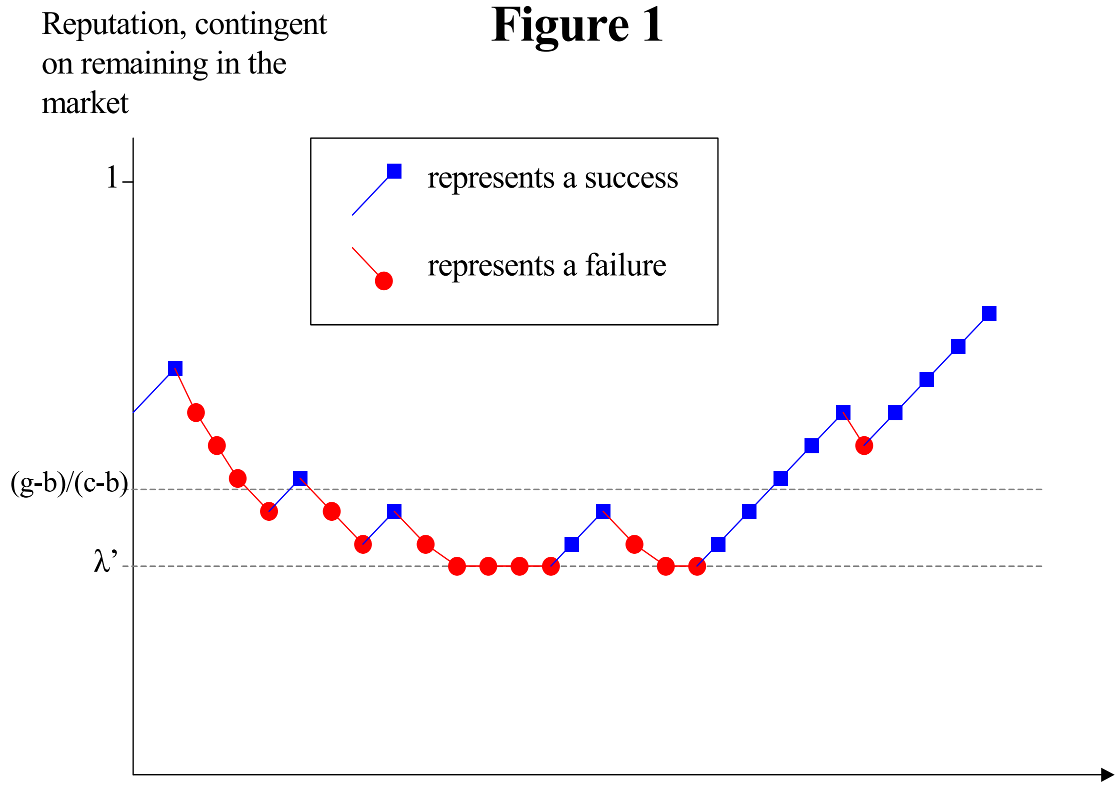

Figure 2