REPUTATION AND PROFESSIONAL SERVICES: SURVIVAL, TEAMS
AND INCENTIVES

A Dissertation Submitted In Partial Fulfilment Of The Requirements For
The Degree Of

DOCTOR OF PHILOSOPHY
IN ECONOMICS

by

Heski Bar-Isaac

LONDON SCHOOL OF ECONOMIC AND POLITICAL SCIENCES,
UNIVERSITY OF LONDON

March 2004

UMI Number: U199015

UMI

Dissertation Publishing

ProQuest

ABSTRACT

Reputation and Professional Services: Survival, Teams and Incentives

Heski Bar-Isaac

Following the opening chapter, which surveys existing literature on the issue of the interaction between individual and group reputation, the remaining chapters each address a simple question to better understand how reputation affects outcomes and incentives. Specifically:

- Even though good and bad luck might affect short term reputations, do agents end up with the reputations that they deserve, that is one that reflects their genuine underlying ability? The central result of Chapter 2 is that, if the agent knows her own ability (though customers can only make inferences by observing history) then eventually the truth will out.

- How can an agent, who has proven her ability, commit to working hard? Once an agent has established a reputation then it is a tautology to say that she has no reputational concern. However, if effort is efficient she may want to commit to exerting effort and get adequately rewarded. Chapter 3 shows that one way that she may be able to do this is by hiring and working with a junior agent of uncertain ability.

- Do teams care any more or less about their reputations than individuals? Joint work with Juanjo Ganuza and introduced in Chapter 4 suggests that an important aspect in answering this question would be to determine whether an agent or team's reputational concerns are primarily about a concern to show itself to be excellent and capable in every task, or about a concern to avoid customers thinking that it is inept.

- How does industry structure affect a firm's ability to commit to producing high quality? The final chapter discusses this question and suggests that a number of different effects are in play so that there is no simple answer.

3

Acknowledgements

that a Ph.D. and academic research is a lonely calling seemed somewhat ridiculous and I thank the numerous colleagues and friends who helped make it seem so.

I gratefully acknowledge the hospitality and support of the Department of Economics and the Financial Markets Group at the London School of Economics and Political Sciences, the Eitan Berglas School of Economics at Tel Aviv University, the Department of Economics at Harvard University, the Center for Mathematical Studies in Economics and Managerial Science at Northwestern University and the Economics Department at the Stern School of Business of New York University.

Preface

And this is the joke.

There are three men on a train. One of them is an economist and one of them is a logician and one of them is a mathematician. And they have just crossed the border into Scotland (I don't know why they are going to Scotland) and they see a brown cow standing in a field from the window of the train (and the cow is standing parallel to the train).

And the economist says, "Look, the cows in Scotland are brown."

And the logician says, "No. There are cows in Scotland of which one at least is brown."

And the mathematician says, "No. There is at least one cow in Scotland, of which one side appears to be brown."

And it is funny because economists are not real scientists and because logicians think more clearly but mathematicians are best.[1]

As a former student of mathematics and a continuing student of economics, these comments by Christopher Boone, the autistic teenage narrator of Mark Haddon's novel *the curious incident of the dog in the night-time*, and the reactions I envisage of various people, both within academic economics and outside, raise somewhat mixed emotions. The approach to which I have aspired in the pages below is to say "There is at least one cow in Scotland, of which one side appears to be brown; let's proceed, albeit skeptically, on the basis that there are some brown cows in Scotland."

---

[1] From *the curious incident of the dog in the night-time* a novel by mark haddon (p.143) 2003, doubleday

Contents

List of Figures

Introduction

Underlying this piece of work is the experience that I have had and observed of friends and colleagues in professional services firms, such as consulting firms, law firms, architect practices and so on. To me at least, a central feature of these industries is that it is very difficult to write outcome-contingent contracts and so customers, in determining their willingness to pay for a service or choice of where to buy this service, rely on any information that they might have available (such as industry directories, trade journals, personal contacts—all of which essentially derive primarily from someone's previous experience) and on the agent's concern for its reputation. I focus on both these aspects—how information is generated and agents' reputational concerns—in the pages below.

In particular, following the opening chapter, which surveys existing literature in a somewhat idiosyncratic way to focus in particular on the issue of the interaction between individual and group reputation, the remaining chapters each address a simple question to better understand the wider overall theme outlined in the paragraph above. Specifically:

- Even though good and bad luck might affect short term reputations, do agents end up with the reputations that they deserve, that is one that reflects their genuine underlying ability? Chapter 2 draws on existing literature to show first that this need not be the case if the agent can take an action (such as leaving the industry) which prevents further learning. The central result of this chapter, however, is that, if the agent knows her own ability (though customers can only make inferences by observing history) then eventually the truth will out.
- How can an agent, who has proven her ability, commit to working hard? Once an agent has established a reputation, so that her actions will not affect others' beliefs about her then it is a tautology to say that she has no reputational concern. However, if her hard work is more valued by customers than the displeasure it causes her, she may want to commit to hard and get adequately

rewarded. Chapter 3, shows that one way that she may be able to do this is by hiring and working with a junior agent of uncertain ability.

- Do teams care any more or less about their reputations than individuals? Joint work with Juanjo Ganuza and introduced in Chapter 4, suggests that an important aspect in answering this question would be to determine whether an agent or team's reputational concerns are primarily about a concern to show itself to be excellent and capable in every task, or about a concern to avoid customers thinking that it is inept.

- How does industry structure affect a firm's ability to commit to producing high quality? The final chapter discusses this question and suggests that a number of different effects are in play so that there is no simple answer.

Enlarging on each of these bullet points in turn, the first chapter briefly reviews alternative modelling approaches to understanding reputation and show how these can be used and have been applied to consider reputations of individuals working in a firm. The focus is on models in which the firm reputation is informative about the type of an individual belonging to the firm. The selection of papers is somewhat idiosyncratic and inevitably incomplete, but illustrates a number of important themes which appear throughout the rest of the thesis.

In the next chapter, I ask whether firms end up with the reputations they deserve in the long run or whether short-term effects are critical. Specifically, I 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 and similar results are generated when the seller's opportunities for strategic behaviour are limited.

Individuals with established reputations by definition have no reputational incentives to exert effort when working alone. However, working as part of a team can introduce such incentives and so allow agents to commit to exerting effort. This is the theme of Chapter 3. In particular, I show that this occurs when a senior acts as a residual claimant of at least some of the future reputation of her co-worker and her choice of action can influence the evolution of her co-worker's reputation. The junior's incentive problem is addressed by an up-or-out contract and the possibility of becoming a senior

in an overlapping generations model. Thus when young, agents are concerned for their own reputations and when old for the reputation of their employees. The robustness of this mechanism in a number of related models are discussed.

Chapter 4, which features joint work with Juanjo Ganuza, begins with the observation that one way to improve the static productivity or efficiency of an agent or organization is to reduce the range of tasks in which it is incapable, or to increase the range of tasks that it can perform effortlessly. The reputational consequences of such improved efficiency, however, depend critically on the kind of reputation considered. In particular, when reputation is about avoiding a reputation for ineptitude, such improved efficiency heightens reputational incentives. However, when reputation is about trying to show excellence (the kind of reputation typically considered in the literature) then such efficiency reduces reputational concerns. We illustrate this formally, relate the observation to previous literature on reputation and briefly discuss applications to training and optimal size of teams.

In the final chapter I go on to consider the interaction between reputational incentives and industry structure. When the degree of competition that a firm faces changes, the profits when committed to good behaviour (such as efficient high quality) and bad behaviour as well as the short-term profits from "cheating" (behaving badly when good behaviour is expected) all change. Moreover these profits change at different rates, so that overall as the degree of competition changes, the effect on the ability of a firm to commit to good behaviour is ambiguous. In particular, the effect of competition on the likelihood that a firm produces high quality is not monotonic and so, for example, greater competition may lead to higher prices. I prove and illustrate this argument (considering two different measures of competition which show qualitatively similar effects) and discuss implications (for example a firm might choose more competitive environments) and related literature.

The appendices contain the majority of the proofs and a number of variations of the models contained in the main text and intended to demonstrate some robustness of the principal qualitative results.

CHAPTER 1

On reputation and the interaction of firm and individual reputations

1.1. Introduction

In this chapter, I seek to draw together and categorize differing models pertaining to firm reputations and in particular I highlight those models that distinguish and discuss the difference between reputations of firms and of individuals. This survey is not exhaustive; however, I hope that it might prove useful.

The motivation for my interest in these models was a broader interest in professional services, such as law and consulting. In such markets which are essentially markets for credence goods (a term introduced by Darby and Karni (1973)), reputation clearly plays a crucial role. However, spin-offs in consulting (see for example Pinault (2001)), the importance of client control (Nelson (1988)) and increasing occurrence of lateral hiring in the legal profession (Trotter (1997)) highlight that individuals can develop reputations which differ from the reputations of the firms to which they belong. Moreover, in such industries, human capital is the critical factor of production. Much of the discussion in this chapter (and no doubt its coverage) may be strongly influenced by consideration of these industries.

Three different theoretical approaches for modelling reputation, which highlight different aspects of reputation are introduced and discussed. These approaches underlie much of the subsequent discussion which focuses on issues surrounding the interaction of individual and collective reputations. After introducing this issue and the relevant actors involved, a number of models are surveyed, and the similarities and points of distinction are highlighted. A final section concludes.

1.2. Three approaches to reputation

In this section, before considering specific models of reputations of firms distinct from individuals, I briefly review different notions of reputation that have been employed in

formal models. Typically these have been applied to a firm as an unitary entity or to an individual and have been developed in a number of different applications (including predatory pricing, central banking, reputation for quality etc.)

In particular, more recent literature, following Fudenberg and Tirole (1991) restricts the term reputation for consideration of the first notion considered below, that is where an agent's reputation is defined as the belief of others about the type of the agent, which is drawn from some distribution. I briefly review this notion and then go on to consider two other notions of reputation. The first of these two views reputation as a pure commitment device and is strongly related to the type-based notion which to some extent predates it. Both this approach and the belief-based approach highlight reputation as a valuable asset in which one can invest and which can be dissipated (often in an extreme and discontinuous way).

Finally, I consider reputation as a coordinating device; later in the discussion on firm reputation, I relate this notion both to the coordination within the firm and among those externally involved in the firm. Note that in a broad sense in the first two notions as well, reputation acts to coordinate actions of different agents (for example in the repeated game notion, an agent with no reputation to maintain does not take the good action as her lack of reputational capital affects the behaviour of other agents with whom she interacts; in the type based notion this kind of coordinated action between the different agents affects only strategic agents rather than commitment types). More generally, any equilibrium chosen when there are multiple equilibria relies implicitly on some kind of coordination—reputation is a natural and explicit device to allow this.

1.2.1. Reputation as belief

This view of reputation begins by supposing that there is incomplete information about a player's type with different types expected to play in different ways. In many models (including Kreps and Wilson (1982), Milgrom and Roberts (1982), Fudenberg and Levine (1992) and Ely, Fudenberg and Levine (2002)) it is assumed that at least one type will behave in some deterministic fashion (these are termed "commitment" types in Fudenberg and Levine (1989)) and in these models only one type behaves strategically. In all such "reputation as belief" models, a player's reputation is summarized by other players' (and possibly her own) current beliefs about her type. These beliefs in turn (and depending on the equilibrium strategies) lead other players to expect particular actions

(or distributions over a set of actions), in equilibrium these expectations are rational or accurate predictions of behaviour.

Note that in these models, for reputation to play a role in affecting behaviour, it must be that there is some uncertainty, or equivalently that the reputation is at risk. If the player's type is known precisely then her action will have no effect on other's beliefs about her and so will not affect her behaviour. In the Kreps, Wilson, Milgrom and Roberts models such uncertainty can be maintained indefinitely as when the strategic type behaves identically to a strategic type, there is no further learning on the part of other players and so uncertainty is perpetually maintained. In the absence of a commitment type that the strategic type would like to mimic, other mechanisms are necessary to maintain the uncertainty required for reputational incentives, as discussed in particular in Chapter 3. Even with such a commitment type, as Cripps et al. (2004) show, reputational incentives cannot be sustained indefinitely if actions are only imperfectly observed.

1.2.2. Reputation as a pure commitment device

The notion of reputation purely as a form of commitment is based on a framework where in effect the type of the agent is known. The approach focuses on self-sustaining equilibrium strategies in infinite horizon models which ensure some kind of good behaviour ("cooperation" in the prisoner's dilemma, for example) with the threat of some kind of punishment should the agent ever defect. This is the usual trigger strategy approach, where so long as everyone has been acting as they should, then everyone cooperates but if ever this trust is breached (or the good reputation is lost) there is punishment. Such punishment need not be permanent – as highlighted in particular in the work on repeated games with imperfect observation of actions of Green and Porter (1984), and Abreu, Pearce and Stachetti (1990). However, a wide range of equilibria is typically admitted and the results rely strongly on an assumption of an infinite horizon in these models. An example of this approach is Klein and Leffler (1981) in which a firm can maintain high quality and high prices with the threat that should it ever deliver low quality products, it would no longer be able to charge high prices.

This approach has been criticized on the grounds that the initial reputation springs from nowhere, though this criticism has been to some extent addressed. In Klein and Leffler (1981), for example, the seller must spend on advertising, marble-clad offices or some other irrecoverable sunk cost in order to establish an initial reputation.

Another criticism of this approach has been that there is little scope for interesting reputational dynamics, at least in the perfect observability case. The common perception of reputation as an asset, which can be slowly built up, dissipated and then rebuilt, is hard to accommodate in this approach to reputation, whereas the type-based approach can allow for this sort of behaviour.[1] Moreover, the type-based approach can also be seen as advantageous inasmuch as it allows for reputation to affect behaviour even in a finite horizon model, it might allow for equilibrium selection or more robust equilibrium, and perhaps most importantly, the type based approach seems to ring true in describing people's understanding of reputation—that it reflects a belief about the type of person that an agent is.

In many cases, however, the belief-based approach to reputation introduces reputation effects in allowing strategic types at given reputation levels to commit to particular actions and in a number of models, the reputation as a pure commitment approach and reputation as belief approach can lead to similar outcomes.[2] For example in the repeated prisoner's dilemma one can allow for the possibility of a type that always cooperates, and this might allow a strategic types also to commit to cooperation in an attempt to maintain a reputation as this cooperating type; a similar outcome can arise in the repeated game with no type heterogeneity where players play the "grim trigger" strategies "cooperate until an opponent defects and then defect thereafter".

1.2.3. Reputation as a coordination device

The third notion of reputation views reputation as a coordination device. As a simple example, consider driving conventions. Most people know that in England, people drive on the left and in the US on the right; either convention would work equally well, but it is important that in each country people drive on the same side, even though this may differ by country. One can think of the US as having a reputation for people driving on the right (albeit one that is maintained to some extent by legal requirements) and that this reputation acts to coordinate people's actions. Thus reputation can establish a "focal point" for coordinated action (Schelling (1960)). Boot and Milborn (2002) for example, consider the role of credit ratings in coordinating behaviour and expectations

---

[1]See for example Diamond (1989) and Benabou and Laroque (1992) who can allow for such reputational dynamics in a type-based model of reputation.

[2]See, in particular, Fudenberg and Levine (1989).

and so in determining a particular equilibrium in an environment where without the use of these ratings, many equilibria are possible.

This kind of reputation underlies some discussions of corporate culture, in particular in parts of Kreps (1990).[3] Kreps discusses corporate culture as some sort of principle or rule that has wide applicability and is simple enough to be interpreted by all concerned together with the means by which this principle is communicated. In many instances, a number of principles may be valid, but that one is chosen and clearly communicated, understood and anticipated is crucial.[4] Defining firm boundaries and activity conducted with and within a firm which has a particular corporate culture (or reputation for up-holding a particular principle) can therefore be of considerable value.

A related literature, which is concerned with the role of reputation in coordinating the behaviour of a firm's customers rather than focusing on behaviour within the firm, considers the coordinating role of advertising. In these models one could consider advertising as establishing a reputation that the product will be popular and thus these models address this aspect of reputation. In particular, Bagwell and Ramey (1994) suppose that the price that a seller charges decreases with the number of customers that the seller expects and Pastine and Pastine (1999) suppose a consumer's valuation for a product is directly increasing in proportion to the number of others that purchase; in both these models coordination takes place through observation of advertising levels. Finally, in Clark and Horstmann (2001) coordination takes place through consumer beliefs about advertising levels.

Advertising campaigns have frequently appealed to this notion of reputation; for example IBM ran a campaign with a famous slogan "no one ever got fired for buying IBM" or the British soap opera *Eastenders* ran a campaign with the slogan "Eastenders: Everyone is talking about it". The IBM campaign, rather than appealing to an intrinsic quality of the computers appeals to a social quality—that others rate IBM computers as good; more directly in the Eastenders campaign it is clear that the appeal is to a social externality.[5]

---

[3]See Hermalin (2001) for a survey on corporate culture.

[4]For example consider a game of matching pennies—coordination on an equilibrium is more important than the particular equilibrium chosen.

[5]See Brandenburger and Polak (1996) for a related idea in which CEOs behave even against their information to satisfy expectations of investors.

1.3. Individual and firm reputations

In this section I ask when and why it may be important to distinguish between individual and firm reputations and highlight that in considering reputation, it is important to step back and ask what is the reputation for and who are the parties interested in this reputation.

In essence, distinguishing between the reputation of a firm and an individual member of that firm, is an exercise in opening the "black box" of the firm and in developing an understanding of firm and individual career dynamics. If the constituent members of a firm never change and have no opportunity of leaving and continue operating in an identical way throughout time, then there is little value in distinguishing between the reputation of individuals within the firm and the contribution of the firm itself (I discuss below what this might mean).

However, since employees may choose to leave and work in other firms or on their own or can threaten to do so as a means of raising their wage, or other new employees may join the firm, developing individual reputations may be of considerable interest and disentangling the reputation of the firm might be valuable. Reputational considerations, in addition to determining wages, can also be useful in terms of sorting—that is ensuring that the appropriate agent is assigned to the most productive job, or rather the job in which they are most valuable.[6] In addition to determining wages and affecting job assignments, and through its effect on these, reputational concern can also affect incentives, task assignment and the internal design of organizations.

1.3.1. Reputation for what?

Firms interact with numerous different agents and in many different ways. Audiences include customers, current employees, potential employees, competitors (in the product market and in the labour market), and investors and large firms may also interact with consumer groups and government. For each of these groups, the firm's reputation may have a different meaning; for example, customers are likely to be primarily interested in the quality of the service provided, and current and future employee may be interested in the firm's informal compensation bonus and promotion schemes, in the training provided

---

[6]See Felli and Harris (1996) in which the value of a worker in a job is determined by the best estimate of their current value in that job and how much information would be gained about their productivity in this job and in an alternative. See also Anderson and Smith (2002) where learning considerations can over-ride considerations of productivity in a decentralized market.

and in the firm's reputation as a good trainer or for its ability to identify good workers (which would affect the employee's future career prospects). To some extent, it may be possible to write explicit contracts to deal with some of these features (for example for some experience goods the firm might be able to provide a warranty) but in many cases no formal contract can be written and reputation is critical.

In highlighting that there are many different attributes of the firm for which it may develop a reputation and that different interested parties have different concerns, we draw attention to the observation that much informal discussion on developing a good reputation is somewhat disingenuous inasmuch as it does not discuss for what and with whom. Moreover, it is far from clear that developing a good reputation with one party entails a good reputation among other interested parties. For example, having a reputation as a generous employer might help to attract and retain good staff, but may lead customers to believe that the goods or services are over-priced. Similarly, a firm might benefit from customers believing that the firm's employees are excellent individuals but if competitors also share these beliefs then the firm might be at risk that its staff are poached.[7] Nevertheless, Fomburn (1996, p.395/6) has argued that it makes sense to speak of a corporate reputation as a single notion, though the firm may have numerous attributes. Fomburn argues that one should consider a company's overall reputational profile though this should be built up by identifying a company's key constituent groups, sampling these and aggregating to obtain a reputational rating on "relevant dimensions".

Chapter 4, for example, highlights that even thinking about reputational concerns in terms of customers' beliefs regarding the quality of service that they will receive can have exactly opposite comparative statics depending on whether the agent's concern is to acquire a reputation for excellence or to avoid one for ineptitude.

In another recent paper, Koszegi and Li (2002) show that when agents differ not only with regard to their ability but also with respect to their ambition, then since a good outcome may reflect high ambition rather than high talent, a good outcome may reflect ambition and so be interpreted as bad news about the agent's ability. This richer view of an agent's motivation has interesting consequences, for example a principal might want to observe measures of the agent's effort (for example hours worked) early but not late in the agent's career.

---

[7]This concern, over giving too much information to rival employers has been explored; for example in Greenwald (1986) and Waldman (1984).

Below and following a large literature in economics, I focus on reputation as customers' belief that they will receive high quality service. This will involve primarily the ability of individual employees of the firm, but also the firm's ability to ensure effort on the part of its workers, its ability to introduce good new employees, and its ability to coordinate the behaviour of its employees and possibly of other actors. The last of these categories—that is the role of reputation in coordinating the behaviour of employees—differs significantly from the others considered and so beyond the discussion of reputation as coordination above we do not return to address it specifically. The other categories apply when each individual's production is independent of other agent's actions, at least in the static sense that a given individual action leads to a given output, independently of whatever anyone else might be doing or whoever else might be in the firm. For the coordinating role of reputation, this is clearly not the case.[8]

Thus my focus is on models of firm reputation which view firm reputation as giving some information about the individuals within the firm, this information may be about:

- productive behaviour of individuals within the firm;
- productive ability of new employees; and,
- productive ability and behaviour of individuals within the firm.

To some extent, these distinctions are arbitrary. For example, if firm reputation is informative about the productive ability of new agents, then in a dynamic model it would also give information on the productive ability of individuals within the firm. However, it is hoped that these distinctions might be useful in highlighting the multi-faceted role of reputation and draw attention to particular aspects (such as the implications for the hiring process).

Each of these categories is addressed separately below; though the last two are at the heart of this chapter.

---

[8]Note that in addition to uncertainty about the ability of employees, there may be uncertainty about a firm's assets, that is a firm's reputation may also depend on non-human capital. To some extent, we allow for this in discussing corporate culture which is an intangible non-human asset, but abstract from allowing for other physical or intangible assets, or equivalently we suppose that these are fixed exogenously and there is no uncertainty concerning them.

1.4. Firm reputation and behaviour

This class of models builds on the notion of reputation described in Section 1.2.2. Before, considering the role of individuals within a firm, first we briefly review this notion applied to firms behaviour with no separate identification of individuals.

In a seminal paper, Klein and Leffler (1981) highlighted the role that a repeat-purchase mechanism can play in ensuring high-quality supply. A necessary condition for such a mechanism to play a role is that information on a firm's current behaviour is at least partially observed by consumers before they make their decisions to buy in the next period; in effect, this condition is equivalent to the requirement that behaviour should be able to change reputation otherwise it is intuitive that reputation would impose little discipline on behaviour. An important contribution of Klein and Leffler (1981) however, is that in itself this is observation of previous behaviour is not sufficient to ensure that a firm will produce high rather than low quality (or in some other way will behave well even though it is more costly in the short run than behaving badly). Klein and Leffler (1981) argue that:

> Cheating will be prevented and high quality products will be supplied
> only if firms are earning a continual stream of rental income that will be
> lost if low quality output is deceptively produced

Equivalently, high reputation firms must earn positive profits, since sellers can always increase profits in the short-run by reducing the quality of their products. Thus it must be the case that behaving well, for example by producing high quality, should earn the firm profits since otherwise it would be dominated by a fly-by-night or hit-and-run strategy of quality reduction.[9] In this context two related questions that arise are how this flow of profits can be reconciled with free entry and where the initial high reputation comes from. Klein and Leffler (1981) suggest that a firm will only obtain such a reputation to begin with following an appropriate level of non-recoverable expenditure, such as advertising—this up-front investment which buys reputation and the revenue flow that it implies therefore works as a bond which is lost if the firm should "cheat". Shapiro (1983), in a model where consumers believe that the firm will produce today at the

---

[9]Essentially since a firm can always withdraw from the industry and earn zero profits, this sets a lower bound for any effective punishment that customers might impose on a firm; such a punishment can only be imposed after any "cheating" has been detected and so a "cheating" strategy always earns positive profits if the firm is expected to behave honestly in the current period. Thus to encourage honest behaviour, such behaviour must be rewarded with positive returns.

quality it produced in the last period, describes how initially customers would believe a firm produces low quality, so that producing high quality (and selling at the low quality price) for an initial period can dissipate profits and act as an investment in reputation.

Thus in both models, early investment and continual good behaviour lead to a stream of returns on reputational capital, though "cheating" might be more attractive in the short-term. In Klein and Leffler (1981) the early investment in reputation is through some irrecoverable expenditure, in Shapiro (1983) this kind of irrecoverable expenditure is more directly observed and indeed the focus of the paper—it is a loss incurred on early sales in an early phase of building up a reputation. These models require that the firm be infinitely-lived—a finitely-lived firm would "cheat" in the final period, and so there would be no sense in maintaining reputation by behaving well in the penultimate period, and similarly such incentives unwind for earlier periods.

The acquisition of a firm's name can constitute a kind of irrecoverable investment as suggested by Kreps (1990).[10] In particular this implies the construct of a firm name can effectively act as a potentially infinitely-lived object and so even though individual agents might be finitely-lived they have the incentive to maintain the firm name, which they can sell to younger agents. Younger agents in turn must purchase such costly firm names as a bond to ensure that will behave well. Thus this construct both creates a potentially infinitely-lived object (the firm-name) for which the repeat-purchase mechanism can ensure a reputational concern for good behaviour and through its costly transfer, it allows finitely-lived agents to maintain reputations (it is valuable for an agent to do so since at retirement she can sell the name that she has maintained).

Moreover, in contrast to the Klein and Leffler (1981) and Shapiro (1983) models which do not distinguish between a firm and an individual, Kreps (1990) does make this distinction, in a model where the firm can be thought of as a label passed from one individual to another; and individuals need such a label in order to credibly promise good behaviour. In a related model, Cremer (1986) also considers over-lapping generations in a model with no direct monetary transfers between members of a firm but where a junior in a firm behaves well as otherwise the firm's reputation would be lost and a new junior would cheat when she is a senior.

---

[10]The investment in the firm name is recoverable, but only in the case that the holder of the name has behaved well, that is it acts as a bond.

1.5. Firm reputation and hiring new employees

In this section, I consider the interaction of firm reputation and the recruitment of new employees. First, as outlined above, inasmuch as a firm affects the behaviour of those within it, this clearly extends to the behaviour of new employees; however, in this section, we focus on reputational concerns and hiring decisions. Essentially, in this section, I ask how reputational concerns affect who is hired.[11]

First note, that in asking the question of who gets hired, we already impose an assumption that there are different kinds of agents who might be hired—that is I implicitly type heterogeneity among the pool of potential employees. Moreover in most of the models discussed below, it is assumed that either employers have better information about potential employees than customers or otherwise that they can acquire it.[12] I review models of the following form:

- firm reputation as commitment to select only good employees (certainty about firm, uncertainty about employee);
- hiring as a costly signal (uncertainty about firm, no uncertainty about employee);
- uncertainty about both firm and employee—umbrella branding, scapegoat and other notions; and,
- hiring to introduce uncertainty.

1.5.1. Firm reputation as commitment to select only good employees (certainty about firm, uncertainty about employee)

As discussed above, in Section 1.3.1, there are different aspects for which a firm may develop a reputation; in particular, even if customers know perfectly the ability of those currently employed by the firm, a firm may seek to develop a reputation for only hiring good employees. Although as far as I know this idea has not been much explored in the

---

[11]In particular I ignore complementarities in production itself.

[12]It is possible to construct a model where hiring can still affect the reputation of a firm when there is no type heterogeneity among potential employees or customers have identical information to the firm, in which hiring in effect acts as a costly signal about the firm or existing employees rather than about new employees. For example, hiring many employees might serve to demonstrate that a firm is confident in its ability to continually win new business, this can lead a customer to try it out—I discuss a similar kind of effect below in a model of type heterogeneity.
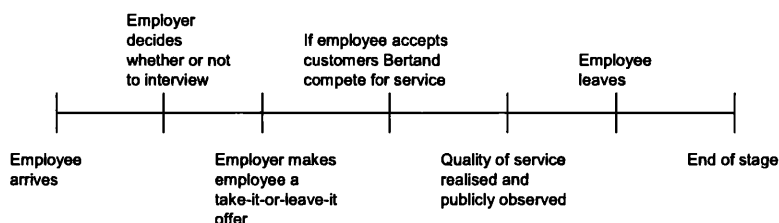
Figure 1.1. Timeline for a stage

context of labour, an analogous idea has been discussed in the context of extending a brand name to new products in Choi (1998).

In order to make this idea clear, consider the following game.

There is a firm or employer, a steady stream of potential employees and a continuum of myopic customers. At each stage, a single employee comes to the employer, the employer may learn the quality of this employee, which may be either low or high, at a cost i (the cost of interviewing).[13] The employer then decides whether or not to make the potential employee an offer, the employee decides whether or not to accept (if not the employee takes up some fixed outside option $R > 0$). If this employee is hired, customers Bertrand compete for the service; only after the service is bought is its quality realized and we suppose that low quality is worth 0 and high quality worth 1. The employee then retires or moves on.[14] Such a stage, as represented in the timeline below, is then repeated.[15] Suppose that *ex-ante*, there is a probability λ that a potential employee will be good and that there is a discount rate between periods of δ.

Standard techniques can be employed to show that so long as $\frac{\lambda(1-R)-i}{1-\delta} \geq (1 - R)$, then the following is an equilibrium:

---

[13]One might ask could customers not directly interview. It may be that the firm has unique access to an interview technology (so for example those with experience are best positioned to judge the potential of employees). Alternatively, one might consider the firm as a delegated monitor, a single employee might work for hundreds of customers, for whom individually it is not worth assessing the quality of this employee.

[14]This assumption is made so that the firm does not grow infinitely and since once the employee's quality is known the profit that the firm may be able to make from this employee may be more limited.

[15]Note that there is some tension in this formulation in supposing that employees stay only one period and yet the firm is infinitely lived; one can think of the employer as holding the firm name, as in the Kreps (1990) model, say. Of course, it also need not be the case that the employee literally only works once, a period may be a length of time that lasts for a while, and it need not be that only one employee is hired at a time.

The firm begins in Regime I and stay there, unless customers observe low quality service in which case switch to Regime II.

Regime I Employer interviews, if the potential employee is good, then employer offers employment at wage R. The service is bought at a price 1.

Regime II Employer does not interview, makes no offer of employment (and if she did then the service would be sold for λ).

Thus in this equilibrium the firm maintains a reputation for hiring only good employees, and should it ever lose this reputation (which it may be tempted into as this would save the cost of interviewing or allow the firm not to have to wait in order to hire) then it loses the option value of future profitable hires. Essentially, then in this model, the firm acts as an intermediary between customers and the labour market, maintaining a reputation for selecting only good employees.[16]

A similar result can be maintained even with no dynamic considerations. In a model of horizontal differentiation, where customers have heterogeneous tastes a firm with better information on workers might have a strict incentive to hire a particular type of worker even within a period. Thus the firm confers a meaningful label on the employees that it admits to the firm. As an example, consider a situation where there are economies of scale in advertising and customers have horizontally-differentiated preferences, say that some customers prefer paintings by left-handed artists and others prefer those by right-handed artists. Then there is a clear benefit for agglomeration, that is for groups to sort themselves separately into organizations of left- and right-handed artists.

1.5.2. Hiring as a costly signal (uncertainty about firm, no uncertainty about employee)

In the discussion above, we have considered a case where there is no uncertainty about the ability of the firm, but there is *ex-ante* uncertainty on the ability of employees. Hiring can also play a role in the opposite case, that is where there is uncertainty about the firm but not about the workers. The costly hiring of a well-known employee can serve as a costly signal as to the type of the firm, or other workers. For example, Hollywood movie-makers clearly take into account in choosing the cast that potential customers

---

[16]For a thorough discussion of different aspects of the firm acting as an intermediary between customers and suppliers see Spulber (1999).

will be influenced by this choice in deciding whether or not to see the movie;[17] similarly academic departments and professional services might display an intention that they are changing strategy with a high profile hire.

This kind of signalling is also known as money-burning.[18] A signal is meaningful inasmuch as it has more value for one type than another, this different value may arise either from different costs or different benefits of the signal (or both). In a money-burning equilibrium, the cost of hiring a star would be identical to either a "good" or "bad" firm, but the benefits differ. Typically this is because customers buy repeatedly and though the signal might cause customers to buy from the firm once (or try the services of a worker within the firm other than the star), they will only buy again if the firm performed well.[19] This in turn can contribute to the wages of those with sufficiently high profiles, and thus contribute to superstars earning stellar wages, which reflect not only a productive ability but also a reputational value.

1.5.3. Uncertainty about both firm and employee

I extend the discussion and now allow for the possibility that the abilities of both the firm and potential employees are uncertain. In this case, a decision to hire (particularly when this is costly, say because the employee needs some basic training in the way that the firm operates) can act as a signal on both. This notion has been explored, though again rather than in relation to hiring, in the context of a firm which must decide whether or not to brand a new product under a new brand name or the existing name. This kind of umbrella branding has been explored, for example by Cabral (2000) and Wernerfelt (1988), who concisely summarizes this mechanism as follows:

> When a firm brands a new product, it is in effect doing two things: it is claiming that the old and new products are both of good quality and it is inviting consumers to pool their experience with the two products to infer the quality of both.

---

[17]Of course, there are a few effects involved, some viewers may be influenced just because they like looking at Arnold Schwarznegger, say, for a couple of hours; others might be influenced because they have expectations of the kind of movies in which he is likely to star.

[18]The notion of costly expenditure as a signal is well-articulated in Milgrom and Roberts (1986) where the expenditure is on advertising rather than hiring a superstar.

[19]In the movie example, knowing the star might lead a movie-goer to look at the reviews, or though it may cause a movie-goer to watch the movie, its quality might then lead her to recommend it or warn others.

This notion thus suggests that good firms are more likely to hire, to keep showing that they are good and are likely to hire good employees. Other authors, however, have suggested that good seniors may hire bad juniors. In particular, Segendorff (2000) and Glazer and Segendorff (2001) suggest that a good senior may choose to work with a junior who is perceived as bad—the rationale is that the incompetent co-worker can credibly be blamed if things go wrong.[20] Such a model relies on the feature that outsiders (whether principal or customer) cannot attribute the contributions of the two agents separately. The model also relies on a non-linearity in rewards to reputation which leads to risk aversion and so an important role for this insurance effect.

1.5.4. Hiring to introduce uncertainty

As discussed in Section 1.2.1, in order for reputational incentives to affect an individual's behaviour, her reputation must be at risk. When an agent has an established reputation therefore, such reputational concerns are muted and so with no other commitment device, an agent would exert no effort when this effort is costly, even though such effort is efficient. In Chapter 3, I argue that hiring and working with a young agent of uncertain ability can introduce this uncertainty and provide reputational concerns to give incentives to the senior (though the reputational concern is for her employee, rather than for herself). This mechanism applies when the established and junior agents can work together in a non-observable production process, there is sufficient uncertainty about the junior's ability and the senior agent is able to gain from the junior's reputation in the future.

1.6. Firm reputation as information on individual ability and behaviour

Although the section above highlighted the role of firm reputation on new employees and indeed on the role of hiring on reputation, in some cases customers may be unaware, or at least partially unaware of an individual worker's history and in particular how long the individual has been associated with the firm. Nevertheless, though customers may have limited information about the individual, they may have better information on the firm or at any rate the information about the firm may confer additional information on the individual. Tadelis (1999) for example distinguishes between the entity and the identity, that is between what constitutes the firm at present and its perception which

---

[20]Segendorff also suggests other motives for competent leaders to choose incompetent workers: to decrease the number of potential rivals and to emphasize their own ability. Glazer (2002) suggests that a leader may not want to work with an able co-worker for fear of theft of assets.

may depend to a greater extent on the past. This distinction, and underlying it the different information that customers have on the firm and its current owner, and the central observation in reputation models that an individual's behaviour is based not only on her type but also on how she is perceived lead to a number of observations which have been highlighted in a number of papers discussed below.

First, an individual's behaviour is influenced by the reputation of the firm to which they belong—since this might change the way that the individual is perceived and so how her behaviour is understood and used to form her own reputation. Second, a firm's reputation is in itself influenced by its past members and their behaviour and so, in particular, the behaviour of current members might depend on the behaviour of past members. Furthermore when the history of a firm's performance is observed better than individual performance and, in particular, when changes in ownership are unobserved, this can lead to a market for names; and, in addition, such changes in ownership can sustain sufficient uncertainty to maintain reputational incentives.

We enlarge on these points below, discussing models according to the following classification:

- the firm consists of a single individual
  - changes to the firm's composition are fully exogenous
  - changes to the firm's composition are partially endogenous
- the firm is "large"
  - changes to firm composition are fully exogenous
  - changes to firm composition are partially endogenous

1.6.1. The firm consists of an individual (exogenous changes in composition)

The simplest model of this sort views the firm as a label or location held by an individual; customers can observe the history of that location (that is whether good or bad products were produced there in the past) but not whether the individual at the location changed. In each period there is a probability that the location holder changes and the type of the new holder is allocated according to some exogenously probability distribution.

In this model, discussed by Phelan (2001) and considered by Holmstrom (1999), the possibility that the type of the location-holder has changed in each period ensures that customers can never become certain about the type of the current holder. Moreover, since there is a probability that the current holder will stay to the next period, her

action can influence her reputation in the next period and this reputational concern will affect her behaviour as she will benefit (or suffer) by behaving well (or badly).

1.6.2. The firm consists of an individual (partially endogenous changes in composition)

A more complicated approach has supposed that the probability with which a location holder is replaced is fixed (or at least constrained as in Assumption (A4) in Tadelis (2002)) but that rather than the new location holder being randomly assigned as above, this assignment results as the outcome of a market for names. Since name transfers (as the changes in the holder of the location) are not observed by customers and cannot always occur (that is there is always some probability that the holder of a name today is the same individual that held it yesterday), good reputations are valuable and can be traded. Moreover, an individual is motivated to build up a reputation not only in case she still holds the reputation tomorrow (as discussed above) but also since even if she does not hold the reputation tomorrow, she may be able to sell the reputation.

Tadelis (1999) develops an adverse selection model (so that there is no behaviour to be influenced) to focus on the market for names and study the economic forces that cause names to be valuable, tradeable assets. The paper shows that an active market for names with either finite or infinite horizons will exist but there can be no equilibrium exists in which only good types buy good names. Incentive effects together with a market for names are considered in both Tadelis (2002), and Mailath and Samuelson (2001), who both highlight that the market for names cannot be fully separating (that is it can never be the case that only good firms buy good names).

1.6.3. Large firm with exogenous composition

Tirole (1996) spells out a theory of individual and collective reputation, where group composition changes exogenously and in a way that is not observed by customers (that is old agents are replaced with new agents and such replacements are not observed). The building blocks of this model, which highlight many of the themes discussed in this chapter, are clearly described and are repeated here:

> (a) *A group's reputation is only as good as that of its members.* Each member is characterized by individual traits such as talent, diligence or

honesty. Past individual behaviour conveys information about these traits and generates individual reputations.

(b) By contrast with group belonging, *individual past behaviour is imperfectly observed.* If past individual behaviour was fully unobserved, members of the group would have no incentive to sustain their own reputations and therefore the group would always be expected to behave badly. Conversely, the collective reputation would play no role if individual behaviours were perfectly observed. Imperfect observability of individual behavior thus underlies the phenomenon of collective reputation.

(c) *The past behaviour of the member's group conditions the group's current behaviour and therefore can be used to predict the member's individual behaviour.* Each member's welfare and incentives are thus affected by the group's reputation.

(d) If we further assume that the age in the group, or the frequency of interactions with the group, or the number of past opportunities for cheating are imperfectly observed, cohorts in a self-generating group are partly pooled and therefore, *the behaviour of new members of a group depends on the past behaviour of their elders.*

In contrast to the models of the firm as an individual, in this model the firm is treated as a large group, so that a single individual's behaviour cannot affect the reputation of the group as a whole. Thus, although an individual is influenced by the collective reputation as it may change the way that her own behaviour is interpreted and her own reputation revised, she does not consider the impact of her behaviour on the collective reputation (or ignores it since it is negligible). The focus, as highlighted by the extensive quotation above is on how partial observability and behaviour leads collective reputation (even when arising from previous generations) to affect an individual's behaviour.

1.6.4. Large firms with endogenous composition

The approach discussed in the paragraph above which considers how the past behaviour of one's peers, by shaping the way that individual achievements are interpreted, could influence present reputational incentives, is further extended in Levin (2001). This chapter considers how decisions about group membership—that is an individual (assumed to know her own type) can decide which group to join. Groups are assumed not be able to observe the types of potential members (or at least not to discriminate) but different

groups can charge a different membership fee (in the labour context this is equivalent to offering a different starting wage, with later wages determined by individual reputations). In particular, Levin (2001) shows that there are equilibria where groups with different characteristics and behaviour coexist, though the model relies on coarse reputations for individuals.

1.7. Conclusions

This chapter outlined three broad modelling approaches to reputation—a type-based view, a view based on commitment in an infinitely repeated game, and an approach based on coordinating action between several agents. Drawing largely on the first such view, the heart of this chapter is a categorization of approaches to the interaction between the reputation of a firm and its employees. A somewhat artificial distinction was made between consideration of a hiring decision and on information about existing employees and a wide range of considerations were introduced.

In the discussion above, there is no productive advantage or loss *per se* in hiring a particular individual, for example (though of course in changing expectations, such a decision may affect beliefs). In general however, firm composition or size will have an important effect on productive efficiency, and this effect may counteract or overwhelm opposition to reputational concerns.[21]

It is perhaps also worth noting that models of reputation and the interaction between individual and firm reputation are to some extent bound to be complicated. These models rely on at least three different kinds of entity (a firm or perhaps employer, an individual or employee who may be associated with the firm and a customer or some other agent who cares about these reputation) and reputation is clearly a dynamic concept. The models described above are necessarily simplifications and abstractions designed to highlight specific mechanisms; however as our intuition for such mechanisms improves, it might be interesting and valuable to consider other aspects, such as the role of risk aversion and precision of beliefs as well as the levels, and understand the interaction between such mechanisms and other factors, such as heterogeneity in prior beliefs[22] or the effects of industrial structure and competition. Further, inasmuch as dynamic problems rely

---

[21]Tension between reputational (or commitment) concerns and static productive efficiency and its implications for firm composition have been highligted, for example, in Levin and Tadelis (2002) and Anderson and Smith (2002).

[22]See for example Van den Steen (2001) on the role of heterogenous priors on corporate culture.

on complex inference problems, the effect of costs in making calculations or other forms of bounded rationality might have a large effect on the conclusions drawn from formal models. Another aspect that has not been addressed in this discussion is how reputation spreads—typically I have assumed that all customers have the same information. In some applications it may be reasonable to suppose that current customers have better information and that the current customer base has a role to play.[23] No doubt numerous other factors can be brought to bear in addressing some of the issues raised above and there is much still to be learned.

---

[23]See for example Ottaviani (1999) and Fishman and Rob (2002) on related questions.

CHAPTER 2

Reputation and survival

2.1. 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 chapter can be read as contrasting the
stylized 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 in the long run, though
Iago's viewpoint might prevail in the short-term.

We consider a model where a monopolist sells single units of a good in discrete time
periods. The quality of the service is unknown to potential buyers but each purchase
yields additional information about the true quality. At each date, the monopolist begins
with a reputation for being a high-quality seller. She then chooses whether to sell.
Buyers then update the seller's reputation on the basis of this decision. We consider the
problem under different assumptions concerning the seller's knowledge about her own
quality, specifically we begin by examining feasible outcomes if she knows no more than
buyers, and if she knows her quality perfectly.

The central result of this chapter is that if the monopolist knows her true type there
is an equilibrium in which a good seller never stops selling. This result stands in stark
contrast to the case where her information is identical to buyers. In the latter case,

similar to Rothschild's (1974) two-armed bandit result, a long sequence of bad outcomes leads a good seller to stop selling. Thus the case where the monopolist knows her own type lends some support to Shakespeare's stylized view, and the case when she has the same information as buyers to Iago's.

Underlying the result is that when the seller knows her own type the market can make inferences from her decision to sell. In equilibrium, bad sellers play a mixed strategy below some threshold reputation level. Thus for reputations below this threshold, the act of selling improves a seller's reputation. Indeed, if the seller's reputation is below this threshold it will jump back to the threshold level if she chooses to sell. By definition the bad seller is indifferent at this threshold and so a good seller, who is more likely to see her reputation improve, prefers to sell. Thus a good seller will always sell. Moreover, the 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.[1]

Note that the bad seller plays mixed strategies at low reputation levels—no short-term fully separating strategy is feasible since this would lead buyers to believe that the seller must be good and so a bad seller would not wish to cease trading;[2] nevertheless, over time this threshold would be breached many times and since each time this occurs the bad seller would drop out with some probability, she would eventually stop trading. Thus asymptotically, there would be full separation.

We proceed to test the appeal and robustness of the signalling equilibrium characterized. 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 the mechanism described 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;

---

[1]The experience of bankrupts in Silicon Valley might provide anecdotal evidence of this kind of self-confidence mechanism in operation. 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.

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 equilibrium we identify is the most appealing one in a well-defined sense—specifically, a natural restriction 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.

2.1.1. 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 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 chapter is work which has taken a similar approach to this chapter 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, which model two-armed bandits, the bandits have a

strategic role through their ability to vary the prices they charge for each trial. The authors are chiefly concerned with the division of the costs and benefits of experimentation through pricing when information is symmetric. Here we are principally concerned with the consequences of asymmetric information and dynamic signalling by the decision to continue trading.

Ottaviani (1999), like this chapter, considers a situation where a monopolist effectively controls the learning process. In common with most learning models, though in contrast with this chapter, 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 cannot arise in their model where there is an implicit assumption that there is always a non-trivial efficient level of trade.

The introduction of private information to the bandit framework and the adverse selection that arises brings this chapter close to a wide literature on reputation and signalling. In particular, the 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. Money-burning through uninformative and costly advertising in their model plays a similar role to the willingness to incur losses in this chapter.

Tadelis (1999) presents a model with a similar framework to the one in this chapter, 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 chapter, we consider a monopolist seller and consider the influence of luck on her survival.

2.1.2. Plan of the chapter

In the following section, we introduce the framework for the models we go on to consider. In Section 2.3, we suppose that the seller has no private information regarding her own type, whereas in Section 2.4, we suppose that the seller knows her own type perfectly and present the central result of the chapter—that there is an equilibrium where a good seller always trades. Efficiency is also discussed; in particular, lowering expected profits at all reputation levels can allow arbitrarily close approximation to full efficiency.

In Section 2.5, we suppose that the seller has more information than buyers but that this information is not perfect. In Section 2.6, the seller can choose to cease trading only with some probability. In this case, we characterize equilibria qualitatively similar to the one described in Section 2.4.1. These equilibria are uniquely robust to a natural restriction on buyers' beliefs—that the reputation contingent on the seller's decision to trade should be non-decreasing in the reputation before this decision is made. Moreover, as the probability that the seller can behave strategically tends to 1 these equilibria tend to a single equilibrium—the one discussed in Section 2.4.1. Finally, Section 2.7 concludes. Details of proofs appear in Appendix A.

2.2. 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

---

[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 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.

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, 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] At the beginning of each period the seller first decides whether or not she will seek to trade.

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 $ex\text{-}ante$ 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 $ex\text{-}ante$ trade with a good seller is socially efficient and trade with a bad seller is inefficient. We suppose that buyers buy only once and that in each period there are at least two potential buyers and so, in each period, given their beliefs and that buyers Bertrand compete for the good, the price that emerges is simply their valuation for the good.[7] In particular, after observing the seller's decision to continue trading, buyers revise their beliefs to μ_t and so the price at which trade occurs will be $\mu_t g + (1 - \mu_t) b$. This framework allows us to focus on the signalling role of the monopolist's decision to trade without any interference from other signals; in particular, Bertrand competition among buyers determines prices, excluding any role for price signalling.[8]

---

[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.

[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]These assumptions are analytically convenient; however it would be sufficient that 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.

[8]Similar qualitative results can be obtained in a framework that allows for price signalling, though many more equilibria can arise, as discussed in the Appendix.

The success or failure of the good can only be determined after it has been bought and consumed. We suppose that this realization is publicly observable. Although this realization 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 realization 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 summarize 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) the seller then produces a single unit at a cost c and sells it, as discussed above, at a price $\mu_t g + (1 - \mu_t)b$ determined by Bertrand competition among buyers;

(5) the buyer then consumes the good and its realization as a success or failure is publicly observed; and,

(6) on the basis of this realization 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 $\lambda_{t+1} = \mu_t$).

A seller's strategy is simply a probability of trading given the beginning of period belief λ_t; buyers' strategies are the prices which they bid for the good, given their beginning of period belief (and the observation that the seller decided to trade and their beliefs about the seller's strategy). All the propositions discussed below characterize Markov Perfect Equilibria, in which the state is the beginning of period reputation λ_t.

Throughout we suppose that the information structure is common knowledge; for example in Section 2.4, 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.

2.3. 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:

$$\lambda^s = \frac{\lambda g}{\lambda g + (1 - \lambda)b}, \text{ and} \tag{2.1}$$

$$\lambda^f = \frac{\lambda(1 - g)}{\lambda(1 - g) + (1 - \lambda)(1 - b)}. \tag{2.2}$$

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:

$$
\begin{aligned}
V^u(\lambda) &= \lambda g + (1 - \lambda)b - c + \beta[(\lambda g + (1 - \lambda)b) \max\{0, V^u(\lambda^s)\} \\
&\quad + (1 - \lambda g - (1 - \lambda)b) \max\{0, V^u(\lambda^f)\}].
\end{aligned} \tag{2.3}
$$

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 she would have to 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 formalized 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.*

Although Proposition 1 is similar to a well-known result (see for example Rothschild (1974) and Whittle (1982)), the details of the proof appear in the Appendix as similar techniques are used for subsequent propositions.[9] The key elements of the proof we introduce 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 this 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.

2.4. Perfectly informed seller

This section begins by presenting and proving the central result of the chapter—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 2.4.1 and in Section 2.6 below, we argue that the equilibria we characterize in this section are appealing and robust equilibria.

---

[9]The framework here is similar to a one-and-a-half-armed bandit 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); however, note that here it is the arm rather than the experimenter who takes decisions and can gain from information.

2.4.1. 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, aware that the seller knows her type perfectly, therefore 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.

$$V^g(\lambda) = \mu g + (1 - \mu)b - c + \beta[g \max\{0, V^g(\mu^s)\} + (1 - g) \max\{0, V^g(\mu^f)\}]; \quad (2.4)$$

and

$$V^b(\lambda) = \mu g + (1 - \mu)b - c + \beta[b \max\{0, V^b(\mu^s)\} + (1 - b) \max\{0, V^b(\mu^f)\}]. \quad (2.5)$$

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 (2.3)), 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 (2.4) and (2.5) 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.[10] 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

---

[10] 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.

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 $\frac{g-c}{1-\beta}$, again this cannot be an equilibrium.[11] Thus it must be the case that, in equilibrium, for some reputation levels a bad seller would drop out with some probability and 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:

---

[11]Note that, it is at this point that there will be a difference in Sections 2.5 and 2.6, where in equilibrium there can be full separation of types (though note that in Section 2.5 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 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 2.5 and 2.6 for further details.

$$W^b(\lambda) = \lambda g + (1 - \lambda)b - c + \beta b \max\{0, W^b(\lambda^s)\} + \beta(1 - b)\max\{0, W^b(\lambda^f)\}. \quad (2.6)$$

Note that in Equation (2.6), since buyers believe that all types always sell, 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 $d(\lambda) = \frac{\lambda(1-\lambda^*)}{(1-\lambda)\lambda^*}$.*

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 (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.[12]

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

---

[12]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 $\lambda g + (1 - \lambda)b - c > 0$.
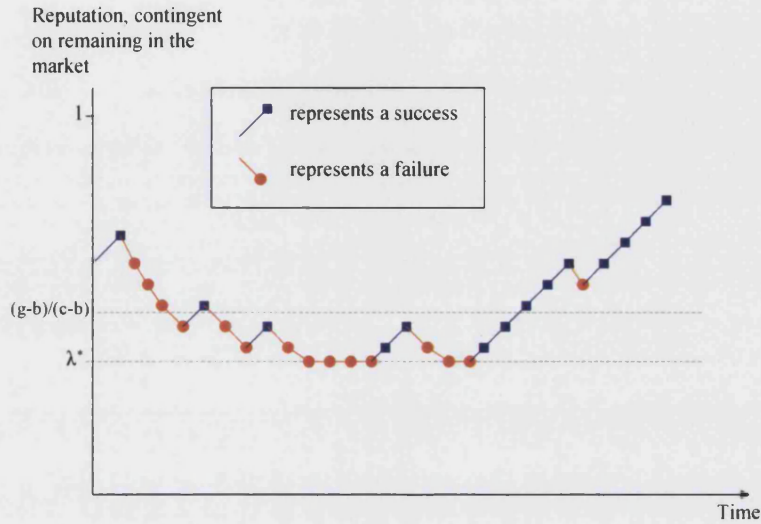
Figure 2.1. Sample path of evolution of reputation of a privately informed seller.

for example, since in the environment of this chapter 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 realized depends only on buyers' beliefs, so that arbitrary extremal 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 the more general environment considered in Section 2.6 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.

2.4.2. 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 2.3) or asymmetry of information (in Sections 2.4.1) can lead to two types of inefficiency, or "errors" in trading: (i) *Type I*: No trade with a good seller; (ii) *Type II*: Trade with a bad seller.

Thus the outcome in Section 2.3, 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 characterized 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 so that, as before, there would be no Type I error. We formalize 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 3.

Proposition 3. *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 since $\lambda^*(\gamma)$ is decreasing in γ and $d(\lambda, \gamma)$ is increasing in γ, 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.

2.5. 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 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 probabilistic sense at low reputation levels. This would imply that the interim reputation level—the reputation contingent on the observation that the seller is 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 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 the buyers' belief of the seller's type is $\lambda_0 = r_0\overline{\lambda}_0 + (1 - r_0)\underline{\lambda}_0$. 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 2.4, 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 realizes. 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 inter-related 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 commonly known.

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. It is clear that a seller with a good signal cannot stay in for longer than if this signal were public and in this latter case Proposition 1 applies. In particular, Proposition 1 implies that there is the set of histories, $\overline{H}$, for which a seller with a high (public) signal stops selling is not empty.[13] Nevertheless, in this environment the signalling role of choosing to trade still operates. We discuss this below.

In 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. 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

---

[13]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 $\overline{H} \subset \underline{H}$.

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 4. *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 self-confident signal provided by the decision to continue selling. 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 influences buyers' beliefs, so that these rise above the prior belief. However, the buyers' belief can be no higher than if the 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 few other aspects of the equilibrium described in Proposition 4 are noteworthy. First note that similar to the equilibria described in Section 2.4, 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 5. *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 2.4) 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.

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.
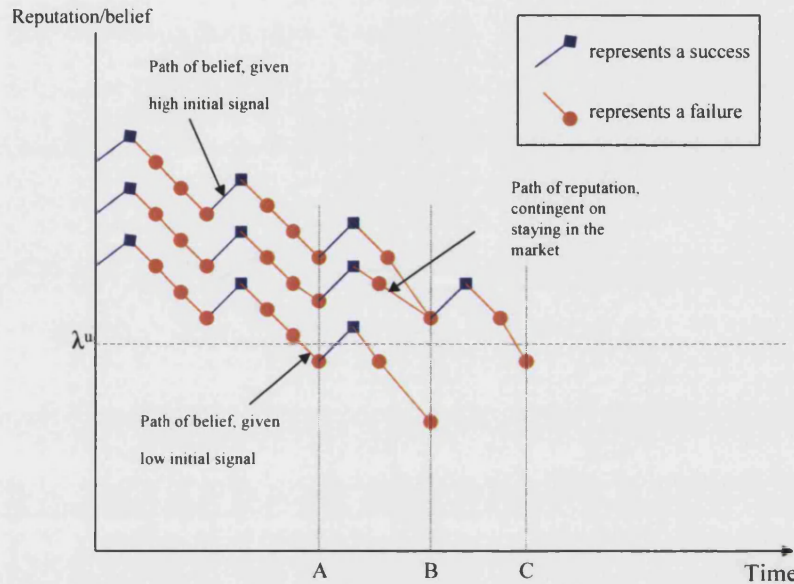
Figure 2.2. Evolution of beliefs of high signal seller, low signal seller and customers based on a sample history.

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).[14]

Figure 2 illustrates the result characterized in Proposition 4. 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

---

[14]See the discussion in the final paragraph of Section 2.4.1 for more on this issue.

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 $\overline{\lambda}$.

2.6. Varying the degree to which the seller can behave strategically

In this section, we further consider the robustness of the mechanism introduced in Section 2.4 and, in particular, the equilibrium characterized 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 characterized in Proposition 2 is the unique limit of equilibria among that class of equilibria. Moreover, in the more general environment of this section the signalling role of continuing to trade continues to operate in a similar way.

We suppose throughout this section that the seller is fully informed as to her type, as in Section 2.4, 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 cannot 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 2.4) 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.[15] 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 2.4. Suppose that $\alpha \in (0,1)$ then we begin by stating the following result, which characterizes a seller's behaviour when given the opportunity to make the strategic decision to continue or cease trading.

---

[15]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.

Lemma 6. *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—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, 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 argue that a natural restriction on beliefs lead to particular classes of equilibria, which in the limit converge to the equilibrium characterized in Proposition 2, though at $\alpha = 1$, this restriction is not sufficient to ensure uniqueness. Specifically we make the following assumption:

Assumption 1: *Buyers' interim beliefs are non-decreasing in their prior (beginning of period) beliefs, i.e. $\mu(\lambda)$ is non-decreasing in λ.*[16]

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

Lemma 7. *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:*

---

[16]First to understand the necessity of this restriction, note that at an intermediate range of reputation, buyers' more optimistic or pessimistic beliefs concerning whether or not good types seek to trade could be consistent with the seller's behaviour. Optimism or pessimism could in principle vary in a discontinuous way and so in some cases there might be greater value in having a worse reputation so that at some reputation levels this might favour a bad type of agent (that is a bad type of agent would have higher continuation value). This might lead to an instance where $\mu(\lambda)$ was actually decreasing in some range. Note also that this assumption places a restriction on an endogenous variable, the interim belief. There may be a more fundamental restriction that delivers a similar result; however, this assumption seems a reasonable one.

- *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 hypothesized 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 8. *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 proof 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 characterize properties of these value functions. Finally, sufficient conditions to ensure equilibrium can be summarized as conditions on l and m, the proof concludes in showing that these conditions can be satisfied with $0 < l < m < 1$.[17]

---

[17]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 $m(\alpha +$

Proposition 8 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 arises. 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.[18]

We argue, however, that as α tends to 1, the limiting case is unique, moreover the unique limiting equilibrium is the one characterized in Proposition 2. Note though, that in the case $\alpha = 1$, there are a multiplicity of equilibria of the form described, consistent with Assumption 1. 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 hypothesized form with $l = m = h = 1$.

Proposition 9. *As α tends to 1, the set of feasible equilibria of the hypothesized form (that is as described in Lemma 2) converge to a singleton. Moreover, this is the equilibrium characterized 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 characterizing an upper bound for this set (trivially the lower bound is 0) and showing that this upper bound tends to 0.

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

$1 - \alpha)/[m(\alpha + 1 - \alpha) + (1 - m)(1 - \alpha)] = m/[m + (1 - m)(1 - \alpha)]$. Further note that m is the lowest prior reputation for which this interim reputation could be attained.

[18]In the finite horizon version of the model the effect can be isolated within a particular time period and so we can explicitly show the multiplicity of equilibria. (See the Appendix for further details)

2.7. Conclusion

Returning to the quote from *Othello* at the beginning of this chapter, 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 chapter 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 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, the signalling role of the decision to trade and its ability to ensure the survival of a good seller has been shown to apply both 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.[19]

The framework considered in this chapter is quite general. There are numerous markets, such as legal services, consulting, garage mechanics and other credence goods, where quality is learned only slowly and luck can play an important role in the development of reputation. In this chapter, we have highlighted the signalling role that the decision to participate in the market can play in the case of a monopolist seller. In more competitive or oligopolist markets, the results of this chapter might also give some insight. In such environments, a seller's trading decision would play a signalling role when she has private information though such a signal, as in the monopoly case considered here, could not be fully revealing. The extent to which this signalling role of the trading decision might lead to the long-run survival of the most efficient firm, paralleling the central result of this chapter, is left for further research.

---

[19]Similar results also hold in a finite horizon framework, where a good seller trades with probability 1 in all but the final period, as illustrated in the Appendix.

CHAPTER 3

Something to prove

3.1. Introduction

In professional services industries, such as law and consulting, products cannot be inspected prior to sale and it is impossible to make fully contingent descriptions of the product.[1] Therefore in these industries developing and maintaining a reputation is crucial. This is true both at the level of the firm and the individual. Indeed, previous literature has suggested that the very existence of a firm might arise as a means to manage reputation and that careers are designed to take into account reputational considerations.[2] This chapter seeks to develop the notion of the coexistence of firm and individual reputations and the implications for individuals' incentives in a simple framework in which young agents are motivated by their own reputations and old, successful agents are motivated by the reputations of the firms which they own (or more specifically the reputation of her employee).

For an agent to be motivated by reputational concerns it is tautological that some reputation must be at risk and that her actions can affect this reputation. When an agent has an established reputation, such concerns for her own reputation are muted and so with no other commitment device, the agent would exert no effort. Loosely for an agent to be motivated by reputation she needs to have something to prove.

We argue that when she is young an agent is motivated by trying to show that she is competent but when she is older and has proven herself, she is motivated by trying to show that the firm which she owns (and in particular her employee) are competent since she profits from the reputation of her firm. Working with a junior allows an established agent to credibly commit to exerting effort since it provides her with something to prove—the ability of her junior. Thus while much literature (in particular following

---

[1]It is relatively difficult for a consultant to show a client the report that she would write if hired or for a client after the event to complain that he did not receive the report that he was expecting.

[2]For examples, see Kreps (1990) on the first point and Gibbons and Murphy (1992) on the second.

Holmstrom (1982)) has focused on the negative effect of teamwork on incentives through the free-rider problem, we show that teams might have a positive impact on incentives and indeed with no opportunity for teamwork all incentives for effort in the models introduced below disappear.

The focus of the chapter is the provision of incentives for a senior whose own reputation is certain and who hires a junior to exploit uncertainty about this other agent. In essence for this to have an effect on the senior's incentives, the senior must be able to affect the junior's reputation—teamwork allows for this—and the senior must gain from the junior's reputation—in this model this is achieved by supposing that the senior controls a client list without which the junior cannot gain from any reputation she has established.

In the simple example of Section 3.3, this intuition is made clear by supposing that rather than working with another agent, the senior buys a machine which produces output of uncertain quality. In particular, she will be the residual claimant of the machine's future reputation directly through ownership and there is no moral hazard or incentive problem for the machine. This simple example highlights the role of non-observability in team production. Customers cannot distinguish between the machine's and the senior's contributions to total output and so the senior's choice of effort will affect the way the machine is perceived and the revenue that the senior receives for its product in the following period. The example thus demonstrates how the strategic introduction of uncertainty (in this case through the uncertainty about the type of the machine which the senior chooses to buy and work with) allows the senior to commit to exerting effort.

This simple example is intended to provide intuition for the richer environment in Section 3.4, which is the focus of this chapter. Here rather than having the opportunity to buy a machine, the senior has the opportunity to hire a junior. Team production, ensures that the senior has the opportunity to affect the junior's reputation. The ability of the senior to gain from the junior's future reputation, arising from the *ex-post* value of access to clients, allows the senior to commit to exerting effort and thereby overcomes her incentive problem. In this environment, however, the junior also has an incentive problem. In order to overcome this, the senior writes a long-term contract with the junior. This contract is not contingent directly on outcomes (if such contracts could be written then there would be no need for reputational considerations to overcome the senior's incentive problem); instead the contract specifies a price at which the junior can buy the firm. In equilibrium only a successful junior will find it worthwhile to buy the

firm at this price, ensuring that the senior has an incentive to work to allow the junior an opportunity of succeeding; moreover the price is set at a level which ensures that the junior has an incentive to work so long as she has the opportunity to hire her own junior, as is the case in the overlapping generation framework considered. An equivalent formulation of the reward structure is an up-or-out contract with promotion leading to the promoted junior taking over the firm.

There are two principal contributions in this chapter. First, while other papers (notably Holmstrom (1999)) have highlighted that uncertainty is necessary for reputational concerns to affect an agent's behaviour, I believe that this chapter is the first to suggest that an agent can strategically choose to introduce uncertainty in order to make a credible commitment to exert effort, even though there may be no uncertainty about her own ability. Moreover, this approach complements previous literature which has considered the interaction of firm and individual reputation. Such literature has largely either assumed that the firm consists of a single individual or of so many that a single individual has no influence on the firm's reputation as a whole.[3] This chapter can be thought of as addressing the case where the firm is a small team of individuals so an agent's decisions could affect both her own and her co-worker's reputation. Viewing the firm as a small team in which reputations interact leads to the model presented where at the beginning of their careers agents are motivated by their own reputations as individuals but as they get older and succeed, they become motivated by concern for the reputation of their firms.

Second, it is then shown that this theoretical contribution might have implications for the structure of professional services organizations such as law firms. In particular, we show that teamwork can create rather than dampen incentives since mixed teams of partners and juniors can provide incentives for partners and that in this framework, for an up-or-out mechanism to be effective promotion must be to partnership—an empirical feature which previous literature has not much addressed.[4] Finally, whereas some have argued that law firm partnerships might exist to diversify risks for individuals, in practice one sees groups of lawyers working in the same or related fields. An established argument for this phenomenon is that it allows for mutual monitoring among the partners in a law

---

[3]See Breton et al (2002, 2003) and Anderson and Smith (2002) for recent exceptions which focus on what types of agents work together. In these papers agents do not have any effort decision to make and so these papers cannot address reputational incentives to exert effort—the focus of this chapter.

[4]For an exception and elaboration on this criticism of previous literature, see Rebitzer and Taylor (2001).

firm (Alchian and Demsetz (1972)). Note in addition that a more homogeneous firm makes it more difficult for clients to identify individual contributions of seniors and juniors and so enables the reputational mechanism highlighted in this chapter to operate and seniors to credibly commit to exerting effort.

In the following section, related literature is discussed in some detail. In Section 3.3, we present a very simple model to illustrate how a senior, with no reputational incentives when working alone, can exploit a joint production technology and uncertainty about a machine in order to credibly commit to more efficient behaviour. This example is intended primarily to illustrate the intuition underlying the richer model in Section 3.4 which is an overlapping generations model in which seniors can hire and work with juniors. An equilibrium in pure strategies is characterized. In particular, the equilibrium, which allows for the mechanism highlighted in Section 3.3, also generates interesting career dynamics. Section 3.5 discusses and provides some support for a number of the modelling assumptions. The final section concludes and discusses the robustness of the results.

3.2. Related literature

3.2.1. Reputation

Kreps, Milgrom, Roberts and Wilson in a series of papers published in 1982 and developed by Fudenberg and Levine (1989) and (1992) pioneered the modern literature on reputation. In these models observers cannot distinguish between a strategic type behaving well and a "crazy" type who behaves in a way in which the strategic would like to be able to commit (Fudenberg and Levine (1989) therefore term this a "Stackelberg" type).[5] That is, since the strategic type has an incentive to convince the public that she is a Stackelberg type, in equilibrium there is uncertainty about an agent who behaves in the way that a Stackelberg type would behave, as this might be a strategic type mimicking. Thus there is no learning about the agent's type over time and the uncertainty necessary to maintain reputational incentives arises naturally and is maintained.

However, the assumption that such a Stackelberg type exists might be unrealistic in some applications of interest where reputational considerations play an important role (for example in the context of lawyers this would correspond to a lawyer who was successful without exerting effort, or enjoyed exerting effort). In this chapter, rather

---

[5]In the model below a Stackelberg type would be a competent agent who always exerted effort.

than suppose the existence of a Stackelberg type, we suppose instead that there is an inept type who is likely to fail even if exerting effort and a competent (strategic) type who seeks to differentiate herself and simultaneously commit to her preferred action, following an approach employed in Mailath and Samuelson (2001).[6] In this environment, if a competent agent has convinced the public that she is competent, then she has nothing to prove about herself and so cannot commit to the preferred action and does not undertake it. It is argued that by working in a team a senior can strategically introduce uncertainty which allows her to commit to the preferred action.[7]

3.2.2. Individual and firm reputation

The role of an individual's concern for her own reputation has been explored through the literature on career concerns beginning with Holmstrom (1999) who shows that with no type uncertainty there are no reputational incentives to induce effort. A recent literature (see for example Tadelis (2002) and Mailath and Samuelson (2001)) has considered the role that a firm's reputation can play in providing incentives, though in a context where individuals cannot be identified so that individual reputation cannot play a role.

In two papers that deal explicitly with the interaction of individual and collective reputations (Tirole (1996) and Levin (2001)), groups are very large so that a single individual's actions can have no effect on the collective reputation and so no agent takes action to try to influence this collective reputation. In this chapter, we consider a small group—consisting of only two individuals—and while the team does not have a reputation separate from the reputations of its constituent members, the actions of each member can affect the reputation of the other; indeed this is at the heart of the chapter. Here, over her lifetime an agent is motivated first by concern for her own reputation and later for the reputation of her firm (effectively her co-worker).

---

[6]Furthermore Cripps et al. (2004) show that even with such a Stackelberg type if there is imperfect monitoring of actions, then in the very long run customers would learn an agent's true type and reputation effects would disappear. They view their results:

> as suggesting that a model of *long-run* reputations should incorporate some mechanism by which the uncertainty about types is continually replenished.

One contribution of this paper is to suggest an *endogenous* mechanism for achieving this.

[7]That it is the senior who introduces uncertainty contrasts with other literature in which a principal (client) who deals with an agent benefits by taking steps to remain relatively uninformed. For papers building on this possibility see Cremer (1995) and Meyer et al (1996).

At the heart of this chapter is the observation that some uncertainty whose resolution is influenced by the agent's actions is necessary for reputational incentives and that agents can *strategically* introduce uncertainty by choosing to work in teams and hire juniors of uncertain ability. Holmstrom (1999) and Phelan (2001) consider the possibility that there is a persistent chance of an exogenous type change which can sustain incentives. In Mailath and Samuelson (2001), the reputational mechanism is maintained since it is firms rather than individuals who have reputations and in each period there is a chance that it is a different individual working under the same name, creating persistent type uncertainty, and thus an agent always has something to prove about herself. In this work, and a wider literature beginning with Tadelis (1999) and further developed in Tadelis (2002, 2003), the probability that an agent is replaced is exogenous or not the focus of analysis; the focus in these papers is rather that the type of the agent replacing the existing agent is endogenously determined in a market for names. Essentially, the persistent possibility that the firm or individual's type changes (or at least the type of the owner currently associated with that name) ensures that beliefs never get so close to certainty that today's outcomes have too little influence on tomorrow's beliefs and so reputational considerations can outweigh the direct costs of exerting effort.

It is further worth noting explicitly that the work discussed above and based on name trading relies heavily on the *intertemporal non-observability* of name transfers.[8] In contrast, this chapter relies on the *contemporaneous non-observability* of joint production and the focus is on the decision of the senior agent's decision of whether or not to hire, rather than the junior agent's decision to join one firm rather than another or buy one name rather than another.[9] In particular, this feature allows equilibria both in the case where agents know their own types and in the case where all information is *ex-ante* identical (since an agent will have private information about her choice of action within the period).

---

[8]This is at the heart of Tadelis's distinction between entity and identity.

[9]Note that in this paper, the decision of whether or not to hire is considered. It is assumed that the senior cannot screen juniors and that all juniors are ex-ante identical, so that the question of who to hire does not arise. This latter question has been considered elsewhere, in particular Segendorff (2000) and Glazer and Segendorff (2001) consider how reputational concerns might, for example, lead good seniors to hire bad juniors when the relationship between rewards and reputation are non-linear.

3.2.3. Partnerships

A number of papers define partnership as an equal sharing of revenue among partners. In particular, such equal sharing can lead to diversification and risk sharing (Gilson and Mnookin (1989)), an inefficient size of the firm (Farrell and Scotchmer (1988), Levin and Tadelis (2002)) and inefficient composition (Sherstyuk (1998)). Largely, these results would apply to any form of fixed profit sharing which implies redistribution among partners. However, it is not clear that fixed shares of revenue characterize professional services. Cotterman (2001), for example, discussing law firms states that:

> Historically a prominent compensation method, the lockstep approach
> is now the least preferred way of allocating compensation.

Cotterman argues that performance-based measures and to some extent compensation based on the revenue that a partner has generated ("eat what you kill") are important features of partners' compensation.

A rather different approach is taken by Garicano and Santos (2004) who focus on how partnership arrangement can help to overcome the problem that individuals might prefer to keep hold of and work for a client even though it may be more efficient to refer the work to another professional who may be better qualified to deal with the problem. Reputation does not arise in their essentially static model, indeed their first assumption is that clients cannot self sort (determine a particular lawyer's ability to help them) and thus can be thought of as a complimentary approach to this chapter.

Rebitzer and Taylor (2001) focus on employees threatening to "grab and leave" with an important client, abstracting from other incentive problems, and suggest that the up-or-out system has evolved as a resolution to this problem. In particular, they highlight that in their model up-or-out implies promotion to partnership—a feature absent from other models of up-or-out, though note that this feature does arise naturally in the model below. Two important assumptions in their model are that an employee can walk away with a customer and that a team has to be sufficiently large in order to be feasible. In contrast, the model in this chapter and in Appendix C are based on very different assumptions, specifically that an employee cannot walk away with her reputation intact and that there are no technological productive advantages or disadvantages to teamwork, the only implication difference between two individuals working alone and working in a team is that in the latter case output cannot be attributed to one or other of the individuals.

In contrast to all these papers, which focus on the role of partnership in distributing revenue among a number of individuals at the end of a single period, this chapter highlights a different aspect, specifically the phenomenon of promotion to partnership, whereby the junior buys into the partnership and takes over the firm from an existing owner.

In a related, independent paper, Morrison and Wilhelm (2003) consider the role of partnerships in ensuring that seniors mentor their juniors. In their paper a partner has an incentive to mentor a junior as only a good junior who has been mentored would be willing to buy a partnership share in the firm, as only in this case will the junior be able to maintain the firm's collective reputation and the value of her partnership stake. Thus the paper echoes a number of the themes highlighted here, in particular, a senior is motivated to work since this affects her ability to sell the firm to the junior. However, the issue at the heart of Morrison and Wilhelm (2003) is the incentive for partners to mentor juniors—an incentive which naturally leads seniors and juniors to work together— and they consider no incentive issues for the juniors. Here instead, seniors and juniors are identical and face similar decisions—it is not the case that the senior affects the productive capability of the junior, as in Morrison and Wilhelm (2003), but rather can affect how this is perceived; moreover we show that the promotion structure described can be successful in providing incentives for the junior as well as for the senior. Further their model has no role for individual reputations.

3.2.4. Other related literature

Cremer (1986) presents an interesting and related OLG model in which younger agents work hard in order to perpetuate a culture of hard work from which they will benefit when older. More specifically, the model is based on a dynamic game with multiple equilibria in which a young agent's cooperative behaviour ensures that future agents coordinate on the cooperative equilibria, from which the agent will benefit as a result of equal sharing among all agents in the organization.

This chapter is also related to a wide literature on how jobs are designed to influence workers' incentives (see for example Prendergast (1999)). Of most relevance, in addition to Holmstrom (1999), is Jeon (1996) in which incentives arise from career concerns and it is assumed that agents must work in teams but the composition of the teams is endogenously determined; in particular it is shown that mixing young and old employees is advantageous since this allows the ability of the young agents to be observed more

accurately and so provides them with greater incentives (in Jeon's model old agents cannot be incentivized).

Finally, the thinking of this chapter generates a number of somewhat counter-intuitive results in the context of other literature. First, in this model when a junior and senior work together reputational considerations may lead to the senior leveraging the uncertainty concerning the junior. Intuition from the literature on umbrella branding (Wernerfelt (1988) and Choi (1998)) suggests in contrast that when juniors and seniors work together the senior leverages her own good name and extends it to the junior. Similarly, while the up-or-out literature has largely focused on the incentives of juniors, here it is argued that the up-or-out mechanism can also help ensure effort on the part of seniors.[10] Finally, in this model the introduction of teamwork (joint production) through the reputational incentives that it allows, creates incentives for effort that would not otherwise be present. Thus dynamic considerations here act in opposition to the free rider problem highlighted in Holmstrom (1982) (though note that here we consider an environment of implicit dynamic contacting rather than an explicit single-period contract) and teams can create rather then reduce incentives—based on dynamic reputational considerations rather than altruistic preferences or other peer effects.[11]

3.3. A motivating example

In this section, we introduce a simple model to demonstrate that a senior, an agent with an established reputation, can exploit uncertainty about something else, when there is a joint production process which does not allow individual contributions to be observed.

Specifically, consider an agent, the senior, who lives for two periods. There are many customers Bertrand competing for the good that the senior produces which may be of high or low quality depending on her effort decision, as discussed below, and for the good produced by a machine of uncertain quality. Timing is as follows.

---

[10]This literature is discussed above. In the context of law firms, see in particular Galanter and Palay (1991). A paper which does highlight the incentive effects of up-or-out on seniors is Carmichael (1988) but the incentives highlighted are those to hire appropriately rather than incentives for production.

[11]Meyer, Olsen and Torsvik (1996) also show that teamwork—precisely because it clouds the inference that can be drawn on a particular agent—might be preferred to solo production. However, in their model, this is to dampen the ratchet effect in a model of explicit outcome-contingent contracting with limited ability to commit to long-term contracts. This is quite different to this paper's model, in which agents are risk neutral and it is not possible to write explicit contracts, and where reputational considerations rather than explicit outcome-contingent contracts are the driving force behind effort. Che and Yoo (2001) consider that agents within a team can monitor each other's actions better than a principal and over a long period dynamic strategies can allow team members to enforce good behavior.

Period 0: There is a machine which either always produces high quality or always low quality products, but can only produce one unit in each period and does so at no cost. The machine owner together with the senior and customers believes that with probability λ the machine is the type that always produces high quality. Thus, supposing no discounting between periods, a risk neutral machine owner who worked the machine independently would have an expected present value of 2λ.

The senior can make a take-it-or-leave-it offer to the current owner.[12] Thus she can offer 2λ to buy the machine for the 2 periods, collecting the revenue for the machine's output.

Period 1: Customers Bertrand compete for the output produced by the senior and the machine. They assign the value 1 to a high quality product and 0 to low quality and compete to buy the unit so that the sale price is exactly the customers' belief that the product will be successful, since sale occurs before the quality is realized and the price cannot be made contingent on the outcome.

The senior decides whether or not to exert effort at a cost c. If she exerts effort, she will produce a single unit which will be a high quality product with probability $g \leq 1$ and otherwise low quality. If she exerts no effort then the product will be low quality for sure. Quality is non-verifiable and not observed prior to purchase so that the price cannot be contingent on quality but will depend on customers' expectations.

Finally suppose that $g > c$ so that the costly effort is efficient. In particular, if effort were contractible or the senior could credibly commit to it and be compensated for it, then she would exert effort.

Period 2: The quality of the goods produced in period 1 is observed. If the senior bought and worked with the machine then the output cannot be directly attributed, that is if a one high quality good and one low quality good are observed, customers do not know which was produced by the senior and which by the machine. Again customers Bertrand compete for the output produced.

This timing is summarized in Figure 3.1.

---

[12]Note that the assumption of a take-it-or-leave-it offer on the part of the senior is not crucial, similar qualitative results could be generated if the machine owner had the ability to make a take-it-or-leave-it offer or under any intermediate bargaining power assumption and bargaining game. The key point here is that the total value to the senior from working with the machine is greater than the value of them working separately, and that senior is the residual claimant of at least some of the machine's reputation in the second period.

```
├──────────┼──────────┼──────────┼──────────┼──────────┼──────────┼──────────┤
```

Period 0 | Machine | Period 1 begins | Senior | Quality of | Period 2 begins | Senior | End

Period 0

Senior makes take-
it-or-leave it offer to
machine owner

Machine
owner
accepts or
refuses
offer

Period 1 begins

Customers
Bertrand
compete for
good(s)

Senior
exerts
effort or
not

Quality of
good(s)
observed

Period 2 begins

Customers
Bertrand
compete for
good(s)

Senior
exerts
effort or
not

End

Figure 3.1. Timeline for example

First suppose that the senior has no opportunity to buy and work with the machine, then the unique subgame perfect equilibrium of this game is clear—the senior exerts no effort in either period. This is her optimal action in the final period, and so in the penultimate period as well.

However, when the senior can buy and work with the machine, she can commit to exert effort in Period 1 (though not in Period 2). First note that the total output when the senior and the machine work together is simply the sum of their output. Thus the expected value of the output when the senior puts in effort and works with the machine is 0 with probability $(1-g)(1-\lambda)$, it is 1 with probability $g(1-\lambda) + \lambda(1-g)$ and it is 2 with the residual probability $g\lambda$. If the senior does not exert effort then the expected output is 0 with probability $(1-\lambda)$ and 1 with probability λ, and there is no chance that the output might be 2. When the senior works with the machine, however, it is assumed that the production process is unobservable in the sense that when a single unit is produced, customers cannot tell whether it was produced by the senior or the machine. In particular, this implies that if customers expect the senior to exert effort in period 1 and she does not, then the machine's reputation will in expectation be lower than it ought to be—this could induce the senior to exert effort and so allow an equilibrium in which the senior credibly commits to exerting effort. This is the intuition underlying the following result.

Proposition 10. If $\frac{g(1-\lambda)\lambda}{(1-g)\lambda+g(1-\lambda)} > c$ then there is a pure strategy subgame perfect equilibrium in which the senior exerts effort in period 1. If in addition $g(1-\lambda) > c$, this is the unique pure strategy subgame perfect equilibrium.

Proof. Consider the following strategy: The senior buys the machine for 2λ. In the first period the senior works with the machine and exerts effort. In the second period (the senior's "retirement" period) the senior works alone and exerts no effort and puts the machine to work alone, collecting the fee for its output. Buyers believe that the senior

behaves in this way and bid for the good produced accordingly, revising their beliefs on the quality of the machine according to Bayes rule after observing the quality of the Period 1 goods at the beginning of Period 2.

It is clear that the price of the machine is an equilibrium price and that the machine owner's behaviour in equilibrium is optimal. For the senior the strategies describe optimal behaviour so long as she could do no better either by working on her own from the first period (the value of which as described above is 0) or by not exerting effort in the first period.[13] Her value from sticking to the strategy is

$$g + \lambda - c - 2\lambda + g\lambda * 1 + ((1-g)\lambda + (1-\lambda)g)\frac{(1-g)\lambda}{(1-g)\lambda + g(1-\lambda)} + (1-g)(1-\lambda) * 0 \quad (3.1)$$

This expression can be explained as follows. Given the anticipated equilibrium strategies the first period revenue is $g + \lambda$ but the senior must incur the cost c of effort and buy the machine for 2λ. In the second period, when the output is 2 (which will happen in equilibrium with probability λg) the public is certain that the machine is good and will pay 1 in the second period; when the output is 0 (which happens with probability $(1-g)(1-\lambda)$) the public is certain the machine is bad and is prepared to pay 0 in the second period; and, when the output is 1 the belief that the machine is good (and so also the second period revenue) is $\frac{(1-g)\lambda}{(1-g)\lambda + g(1-\lambda)}$. The expression in (3.1) can be simplified to $g - c$; it is simply the value of the senior committing to exert effort in the first period and buying the machine for exactly its expected output.

The value when the senior defects and exerts no effort is given by:

$$g + \lambda - 2\lambda + \lambda\frac{(1-g)\lambda}{(1-g)\lambda + g(1-\lambda)}. \quad (3.2)$$

This follows since she receives $g + \lambda$ in the first period as customers expect her to exert effort but does not exert effort or incur its cost, but now there is zero probability of the realized output in the first period being 2 and the probability of it being 1 is λ. So the total value of deviating and not exerting effort when it is anticipated is as appears in (3.2). The value when she defects by not hiring is 0.[14]

---

[13]It is clear that there is there is no second period strategy that dominates the one described.

[14]We suppose that on this off-equilibrium path customers believe that the senior will not exert effort, this would have to be the case, for example, if imposing a trembling-hand refinement.

So the strategies do indeed describe an equilibrium so long as $g - c > 0$ (following (3.1)) which is assumed to be true—otherwise effort would be inefficient and so long as the senior prefers to exert effort in the first period so that (3.1)>(3.2):

$$g - c > g - \lambda + \lambda \frac{(1 - g)\lambda}{(1 - g)\lambda + g(1 - \lambda)} \tag{3.3}$$

or equivalently

$$\frac{g(1 - \lambda)\lambda}{(1 - g)\lambda + g(1 - \lambda)} > c. \tag{3.4}$$

Uniqueness in pure strategies:[15] Now suppose that it is believed that the senior exerts no effort even when buying the machine. Suppose that the senior buys the machine and exerts no effort, the expected value of this strategy is 0. Suppose that the senior deviates by exerting effort, this yields the value:

$$-2\lambda + \lambda - c + g(1 - \lambda) + \lambda(1 - g) + g\lambda, \tag{3.5}$$

where this expression can be explained as follows. The senior buys the machine at a cost 2λ and since the machine is expected to produce high quality with probability λ and she is not expected to exert effort, the first period revenue is λ, in addition she incurs the cost of effort. In the second period, whether they see one or two high quality products, customers will believe that the machine always produce high quality. The above expression can be re-written as $g(1 - \lambda) - c$ and so the senior deviates and exerts effort when $g(1 - \lambda) > c$. □

It is clear that parameter values exist for which condition (3.4) can be satisfied, and that it is more likely to be satisfied the smaller is c and the larger is g—this follows naturally, since the larger g or the smaller c the greater the gain from taking the costly action. However the comparative statics with respect to λ are not monotone. The intuition underlying this non-monotonicity reflects the discussion in the introduction of the need for "sufficient" uncertainty, or having something to prove. At extremal values of λ whether or not the senior exerts effort would not change the public's belief about the machine by much.[16] Note in particular that for λ close enough to 0 or 1 then the left

---

[15]It can be shown that there is a mixed strategy equilibrium where the senior buys the machine and exerts effort with probability $q = \lambda \frac{c + \lambda g - g}{cg(2\lambda - 1)}$ so long as $q \in (0, 1)$ and $(2 - \lambda)cg - g^2(1 - \lambda) - c^2 \geq 0$.

[16]Specifically, note that in the case following an observation that the output is 1 the public belief about the machine in the equilibrium is $\frac{(1 - g)\lambda}{(1 - g)\lambda + g(1 - \lambda)}$; for λ close to 0 or 1 this is close to λ.
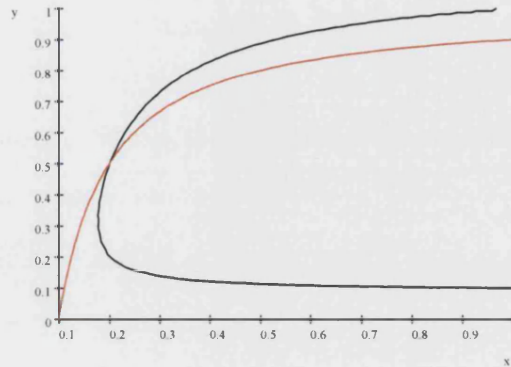
Figure 3.2. A Plot of lambda against g: Existence and Uniquencess at c=0.1

hand side of inequality (3.4) is close to 0 so that the condition fails. This is illustrated in Figure 3.2, which illustrates combinations λ and g for the case $c = 0.1$: an equilibrium in which the senior exerts effort exists to the left of the black line line (this is the unique equilibrium in pure strategies if in addition the (λ, g) pair is below the red line).

Note that the senior's profits in an effort-inducing equilibrium are independent of λ—so long as it is in the range that supports the equilibrium—the senior's profits over the two periods are given by $g - c$ which is simply the value of the senior committing to exert effort in the first period, (the other equilibrium revenues correspond to the senior buying the machine for exactly the value of its expected output which she sells).

As a final observation in this section, note that non-observability in joint production is crucial. Here non-observability arises in the sense that when the joint output is 1, the public does not know whether the single unit was produced by the machine or the senior. If instead this is observed, then the equilibrium described above breaks down—the senior's action would have no influence on the reputation of the machine and on the price for which its product is sold in the second period. In particular, if the senior exerts no effort and it is fully observed that the machine produces a single unit in the first period, then the machine's output can still be sold for 1 in the second period so that the senior loses the incentive to exert effort in the first period.

3.4. A richer example with OLG agents

The above example demonstrates that a senior can exploit the uncertainty about a machine in order to commit to an efficient, though costly action through concern

about the reputation of the machine. However, the above example, though suggestive, is inadequate for many applications. In particular, in trying to apply this reputational mechanism to professional services, it is of more interest to think about introducing uncertainty through working with other agents who also need incentives to exert effort rather than machines with no such need. In this section, we consider a richer framework which treats seniors and juniors in a symmetric and consistent way.

We introduce a framework with overlapping generations of agents with a two period working life then we characterize an equilibrium, which is the focus of this chapter, where agents exert costly and efficient effort in the first period of life and in the second period only if their "firm" is successful in the first period.

3.4.1. Model set-up

In an infinite period framework, agents work for two periods, and can consume in a third "retirement" period. A new generation of equal size is born in each period. An agent may be either competent or inept.

Production Inept agents, whether exerting effort or not, produce low quality products for sure as do competent agents who exert no effort. A competent agent who exerts effort at a cost c produces a high quality product for sure.[17]

An agent can either work on her own, producing as above, or alternatively, an agent can work in a team with another agent, in which case the total output is the sum of the output of each agent.[18] However, in the case of joint production, if a single unit is produced then the public cannot tell which of the agents produced it—this assumption ensures that when working together the senior's choice of effort can affect the reputation of the junior.[19]

---

[17]Assuming that inepts and competents who exert no effort produce high quality with probability $b > 0$ and competents exerting effort produce high quality with probability $g < 1$ leads to qualitatively similar results. See also Footnote 30.

[18]Allowing agents to work alone as well as in teams is a convenient modelling assumption. Without it, in the case where a competent agent exerting effort succeeded with a probability less than 1, the number of successful agents (that is in teams with two successes) in each generation would fall. This assumption, thus, allows us to consider a steady state with a positive fraction of successful agents.

[19]Note that this assumption can be weakened, in particular partial non-observability would do. For example, the case that the senior's output could be seen independently of the junior's but a junior only gets chance to prove herself if the senior exerts effort would lead to similar results.

All customers are willing to pay 1 for a high quality product and 0 for a low quality product. All agents are risk neutral, maximise their expected lifetime earnings and have a discount factor of 1.[20]

Information At birth, an agent does not know her own type and believes that there is a probability λ that she is competent. Potential employers and customers also believe that there is a probability λ that any new born agent is competent. It is assumed that $\lambda > c$ so that inducing effort from new born agents is efficient.[21]

There are very many *ex-ante* identical locations; in particular, a new born agent can always find a location where there is no existing agent at which to found a new firm. At each location there are many identical customers who Bertrand compete for the product of the agent or agents at that location, so that the revenue that an agent generates when working alone, for example, is equal to the customers' belief that the agent will produce high quality.[22] An agent's productive history (that is the sequence of high or low quality produced by the agent working alone or by the agent and her co-worker when working in a team) is assumed to be non-verifiable but observable, though only at the location where it is produced. The assumption that the quality of the agent's output is non-verifiable prevents the use of contracts contingent on this.

The assumption that the agent's history is observable only at the location where she has worked has two implications. First, when one agent has control over a location and can exclude the other agent from its use, she can prevent the other agent from simply moving to another location with no ill consequences, since in moving she would also lose any positive reputation built up through previous high quality output. Thus a junior, in order to keep her reputation, must buy the location (the firm) from the senior. Note, however, that since there are infinitely many such locations, this control is valuable only after an agent has worked there and built up a reputation, but not before. Second, an agent who has failed and consequently has a worse reputation than a new-born agent can move to another location. Then an agent who failed in the first period can pretend

---

[20]The assumptions that agents are risk neutral and value the present and future equally are made for ease of exposition and are not important for the qualitative results.

[21]In Appendix C.1, we consider the case where the junior knows her own type from birth.

[22]This assumption can be relaxed, similar qualitative results would hold so long as the price for output is an increasing function in this belief.

to be a new-born agent. This implies that an agent cannot be severely punished for past actions.[23]

Contracting As mentioned above, outcomes are observable at locations, so that customers' beliefs can change over time, but these outcomes are not verifiable so there is no outcome-contingent contracting. Similarly, a senior cannot write an outcome contingent contract with a junior. Employment contracts will be of the form (w, P) where w specifies the wage that the junior is paid and P specifies the price at which the junior will be able to buy the firm (or equivalently control of the location) at the end of the period.[24]

An agent thus has a rather complicated strategy. Specifically, her strategy consists of:

- choice of mode of work in period 1 of career: work alone, work as an employee (if an appropriate position is offered), hire an employee;

- choice of mode of work in period 2 of career: move location or stay, hire an employee (in which case the strategy will include a decision of the wage contract offered), buy the firm (if previously an employee), work as an employee (if an appropriate position is offered);

- effort decision (in both periods of life).

An agent's strategy will of course depend on outcomes and decisions in previous periods.

---

[23]The case where an agent's age is observable appears in Appendix C.2. In that framework and in an equilibrium characterized, customers would hold the belief that only agents who had failed would move and so observing a second period agent in a new location effectively reveals her history. In that case too, an equilibrium in which agents exert effort in the first period of life and successful agents exert effort in the second period of life can be sustained.

[24]Similar results can be obtained if instead the senior offers an employment contract of the form (w_1, w_2) where w_2 denotes a second period wage for the junior if retained and the senior keeps all revenues generated in that period. In this case, a constraint on equilibrium strategies would be that the senior only promotes a junior who succeeded, thus providing incentives for the junior to work, in the spirit of the up-or-out models of Kahn and Huberman (1988) and Prendergast (1993). In this case, for the promoted junior to have incentives to exert effort, she must gain from the future reputation of her own junior—this can naturally be interpreted as promotion to partnership. With the (w, P) contract the senior always would want to sell the firm, but it is worth more to a successful junior than an unsuccessful one, providing incentives for the junior to work. Since in equilibrium only a successful junior would find it worthwhile to buy the firm, the senior has incentives in the first period of the junior's life to ensure that the junior has the opportunity to succeed. This is similar to the motivation for a partner to mentor in Morrison and Wilhelm (2003).

In each period, customers Bertrand compete for the goods that are produced before observing quality. The quality is observed between periods and customers revise beliefs. An agent's strategy (such as whether to hire, buy the firm, exert effort, relocate) for period 2 of course depends on the observable outcomes in previous periods.

Note that in the case where agents can only work alone and there is no hiring, the unique perfect Bayesian equilibrium outcome is that no agent exerts effort. With no reputational incentives in the final period of her career whatever the belief about the agent at that time and no mechanism for contingent payments, there are no incentives for effort in the final period of the career. It follows by backward induction there are no such incentives in the first period either. Moreover, even when allowing for joint production, there is always an equilibrium where no agent exerts effort. Suppose that the public believe that no agent ever exerts effort and would continue to believe this even after observing an agent producing high quality, then this belief would be upheld in equilibrium since no agent will exert effort at a cost c when whether or not they do so, they sell their service at a price 0.

3.4.2. An equilibrium with effort

Despite the result that an equilibrium always exists where no agents exert effort, allowing for joint production can allow other equilibria which do induce effort.

Proposition 11. *If* $1 - \lambda - c + \frac{c\lambda}{2} - \frac{3c(2-\lambda)}{2\lambda} \geq 0$ *then there is a Perfect Bayesian Equilibrium in which all agents exert effort in the first period of their careers, and competent agents exert effort in the second period.*

Proof. We prove this result by construction, outlining equilibrium strategies and verifying that these strategies do indeed characterize an equilibrium.

Specifically, the strategies are as follows:

In the first period of life, either an agent founds her own firm, working on her own in an unoccupied location or accepts a position as an employee if offered one at sufficiently attractive terms; in either case she exerts effort in the first period of her own life. If she founds her own firm and fails in the first period, then she poses as a new-born agent. If she succeeds in the first period of life after founding her own firm, then in the second period she hires a junior (offering a wage contract that pays w and offers the firm at P) and they work together with the founder exerting effort.

Alternatively an agent might begin life as an employee in an established firm with a (w, P) contract, so long as such a contract offers her in equilibrium at least as much lifetime earnings as founding her own firm. She exerts effort in the first period of life. If the firm as a whole produced one or no high quality outputs, she chooses not to buy the firm but instead poses as a new born agent. If the firm produced two high quality outputs then she buys the firm at the specified price P, hires her own junior offering her junior the contract (w, P) and she works together with her junior and exerts effort.

A second period agent posing as a newborn either works alone or works as an employee if offered a position and if w is at least as great as the wage she could earn when working alone.

In the equilibrium described below, all agents exert effort in the first period of life and so all competents succeed in the first period of life. Thus the population of those who appear to be new borns consists of a measure 1 of true new borns and a measure $1 - \lambda$ of second period inepts posing as new born, thus the probability that an agent posing as a new born is competent is $\mu = \frac{\lambda}{2-\lambda}$.[25]

It is common knowledge among the public and all agents that these are the equilibrium strategies and prices are set appropriately, with the industry capturing the full consumer surplus. Thus the price of the output of a new-born agent (or of an agent pretending to be new born) working alone is μ since there is a probability $\frac{1}{2-\lambda}$ that she truly is newborn, in this case she exerts effort in equilibrium and so generates high quality output with probability λ (the probability that the new born is competent). When an agent has succeeded then it is known that she is competent and will produce high quality products with certainty in the following period when exerting effort. The on-equilibrium path behaviour and strategies are summarized in Figure 3.3.

To complete the characterization of the equilibrium, one must also consider off-equilibrium beliefs:

---

[25] In the case where a competent agent exerting effort succeeds with probability $g < 1$, and an inept or competent exerting no effort with probability $b > 0$, things are a little different. An issue which arises in this case, is that the probability of someone who appears to be new-born truly is new-born will depend on the fraction of second-period agents who had previously been working alone rather than as employees. Solving for a steady state, where the proportion of those starting their working lives as employees closes the model, which appears in Appendix C.4.
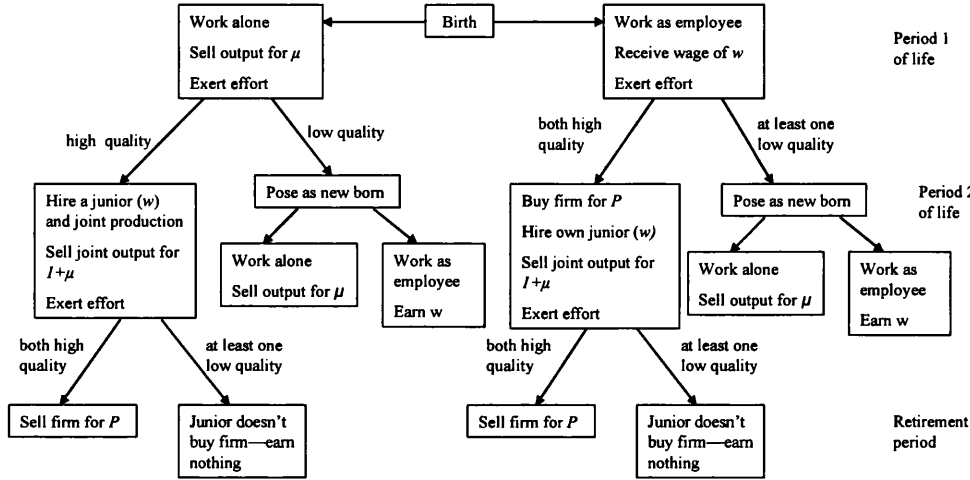
Figure 3.3. Equilibrium behaviour

- if an agent does not offer the contract (w, P) then any potential juniors who will not have seen her history assume that she has previously failed;[26] and finally,

- when an apparently new born agent hires, in this off-equilibrium action, customers suppose that the agent must be a second-period inept agent.

The value of a new born who begins her career by founding her own firm is given by:

$$V_f = \mu - c + \lambda(1 + \mu - c - w + \mu P) + (1 - \lambda)(\mu w + (1 - \mu)\mu). \qquad (3.6)$$

This expression is built up as follows. On observing an agent who appears to be new-born working alone, customers believe that she is a competent new born agent with probability μ. The new born agent, who exerts effort and knows that she is new born expects success with probability λ and in this case she is revealed as competent, she hires another agent as a junior whose product is expected to be worth μ and so can charge $1 + \mu$ for the joint service (hiring the junior costs her w, further she exerts effort

---

[26]Note that although this does not imply that an agent will reject any offer (in particular, she would accept any offer with very high w and/or low P); however she would reject any offer that an inept senior could make to her and which would allow such a senior to make a non-negative profit by hiring her. Further note that as discussed below, by offering (w, P) a competent senior in effect offers the junior a contract which compensates her for her outside option (working alone); the senior can therefore offer no less in expected terms.

at a cost c) and in case of two successes, which occurs with probability μ, she sells the firm earning P. Following a failure in the first period, the agent can pose as a new-born agent and so receive the revenue w if hired, which occurs with probability μ since in equilibrium there is a measure λ of hiring agents and $1 + 1 - \lambda$ is the measure of agents claiming to be new born; if not hired then the agent works alone and earns μ.

Similarly, the value for a new born agent who works as an employee is given by the following expression:

$$V_e = w - c + \lambda(1 + \mu - c - w - P + \mu P) + (1 - \lambda)(\mu w + (1 - \mu)\mu). \qquad (3.7)$$

In equilibrium, a new born agent is willing to become a junior or equivalently $V_e \geq V_f$. In particular, this implies that:

$$w - \mu - \lambda P \geq 0. \qquad (3.8)$$

In addition, a second period agent posing as a new born should be willing to be hired as a junior.[27] Specifically, this condition is given by:

$$w \geq \mu \qquad (3.9)$$

Given that there is a scarcity of junior slots available, hiring seniors will drive down the value that the contracts deliver to employees to the point where they are indifferent. In particular this implies:

$$w = \mu + \lambda P.^{28} \qquad (3.10)$$

Conditions (3.9) and (3.10) can be thought of as individual rationality conditions. The remaining deviations can be categorized into a number of separate groups (i) exert effort (incentive compatibility) (ii) hiring policy and (iii) buying the firm. Each group is considered in turn.

(i) exert effort (incentive compatibility)

For a second period agent who had been successful there is an incentive compatibility constraint—which is identical in both the case that she worked alone in the first period and the case when she buys the firm from her employer—specifically, this constraint is as follows:

$$\mu P \geq c \qquad (3.11)$$

---

[27]It can easily be shown that no equilibrium exists in which only new-borns are willing to work as employees and only the population of those working alone consist of true new-borns and posers.

In the first period in both cases, that the agent works alone or as a junior, it must be worthwhile to exert effort. The corresponding incentive compatibility conditions are the following:

$$1 + \mu - c - w + \mu P - (\mu w + (1 - \mu)\mu) \geq \frac{c}{\lambda} \qquad (3.12)$$

and

$$1 + \mu - c - w - P + \mu P - (\mu w + (1 - \mu)\mu) \geq \frac{c}{\lambda} \qquad (3.13)$$

Note that by (3.11), it follows that (3.13) implies (3.12).

(ii) hiring policy

Without hiring in the second period of life, the agent could not commit to effort and so the best that she could do is pose as a new-born and earn $\mu w + (1 - \mu)\mu$ and so the conditions that a second period agent who had success when working alone in the first period does indeed prefer to hire a junior is given by:

$$1 + \mu - c - w + \mu P \geq \mu w + (1 - \mu)\mu \qquad (3.14)$$

which is implied by (3.12). The same condition ensures that a second period agent who had been an employee and bought the firm, prefers to hire.

Suppose that an agent failed in the first period of life, then she must be inept and a single success would prove that her junior was competent and so a sufficient condition that would ensure that she does not hire is given by:

$$\mu w + (1 - \mu)\mu \geq \mu - w + \mu P. \qquad (3.15)$$

Finally, since customers hold the off equilibrium belief that any apparent new born attempting to hire must be a failed second period agent, the condition that ensures no second period agent posing as a new born hires is also given by (3.15). A true new born agent would not hire when $V_f = V_e \geq \mu - w + \mu P + \mu w + (1 - \mu)\mu$ (note in the first period of life, customers would hold the belief that she must be a failed second period agent). Note that $V_f \geq w + \mu w + (1 - \mu)\mu$ by (3.13) and so (3.10) and $\mu = \frac{\lambda}{2-\lambda} < \lambda$ ensures that a true new born would not hire.

(iii) Buying the firm

An employee who had succeeded would indeed buy the firm so long as this generates more value than her alternative—posing as a new born. This is the case when:

$$1 + \mu - c - w - P + \mu P \geq \mu w + (1 - \mu)\mu \qquad (3.16)$$

This is implied by (3.13).

An employee who had failed is revealed as inept, and so rather than spending $P > 0$ to remain in this location, she would rather costlessly move to another location and pose as a new born agent where she would have a higher reputation.

Thus sufficient conditions are (3.10), (3.11), (3.13), and (3.15). Substituting for $\mu = \frac{\lambda}{2-\lambda}$ and for w from (3.10), these conditions can be reduced to $1 - \lambda - c - \frac{(2-\lambda)c}{2\lambda} + \frac{c\lambda}{2} \geq P \geq \frac{c(2-\lambda)}{\lambda}$, or equivalently:

$$1 - \lambda - c + \frac{c\lambda}{2} - \frac{3c(2 - \lambda)}{2\lambda} \geq 0 \tag{3.17}$$

This concludes the proof. □

The characterized equilibrium relies on the senior's ability to commit in advance to the price at which she will sell the firm.[29] The junior either buys the firm (and goes on to her own junior) or else leaves. Since the senior can commit to a price for the firm, the junior's incentive problem is resolved since there are rewards to working hard, being revealed as competent and buying the firm. The senior's incentive problem is resolved since her action affects the junior's reputation, and only a junior with a good reputation would be willing to buy the firm.

A junior with accumulated reputation buys the firm since locations are assumed to be informationally separate and the senior controls access to the existing location, so that a junior who leaves would have to forego her accumulated reputation.

Note that the result that all competent agents exert effort in the second period of their careers relies on the simplifying assumptions that a competent agent who exerts effort produces high quality for sure and that an inept produces low quality for sure. Thus in the equilibrium described, a competent is always recognized as such. Relaxing these assumptions, would suggest that mistakes are possible (that is a competent agent may not be recognized as such by the end of the first period), but a qualitatively similar result holds whereby under parameter restrictions all agents exert effort in the first period of their careers and successful agents in the second.

Further note, that as in the illustrative example of Section 3.3, uncertainty plays an important role. The equilibrium condition (3.17) fails for λ close to 0 or 1, as illustrated in Figure 3.4, which plots c against λ where the condition is satisfied below the line.
$$1 - \lambda - c + \frac{c\lambda}{2} - \frac{3c(2-\lambda)}{2\lambda} = 0$$

---

[29]In general, many ex-post bargaining schemes which share the benefits of the junior's accumulated reputation between the senior and junior would lead to qualitatively similar results.
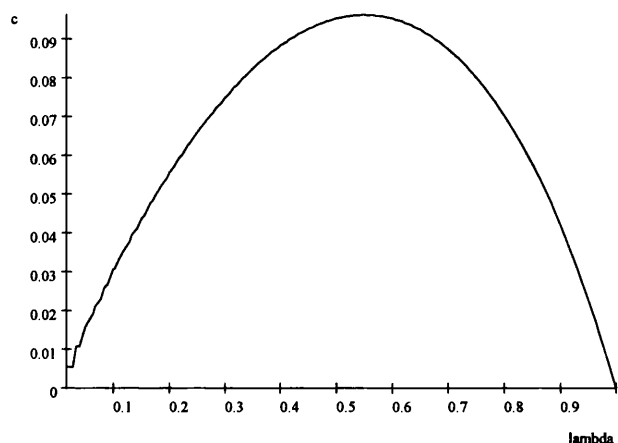
Figure 3.4. Equilibrium condition: c against lambda

3.5. Discussion

The model builds on a number of assumptions which are worthy of further discussion.

First, we assumed that there was no static advantage or disadvantage in having agents work in teams rather than as individuals. A team's production was assumed to be simply given by the sum of individual members' contributions. Although, it is perhaps more reasonable to suppose that there may be important complementarities in team production, perhaps that a team's output is determined by its weakest member, the model deliberately ignored such considerations. First, it can be shown that similar qualitative results (that seniors hire juniors to create a reputational concern to overcome a commitment problem) can be obtained with other joint production functions. More importantly, focusing on the additive case is intended to demonstrate that the effects identified in this chapter are informational and not the result of a joint production technology which in itself is more or less efficient than solo production.

This assumption that agents are either competent or inept, rather than competent or a Stackelberg type, means that over time uncertainty is resolved and an agent's type is learned. In the stark models considered here, this implies that an agent would have no reputational incentives.[30] Hiring juniors as a mechanism to introduce uncertainty

---

[30]In Holmstrom (1999) a "good" agent performs better even when exerting no effort and so a young agent may have some incentives to work but such incentives are diminished or disappear as the principal

and allow seniors to overcome this lack of incentives relies on a couple of assumptions. Specifically, that the senior's choice of actions can affect customers' information about the junior and that the senior cares about how customers will perceive the junior or equivalently that the senior is the residual claimant on at least some fraction of her co-worker's reputation in the next period. The former is implied by the much-noted and intuitive property of team production that it is difficult to attribute the specific contributions of individuals within teams (see particularly Holmstrom (1982)). The latter is addressed in this chapter by an assumption that a junior cannot leave the employment of the senior with her reputation intact.

There are a number of reasons which can explain why the senior might be able to profit from a co-worker's future reputation. In the central model, we assumed that the senior controls access to the location where they work (for example in the context of professional services it seems natural to think of this as access to clients or a non-compete clause in the junior's contract) and that agents build reputation only at the location where they work. Then if the co-worker leaves, she cannot take her reputation with her. This seems a reasonable assumption for law-firm associates for example and given the use of non-compete clauses is likely to apply to junior partners. Thus, although there are many locations where an agent might choose to work at the beginning of her career which are *ex-ante* identical, control of the location acts as a valuable asset *ex-post*.

Levin and Tadelis (2002) cite evidence on the prevalence of non-compete clauses in the US. In the UK Turnor (2001), for example, states that most professional service firms employ restrictive covenants binding outgoing members and describes that:

> "Typical restrictive covenants include, for example:
>
> (a) a clause preventing the outgoing partner from acting for those who have been clients of the firm, or of the particular partner in a specified period (say two years) before his retirement, such restriction applying for a period of two, three or even five years after retirement;
>
> (b) a clause preventing the outgoing partner from practising at all for a specified period in a specified geographical area;
>
> (c) clauses preventing any activity within a specified period which involve the provision of services in a way which competes with the business of the firm"

---

or customers learn the agent's type. We abstract from this effect by assuming that effort and ability are complimentary.

However, the effectiveness of such restrictive covenants is imperfect, as discussed for example in Rebitzer and Taylor (2001). The qualitative results of this chapter would apply even if such covenants were imperfect as long as they had some sufficient effect.

In this chapter, the senior can act as the residual claimant on all of the junior's reputational gains and so the junior might have little incentive to exert effort for the sake of her reputation unless the senior can reward her for such effort. The senior cannot contract with a junior directly on the basis of the output produced as this is assumed to be non-verifiable (or else there would be no need for reputational incentives as explicit contractual incentives would suffice). To overcome this problem, we suppose that the senior can commit to a price at which she would sell the firm to the junior. The price is set in equilibrium at a level which would allow some returns to a successful junior providing the junior with incentives and at which only a successful junior would be willing to buy, ensuring that the senior has an incentive to exert effort and so give the junior an opportunity of succeeding. An equivalent formulation is to suppose that contracts can be made contingent on the task to which the junior is assigned in the second period—that is a contract could allow for a different second period wage (a severance payoff) for a junior who is fired to one who is retained. As discussed by Kahn and Huberman (1988) and Prendergast (1993), promotion or task assignment can be effective in providing incentives when task assignment is at the discretion of the employer only when the expected marginal productivity of the employee differs in the different jobs.[31] In particular when this property holds, then task assignment can act as a quasi-contingent contract—an aspect which is important for the equilibrium characterized.[32]

3.6. Conclusions

This chapter highlights the pivotal role of uncertainty for robust reputational incentives and that an agent can strategically introduce such uncertainty without damaging her existing reputation. We suggest that hiring and working with a junior with uncertain reputation as a strategic choice can endogenously introduce the uncertainty required for

---

[31]This would be the case here—a successful agent is worth more to the senior in the promoted position and by committing to the wage in this position, the senior can commit to reward the junior for success.

[32]Alternative treatments of up-or-out promotion rules include Harris and Weiss (1984), which generate up-or-out rules based on learning with finite lifetime and risk aversion, Carmichael (1988) which introduces a model of tenure built on the insight that current workers are best informed to select new workers, and O'Flaherty and Siow (1992) and (1995), which are based on the scarcity of junior slots used to assess suitability for a senior slot.

reputational incentives and in this way allow an agent to commit to exerting efficient though costly effort. In particular, it is perhaps worth re-iterating that whereas typically teamwork is thought to reduce individuals' incentives and lead to a free-rider problem, here if only solo production were possible (and even if this were more efficient in a complete contracting world) there would be no incentives for effort. The crucial aspects in establishing this result are that a senior and junior can work together in a non-attributable production process and that the senior can gain from the future reputation of the junior. Incentives for the junior are provided by committing to a price at which the junior can buy the firm or equivalently to an up-or-out promotion scheme in which for the promoted agent to have incentives to exert effort, promotion must be to a position where the promoted employee lays claim to future revenue generated by her own junior. This has a natural interpretation as promotion to partnership.

In the context of professional services, the inspiration for this chapter, the conclusion that agents are motivated by a concern for their own reputations when young, and when old (and successful) for the reputation of the firm that they own does not seem unreasonable. An important aspect in this chapter is that firms or teams are small—often a reasonable assumption in such industries—and so other models which have considered the interaction of individual and collective reputations and assumed that firms either consist of a single individual or very many individuals, are unable to consider such as a result.[33] Other instances where agents may be choosing to endogenously introduce uncertainty or choose to allow the reputations of a small number of agents or good to interact might include a firm's turnover policy (hiring and firing at a more aggregate level) and some instances of product bundling.

This chapter has suggested direct ownership (in Section 3.3) or ex-post control of location, perhaps through restrictive covenants (in Section 3.4), as a means by which a senior can gain from the reputation of the machine or a junior.[34] Another alternative might be to suppose that the agent is exploiting uncertainty about another dimension of her own capability; as a specific example suppose that there is certainty that the agent is competent but uncertainty about her ability to select good candidates, then without recourse to directly gaining a benefit from the revenues that the good reputation of a

---

[33]Note that even within large law firms, individual departments or groups which are typically small have their own reputations.

[34]Another possibility is explicit bonding—a junior buys a share in a partnership which is returned only on retirement.

junior can generate, the senior would have a reputational concern but over a different dimension of her ability.

Other aspects of the model which might be developed or extended are the restrictions to a binary effort choice and two types of agent. Relaxing these assumptions may yield interesting results. In particular it is not hard to imagine that in a more sophisticated model seniors would differ according to the history of their previous employer (and their employer's previous employer and so on). This would suggest that a firm's age is important and further, that working for a different type of senior (or a firm of different age) would entail a different employment contract. Another aspect which might be profitably relaxed is the form of the joint production function, the choice of an additive production function in which joint production is the sum of individual contributions is useful in highlighting the purely informational role in obscuring those individual contributions, but introducing some complementarity in joint production would perhaps be more realistic and could enrich the model.

Static efficiency and reputation incentives (with Juanjo Ganuza)

4.1. Introduction

In numerous examples discussed at some length in the chapters above reputational concerns can have a significant effect on incentives and through them on outcomes. In such cases, therefore, firm owners, managers, regulators and other planners have an interest in manipulating reputational concerns through their choice of policies or through the organizational design. We highlight in this chapter that the kind of reputational concern that an agent has might have an impact on a planner's preferred policy.

We begin by noting that most of the economic literature on reputation has focused on a reputation for excellence (trying to show that you're a type who always does well, or where reputation is about "who you'd like to be").[1] More recent literature and common intuition suggests that often, reputational concerns might also relate to avoiding a reputation for ineptitude (trying to show that you're not a type who always does badly or where reputation is about "who you're not"). This distinction has been forcibly made recently by Mailath and Samuelson (1998) who highlight, in particular, that the latter view of reputation leads to increasing certainty about the agent's type over time and so reputational incentives disappear over time unless type uncertainty is continually introduced.[2]

---

[1]Following Kreps and Wilson (1982) and Milgrom and Roberts (1982) and later Fudenberg and Levine (1989); the formal economic literature on reputation has been used primarily to discuss beliefs about the type of the agent. Previous literature and a great deal of intuition has also used the term in a somewhat looser fashion to consider sustaining certain actions in infinitely repeated games. As highlighted in Fudenberg and Levine (1989) this corresponds closely to the notion of reputation where reputation is a concern to show that you're a "Stackelberg" type—that is a type whose behavior a strategic agent would like to promise to commit to—similar to what we term later in this chapter a reputation for excellence.

[2]Further in Mailath and Samuelson (1998), the model is constructed in such a way that there is unravelling so that if there are no reputational incentives at some point, there are no such incentives throughout.

In this chapter, we highlight an important distinction between the two approaches in a simple two-period model and focus on short-term reputation effects. Specifically, we consider making the strategic agent more efficient in the sense of increasing the range of tasks that she can undertake without effort or decreasing the range of tasks in which she is incapable.[3] We show that increasing the agent's efficiency in this way diminishes reputational concerns (reducing effort) when reputation is about excellence but increases reputational concerns when reputation is about ineptitude.

The intuition for this result is clear. When reputation is about trying to demonstrate excellence as opposed to competence, then the better a competent agent is the less valuable it is to try and convince customers that she is excellent and the less convincing is a success in trying to demonstrate excellence rather than competence. In contrast, when reputation is about trying to show competence rather than ineptitude, then the more efficient a competent agent is, the more valuable it is to show that she is not inept and a success is a more convincing proof of competence.

Thus this chapter highlights that different plausible ways of thinking about reputation can lead to very different, indeed contradictory, conclusions. We consider a very stylized model to clearly highlight this point; however, we believe that this result has wider application. In particular, we briefly relate these intuitions to the questions of training, certification and the optimal size of an organization, where if reputation is for excellence then there might be a trade-off between the statically optimal size (which might be large) and dynamic considerations. We then discuss the results in this chapter in the context of the wider literature on reputation.

4.2. Model

There are three possible types of agent. An inept type fails in every task that they undertake whether exerting effort or not. An excellent type succeeds whether exerting

---

The more general point on reputational incentives disappearing over time without some kind of replenishment of type uncertainty applies more widely. Indeed, Cripps, Mailath and Samuelson (2004) show this to be the case unless actions are perfectly observable, even in the case when reputation is about excellence and a competent agent can perfectly mimic an excellent agent (though incentives may disappear only in the very long run).

Chapter 3 suggests an endogenous mechanism to maintain type uncertainty by allowing agents to choose to work in teams.

[3]One could think about such an improvement coming through better screening on the part of a certifcation authority or such an improvement coming over time as an agent learns by doing, for example. Another application that we consider below is teams—larger teams with agents of complimentary ability ought to be capable in a wider range of tasks.

effort or not with probability 1. Finally a competent agent, the focus of our analysis, has different abilities in different tasks. She has facility in some tasks, in which she succeeds with probability 1 without exerting effort. In other tasks, in which she has no facility but is capable, she exerts effort $e \in [0,1]$ at a cost $\frac{e^2}{2}$ and she succeeds with probability e.[4] There are also some tasks in which she is incapable—she fails in such tasks whether exerting effort or not. Specifically, suppose that a competent agent succeeds with probability 1 in a fraction $\alpha > 0$ of tasks at no cost, but on a fraction $(1 - \alpha - \beta) > 0$ of tasks she must exert effort in order to succeed, finally on a fraction β of tasks she fails for sure. We further suppose that agents know their own types and that competent agents learn the type of the task (that is whether it is the kind of task for which they need to exert effort) only after they have accepted to take it on.[5] Customers never learn the type of the task and learn about the type of the agent only by observing her successes and failures.

Note that from an *ex-ante* perspective (that is before the task is realized), regardless of the competent's effort decision, it is always the case that the excellent agent is more likely to succeed than a competent agent, who in turn is more likely to succeed than an inept agent.

Customers are risk neutral, value a success at 1 and a failure at 0 and they Bertrand compete for the agent's service in each period.[6] Moreover, outcomes are observable but effort is not observable and contracts are incomplete, so that in effect an agent is paid in advance at a price which is simply the customers' common belief that the agent will produce a success. It is common knowledge that customers start out believing that there is a probability $p \geq 0$ that any agent is excellent, $q \geq 0$ that she is inept and $r > 0$ that she is competent, where $p + q + r = 1$.

---

[4]This specific form for the cost function is assumed solely for presentational convenience. Qualitatively similar results apply for a more general cost function $c(e)$, where $c' > 0$, $c'' > 0$, $c''' \leq 0$, $c'(0) = 0$ and $c'(1) > 1$.

[5]Alternatively, we could have set the model up by assuming that there is only one kind of task, in which the competent agent succeeds with probability $\alpha + \frac{(1-\alpha-\beta)e}{1-\alpha-\beta}$ where $e \in [0,1]$. The two models are similar, the approach adopted in the main text reflects our interest in reputation for teams and perhaps makes the discussion in Section 4.5.2 a little clearer.

[6]The assumption that customers Bertrand compete for the product is not crucial, similar results would hold so long as the price paid was increasing in the customers' expected likelihood that the agent will be successful.

There are two periods of trade, and outcomes are observed (and beliefs revised) in between the two periods.[7] Specifically timing is as follows:

(1) Period 1

    (a) customers Bertrand compete for the agent's service and the winning customer assigns his task to the agent

    (b) the agent learns the type of the task[8]

    (c) the agent decides the level of effort if appropriate (that is if she is competent and it is a task in which she is capable)

    (d) success/failure commonly observed

    (e) customers update beliefs according to Bayes rule

(2) Period 2

    (a) customers Bertrand compete and winner assigns task to agent

    (b) success/failure observed

Notice that in period 2, we could allow the agents the opportunity to exert effort but no agent would do so. Further, note, that neither the excellent nor the inept type has any strategic decisions to make—the focus of the analysis and of reputational concern is on the competent agent.

We suppose that agents weigh the two periods equally and that agents maximize the sum of profits for the two periods and we solve for the Perfect Bayesian equilibrium.[9]

4.3. Equilibrium analysis

In the second period, there are two possible prices S and F, corresponding respectively to the price a customer pays after observing a success in the first period and after observing failure in the first period. Such prices will of course depend on equilibrium beliefs about the type of the agent and of equilibrium effort. Second period beliefs about the type of the agent are simply updated according to Bayes' rule, taking into account the equilibrium effort decision of a competent agent for a task in which she's capable.

---

[7]One need not take the two periods of the model literally, rather the second period can be thought of as a reduced form payoff for a given reputation level.

[8]That the agent learns the type of the task only after accepting it is not a crucial assumption. First note that no such assumption is necessary in the version of the model suggested in Footnote 5. Moreover in rejecting a task, an agent might reveal herself not to be excellent.

[9]Allowing for discounting between periods or indeed allowing profits in the second period to be more valuable than in the first (consistent with an interpretation of the second period as a reduced form for the future) does not affect the qualitative results.

Both successes and failures could be observed on the equilibrium path, so that there is no possibility of an off-equilibrium observation.

We focus on a competent agent when she faces a task in which she has no facility but can improve her probability of success by exerting some effort—this being the only case in which an agent has a decision to make. In those circumstances she chooses a level of effort to solve the following problem:

$$e^* \in \underset{e}{\operatorname{argmax}} \, eS + (1 - e)F - \frac{e^2}{2}. \tag{4.1}$$

Thus, the optimal level of effort e^* is the implicit solution to the equation $e^* = S - F$

Now in equilibrium and using Bayes rule we can compute, S and F as below:

$$\begin{aligned} S \;&= \frac{p}{p + r(e^*(1-\alpha-\beta)+\alpha)} + \alpha \frac{r(e^*(1-\alpha-\beta)+\alpha)}{p + r(e^*(1-\alpha-\beta)+\alpha)} \\ &= \alpha + \frac{p(1-\alpha)}{p + r(e^*(1-\alpha-\beta)+\alpha)} \end{aligned}. \tag{4.2}$$

The expression above is built up as follows. If customers believe that success can arise either from an excellent agent, or from a competent agent with facility in the task or a competent agent with capacity in the task but exerting effort e^*, the probability of a first period success is $p + r\alpha + re^*(1 - \alpha - \beta)$. Therefore, given a success, the probability that the agent is excellent, according to Bayes rule, is $\frac{p}{p+r(e^*(1-\alpha-\beta)+\alpha)}$ and the probability that the agent is competent is $\frac{r(e^*(1-\alpha-\beta)+\alpha)}{p+r(e^*(1-\alpha-\beta)+\alpha)}$. An excellent agent is expected to generate a value of 1 in the second period and a competent agent, who would exert no effort in period 2, is expected to generate α, leading to the value of S as above.

Again, following Bayes rule and noting that an inept agent generates no value and that an excellent agent always succeeds so that no success in the first period suggests that the agent must be either competent or inept, we can derive an expression for F explicitly as below:

$$\begin{aligned} F \;&= \alpha \frac{r((1-e^*)(1-\alpha-\beta)+\beta)}{q + r((1-e^*)(1-\alpha-\beta)+\beta)} \\ &= \alpha - \frac{q\alpha}{q + r((1-e^*)(1-\alpha-\beta)+\beta)} \end{aligned}. \tag{4.3}$$

From the first order condition of the maximization problem (4.1) and substituting the expressions for S and F above, we obtain:

$$e^* = \frac{p(1 - \alpha)}{p + r(e^*(1 - \alpha - \beta) + \alpha)} + \frac{q\alpha}{q + r((1 - e^*)(1 - \alpha - \beta) + \beta)}. \tag{4.4}$$

We prove in Appendix D that Equation (4.4) has a unique solution which is feasible (that is $e^* \in [0,1]$). Given this result, a number of properties and comparative statics exercises can be explored. Note, in particular, that e^* is lower than the efficient solution $e^{fb} = 1$.[10]

A comparative static exercise in which we are particularly interested (and which we motivate in discussion of applications below) is an increase in α or decrease in β—that is increasing the static expected efficiency of a competent agent.[11] In particular, it can be shown that:[12]

$$sign(\tfrac{de^*}{d\alpha}) = sign(-p\tfrac{p+r(1-\beta e^*)}{(p+r(e^*(1-\alpha-\beta)+\alpha))^2} + q\tfrac{q+r(1-e^*)(1-\beta)+r\beta}{(q+r((1-e^*)(1-\alpha-\beta)+\beta))^2}), \text{ and} \qquad (4.5)$$

$$sign(\tfrac{de^*}{d\beta}) = sign(p\tfrac{r(1-\alpha)e^*}{(p+r(e^*(1-\alpha-\beta)+\alpha))^2} - q\tfrac{r\alpha e^*}{(q+r((1-e^*)(1-\alpha-\beta)+\beta))^2}). \qquad (4.6)$$

4.3.1. Reputation for ineptitude

If there is no possibility that the agent is excellent then all reputational concerns for the strategic, competent agent are to show that she is not inept. No possibility of excellence is equivalent to setting $p = 0$. In this case, given the expressions (4.5) and (4.6) we can conclude that $\tfrac{de^*}{d\alpha} > 0$ and $\tfrac{de^*}{d\beta} < 0$.

Thus, when reputation is about showing that she is not inept, increasing the efficiency of the agent (either by increasing the range of tasks in which she has facility or reducing the range of tasks in which she is incapable) ensures that the agent exerts more effort, that is it increases her reputational concern.

Underlying these results are two effects.

---

[10]This follows as the proof shows that $e^* \in (0,1)$ and so in particular $e^* < 1$. A similar result (that $e^* < e^{fb}$) can be proven for the general case with a cost function as described in Footnote 3 and so long as period 2 profits are not weighted more heavily than period 1 profits.

[11]Note that another plausible way to view an improvement in a competent agent's static efficiency is to leave α and β unchanged but to alter the cost of effort $c(e)$. In this latter case, for example reducing cost of effort proportionally so that new cost of effort are given by $\tfrac{e^2}{3}$ leads trivially to greater effort as can readily be seen by considering Equation (4.4). However, whereas improving efficiency by increasing α does not change the efficacy of effort (and the decision to exert effort is taken only after the agent observes that the task requires effort), it is hard to interpret a change in the cost function as one that affects reputational concerns.

[12]We prove the first of these two results in Appendix D, the second can be proven similarly.

First, for a higher α or lower β, for any given level of effort, a failure is relatively more informative about the type of the agent (although with no possibility of an excellent type a success always guarantees that the agent must be competent). Formally, this can be observed as the probability that the agent is believed to be inept, given a failure and for a fixed level of effort e is given by $\frac{q}{q+r(e(1-\alpha-\beta)+\beta)}$ which is strictly increasing in α and strictly decreasing in β. Second, for a higher α, it is more valuable to be thought of as competent rather than inept in the second period. Somewhat more formally, for any given belief μ that the agent is competent rather than inept, the second period profit $\mu\alpha$ is strictly increasing in α (and independent of β).

Thus for a higher α, for any given effort level, failure will be worse news, and its more valuable to be thought of as competent. Both effects push towards greater effort, that is $\frac{de^*}{d\alpha} > 0$ as proven above. For lower β it is only the first effect that applies but has the proven effect that $\frac{de^*}{d\beta} < 0$.

4.3.2. Reputation for excellence

In the case where there is no possibility that the agent is inept then $q = 0$ and all reputational concern is about trying to convince customers that the agent is excellent rather than competent. In this case, given the expressions (4.5) and (4.6), we obtain $\frac{de^*}{d\alpha} < 0$ and $\frac{de^*}{d\beta} > 0$ and so when reputation is about showing that she is excellent, then increasing the efficiency of the agent (either by increasing the range of tasks in which she has facility or reducing the range of tasks in which she is incapable) ensures that the agent exerts less effort, that is it decreases her reputational concern.

The intuitions underlying these results are analogous to those considered for the case where reputation is purely a concern about appearing inept. In the case where reputation is about trying to show excellence rather than ineptitude then, for any given level of effort, when α is higher or β is lower then a success is less informative about the type of the agent. Further for any given belief ν that the agent is competent rather than excellent, the additional second profit in successfully convincing customers that the agent is excellent ($= 1 - \nu\alpha$) is strictly decreasing in α (and independent of β).

4.3.3. General case

The results in the two extreme cases, where reputational concerns increase with the competent agent's efficiency (through an increase in α) in the case when reputation is entirely about a reputation for competence rather than ineptitude but decrease when

reputation is about excellence rather than competence, suggest more general results. Note, first, that these are not knife-edge results since for example the expression on the right hand side of Equation (4.5) is continuous in both p and q and e^* is continuous and takes values in $(0,1)$, then it must be the case $\frac{de^*}{d\alpha} > 0$ for p small enough and $\frac{de^*}{d\alpha} < 0$ for q small enough.

Thus even if reputation is not entirely about ineptitude but primarily so then it should still be the case that increasing the competent agent's efficiency should lead to greater reputational concerns. It seems reasonable to suppose that the more reputation is about a concern for ineptitude rather than excellence, in comparison to a prior as competent, the more increasing efficiency should make reputational concerns more pronounced and so lead to higher effort.

Making this conjecture more precise and verifying it is not trivial. The first problem is that we need to specify, for example, how changes in the probability of being excellent changes the probability of being competent and inept, given that these three variables, q, p and r, are linked. In order to do so, we assume that r is constant, and consider $q = 1 - p - r$, so that raising p—loosely increasing concern about excellence—simultaneously lowers q—decreasing concern about a reputation for ineptitude. We make this modeling choice since, given that the only reputational concern of interest is that of a competent agent, it seems reasonable to analyse the case in which we hold constant the prior belief that the agent is competent.

Making this assumption, one might conjecture that $\frac{d^2e^*}{d\alpha dp} > 0$. Such a conjecture would be false;[13] however, the examples that we have been able to find where this conjecture fails have a natural and appealing intuition. In particular, if r is close to 0 but p and q are not, then reputational concern is not so much about a reputation for excellence rather than competence or competence rather than ineptitude, but excellence rather than ineptitude. In the case where r is large enough then customers begin by believing that it is very likely that the agent is competent and so the reputational concerns for a competent agent are to a great extent about showing that she is excellent rather than competent, or competent rather than inept and so these considerations and the effects

---

[13]Specifically, there is a counter-example at $q = 0.98$, $r = 0.01$, $p = 0.01$, $\alpha = 0.01$, $\beta = 0.98$. At these parameter values, solving Equation (4.4) implicitly (and taking the root with a value between 0 and 1) yields $e^* = 0.980\,49$. Taking appropriate derivatives of Equation (4.4) and solving, one can find $\frac{de^*}{d\alpha} = -8.967\,3 \times 10^{-3}$ and
$\frac{de^*}{dp} = 1.867\,4$
and so finally that $\frac{d^2e^*}{d\alpha dp} = 91.466 > 0$.
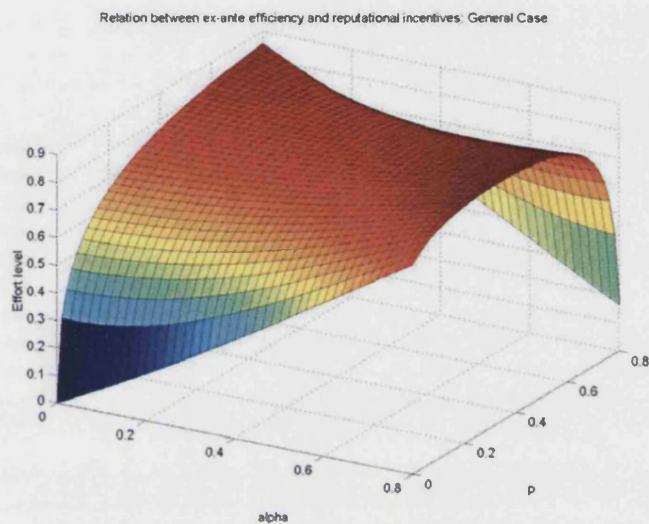
Figure 4.1. Efficiency and reputational incentives

outlined above are the relevant ones, leading to the following result, which is proved in Appendix D.

Proposition 12. *For r large enough $\frac{d^2 e^*}{d\alpha dp} < 0$.*

Although the proof is a limiting result, simulations suggest that the parameter values, such as in the example described in Footnote 13, where this result fails are unusual and that the case where $\frac{d^2 e^*}{d\alpha dp} < 0$ is much more common. For example, Figure 4.1 plots e^* for $\beta = 0.2$ and $r = 0.5$, and as p and α vary over the range of feasible parameters (that is with $p \in [0, 0.8]$ and $\alpha \in [0, 0.8]$). The figure shows that effort is increasing in α for low values of p, but decreasing in α for high p.

Substituting $q = 1 - p - r$ in Equation (4.6), we could obtain similar results for the sign of $\frac{de^*}{d\beta}$ (in this case, $sign(\frac{de^*}{d\beta})$ is negative for low values of p while it is positive for large values of p.)

4.4. Discussion

In the model above, we have interpreted an increase in efficiency to imply a reduction in the range of tasks in which the competent agent is incapable or an increase in the range of tasks which she can undertake costlessly. While both suggest that the agent

is more efficient in expectation (that is before she is assigned a task, there is a greater likelihood that she can be successful in the task), such improvements in efficiency affect the probability of encountering a state in which the agent makes a decision regarding how much effort to exert, but not the cost of exerting effort in those states.

In general, in considering models of reputation, one can perform comparative static exercises by altering parameters which can affect either the costs of maintaining reputation or the gain in maintaining reputation (or both).[14] In the model above, in which changes in efficiency for the static production function are in expectation, the costs of taking the action which can help maintain reputation do not change as we vary efficiency, leading to a relatively simple model, in which such changes in efficiency affect only the value of maintaining reputation and the efficacy of the action in altering beliefs. As we have demonstrated above, the way in which such changes in efficiency affect the value of maintaining reputation depends crucially on the extent to which reputation is about ineptitude rather than excellence.

An alternative way to consider improving static efficiency is to consider changing the costs of exerting effort in tasks in which the agent is capable from $\frac{e^2}{2}$ to $\mu\frac{e^2}{2}$ with $\mu < 1$. Such improvements in efficiency lead to greater effort irrespective of the kind of reputation considered.[15] Further, this is the typical modeling approach to assess an improvement in efficiency and so one might naturally question the relevance of the above results.

In reply, one can make a number of points. First, it is perhaps worth noting that for some applications and in particular where all considerations are static, then increasing efficiency in expectation (as we consider in the model above) or increasing efficiency in absolute terms would be equivalent, and so previous researchers would have had little reason to make this distinction. Most importantly, however, we believe that there are applications in which it may be more reasonable to consider efficiency in expectation. In particular, in services industries agents typically face a wide variety of different tasks and

---

[14]See for example Chapter 5, which considers the effect of changing the degree of competition on reputational incentives. This effects both the cost of maintaining reputation, inasmuch as the market share might be different as the degree of competition changes as well as the price charged with a "good reputation" and also the value of maintaining reputation.

[15]This is trivially the case. Replacing Equation (4.4), the implicit equation for optimal effort would be:
$$\frac{1}{\delta}\mu c'(e) = \frac{p(1-\alpha)}{p+r(e(1-\alpha-\beta)+\alpha)} + \frac{q\alpha}{q+r((1-e)(1-\alpha-\beta)+\beta)}$$
which suggests that the optimal effort in the case where the cost of effort is $\mu c(e)$ is strictly higher than when it is $c(e)$ if and only if $\mu < 1$.

it can be difficult to predict in advance which kind of task they might be facing. In such cases it is far from obvious whether to think about a more efficient agent as one who can perform every task more efficiently or can perform a wider variety of tasks more efficiently. In practice, both aspects are likely to be important, though for simplicity of presentation and to make the point starkly, we highlight the latter aspect. In a related application, on organizations with members who may have complimentary abilities (which we discuss in greater detail below) thinking of efficiency in expectation seems a much more reasonable approach. In this case it seems reasonable to think of a more efficient organization or network as one that is more likely to contain an agent with facility or capacity for any given task and it seems hard to justify an assumption of absolute efficiency.

4.5. Applications

4.5.1. Training

In this subsection we suppose that training has a role in improving expected efficiency of an agent. We suppose more precisely that training will improve the efficiency of a competent agent by increasing α or decreasing β. We further suppose that training has no effect on either inept or excellent agents. Though this might appear a somewhat extreme assumption, it is not unreasonable to suppose that training programs are targeted at average people and might have more effect on them than either the particularly talented or inept.

Note that in addition to the static benefits of training in improving (at least in expected terms) the static productive efficiency of an agent, there are also dynamic effects.[16] In particular, it follows immediately from the results of Section 4.3 that, at least when training has no signalling role, these reputational considerations imply that a competent agent will exert more effort following training when reputational concerns relate primarily to avoiding being seen as inept rather than being about excellence.

In itself this result might have implications. For example, consider a government considering imposing the introduction of a mandatory training program, or an industry association, such as a bar association or similar, requiring additional qualifications for accreditation. In these cases, additional training would have no signalling role and

---

[16]If outcome contingent contracts can be written or there are no moral hazard concern—assumptions which have often been made in discussions of optimal investment in training—then all effects are essentially static. However, in many applications, for example in law, consulting, architecture and other services, these assumptions do not hold and dynamic reputation effects can play an important role.

the decision maker would need to take into account the effects arising from dynamic considerations as outlined above. In particular, the correct policy choices might be substantially different depending on whether reputation was primarily about excellence or ineptitude.[17] Furthermore, even in cases where training does have a signalling role, then the above analysis would still be valid, though the effects might be confounded with signalling effects, and so the overall impact from allowing or mandating training on incentives for effort and welfare might be difficult to assess.[18]

A related idea is a normative question on how much one should invest in screening for certifying doctors, for example, or other professionals. Inasmuch as this affects the initial priors, the analysis above suggests that in finding the optimal level one should take account that reputational incentives would likely be affected.

4.5.2. Organizations

This chapter had its origins in conversations between its coauthors on whether reputational concerns for organizations perfectly able to allocate tasks (or refer them) internally were greater or more muted than for individuals.[19] After each modelled the problem somewhat differently and came to exactly opposite conclusions, it soon became apparent that results depend sensitively on whether one thinks of reputation as concern for demonstrating excellence or avoiding the impression of ineptitude. In particular, in considering a group or team where members share a reputation (for example if individual contributions cannot be separately identified) against an individual, a group inasmuch as

---

[17]This might appear to be somewhat disingenuous inasmuch as implicitly, we suppose that government or the industry is either mandating training or else ensuring that no training occurs. A thorough analysis ought also to consider the alternatives of voluntary training, where signalling roles might arise. In practice, the informational requirements underlying the signalling models and appropriate inference might be more credible under circumstances where the decision maker endorsed a particular training program even if not mandating it. If there is no such provision for endorsement, then the only effective policies might be mandatory training and no training.

[18]As Feltovich et al. (2002) have shown, in an environment with three types, one endogenous observation and one exogenous observation correlated with type, countersignaling can occur. Note that in the case where training is voluntary, there are three types and two observations from which inferences can be drawn. In particular these observations are the decision to undertake training or not, which is a fully endogenous decision and the observation of the first period success or failure, which for the competent type in some states is an endogenous decision, but which will be correlated with type. It seems reasonable to expect that counter-signalling might occur in this model, and the analysis might be involved.

[19]For an insightful paper which deals with a static model where work is not efficiently allocated exogenously, but rather there is a trade-off between efficient referrals and agents desire to realize gains from the cases which they handle see Garicano and Santos (2004).

it can consist of agents with complimentary ability can be more efficient in the expected static sense discussed above. However, following the results of Section 4.3, it is clear that the effects of such improved efficiency and incentives for effort depend on the extent to which reputation is about excellence rather than ineptitude. Thus, for example, when reputation is about excellence, and greater efficiency (larger groups) dulls reputational incentives for effort, one can imagine that this might limit the optimal size of groups; on the contrary, if reputation is primarily about ineptitude, then reputational considerations might lead to larger groups than might otherwise arise.[20]

Clearly, in considering effort and the formation and organization of groups or firms, numerous other factors play a role. There may be free-rider problems associated with production in a team, the size of the team may in itself serve as a signal, the possibility that agents join or leave teams can serve as a signal, the process of group formation, and numerous other factors can play a role. A thorough analysis would require further specific modeling assumption and exploration that could not be captured in a couple of paragraphs, or perhaps even a couple of papers. However, the paragraph above does introduce a consideration somewhat orthogonal to those previously discussed in the literature and highlights that the effect of reputational concerns is likely both to be subtle and to depend importantly on the kind of reputational concern assumed.

4.6. Short and long term reputations

First it is worth restating that this chapter describes a model of short-term reputation. Although one could consider extending the model beyond two periods, over time the type of the agent would be perfectly revealed. In particular, the first time that a failure was observed, customers could rule out the possibility that the agent is excellent and the first time that a success was observed, customers could rule out the possibility that the agent is inept; since (whether exerting effort or not) a competent agent would at some time generate successes and failures, eventually the agent's type would be learned.[21]

---

[20]I thank Rona Bar-Isaac for the observation that amongst professional services firm those with the largest groups are audit firms, where it is easy to believe that the primary reputational concern is to avoid a reputation for ineptitude rather than gain one for excellence.

[21]Even in the less extreme case where excellent agents occasionally failed and inept agents occasionally succeeded, eventually the type would be learned. In this case, as Cripps et al. (2004) show this is true even in the case where a competent agent could perfectly imitate an excellent one, that is in the case where there is imperfect monitoring even if there is a Stackelberg type.

Note, however, that the problem in this paper is somewhat more severe inasmuch as the excellent type considered is not a Stackelberg type. (A Stackelberg type, in the sense of Fudenberg and Levine (1989)

Since almost tautologically reputation incentives rely on some type uncertainty (without that there is nothing to prove as discussed in Chapter 3 for example), with the types described in the model, therefore, there can be no long run reputation effects.

Furthermore, in those papers that have considered something akin to the inept type we consider in this chapter (in particular Diamond (1989), Mailath and Samuelson (1998), and Chapter 3), there are no reputational incentives from trying to avoid a reputation for ineptitude or gain a reputation for competence *per se*. In particular, in Mailath and Samuelson (1998) effort and ability are perfectly complimentary and so proving oneself to be competent has no value unless one can also commit to exerting effort; thus either reputational incentives unravel or else there needs to be some kind of continual replenishment of type uncertainty.[22]

Thus the literature to date has been ill-equipped to consider and has passed over the simple observation we make in this chapter. Note also that perpetual replenishment of type uncertainty in some sense might also be thought of as ensuring that all reputational concerns are short-term (though in addition the analysis in Mailath and Samuelson (1998) and Chapter 3 show that constant replenishment of type uncertainty can lead to reputational incentives that would not arise in a finite horizon model) and so our short-term analysis might have some bite even in the long-run where there is continuous introduction of type uncertainty, for example through name-trading, overlapping-generations of juniors and seniors, obsolescence of skills and so on.

---

in the model would be one who had facility in a proportion α of the tasks, was incapable in a fraction β of tasks and in the remaining tasks succeeded with probability $\hat{e}$ where $c'(\hat{e}) = 1$.)

[22]This is also the case in Chapter 3, though the finite horizon version of the model which appears in the Appendix relies on relaxing this extreme assumption. In Diamond (1989) there is a difference between the strategic agent (BG) and the "bad" agent (B). However, although the bad agent allows for some learning which leads to reputational dynamics, the focus of this beautiful paper, there are no reputational incentives arising from the possibility of this bad type, as Lemma 1 in that paper makes clear there is no advantage in any static sense of being thought of as a BG agent rather than a B agent.

CHAPTER 5

Imperfect competition and commitment

5.1. Introduction

In numerous contexts, such as in professional services industries, it is difficult to write explicit contracts contingent on outcomes. In such cases implicit contracts and reputational considerations can play an important role in ensuring efficient actions. There are numerous exogenous factors which might affect the feasibility of such implicit contracts, or similarly of the possibility that a reputational concern can lead to an efficient action; in particular Baker, Gibbons and Murphy (1994) show that as explicit contracting becomes more effective this can render implicit contracts unfeasible. Furthermore, parties may make strategic decisions in order to affect the viability of an implicit contract.[1] In this chapter, we explore the effect of the degree of competition on a firm's ability to make a credible commitment to exert high effort.

In general whether reputational considerations will motivate a firm to perform some costly action depends on the trade-off between the short-term gain or saving in not performing the action and the long-term effects of beginning the following period with a relatively low reputation. Supposing that taking the costly action maintains a high reputation, and not undertaking it ensures a low reputation (as is typically assumed to be the case in the literature on relational contracts or equilibria supported by trigger strategies) this trade-off can be summarized by the following inequality which ensures that the firm exerts the costly action:

short-term cost of action ≤ discounted value of high reputation - discounted value of
low reputation.

One might expect increased competition to reduce profits, both for a high reputation firm and a low reputation firm. In addition inasmuch as increased competition might

---

[1]For example, Baker, Gibbons and Murphy (2002) show that decisions on the allocation of property rights can affect relational contracts and in a related idea, Bernheim and Whinston (1990) show that multi-market contact can make collusion easier to sustain.

lead to a smaller market share it might also lead to a lower total cost of producing at high quality—the short term cost. Thus the effect of competition on reputational incentives, as summarized by this inequality above seems ambiguous and will depend on the rate at which the degree of competition affects these profits and costs.

However, it seems clear that if competition is sufficiently severe then the discounted value of having a high reputation might be driven close to zero, which would make taking the costly action unattractive.[2] In this case, a firm's temptation to pursue a fly-by-night or hit-and-run strategy whereby the savings from not undertaking the action outweigh the reputational gains would preclude any expectation that a firm would ever take such action and any reputational incentives. This intuition, that reputational incentives can only be maintained if the firm enjoys some premium above costs has been long recognized and has provoked some discussion on how this observation might be squared with free entry into markets—in essence the resolution was some form of loss-making period of costly signalling in order to establish the reputation initially and which then leads to a period of maintaining reputation and enjoying a price premium (see in particular Klein and Leffler (1981) and Shapiro (1983)).[3]

Given this strong intuition that "too much" competition can damage reputational incentives, it is of some interest to determine whether "some" competition can ever help. Again, there is strong intuition to suggest that it might; in particular a firm in competition faces the prospect of a loss in market share as well as a price drop on losing reputation and so the "punishment" for not exerting effort is more severe.

Below, we show that both these intuitions have some force and so overall the degree of competition has ambiguous effects on a firm's ability to commit to high quality. In the final section, we discuss related literature and consider implications of this result, in particular that a firm may choose a more competitive environment, and more competition may lead to higher prices. It should be noted that the analysis in this chapter is partial inasmuch as we largely take the market structure as exogenous without addressing sunk costs, barriers to entry or other reasons why competition might be limited.

---

[2] In addition a low reputation would be driven out of the market and so the value of having a low reputation would be zero—it cannot fall below this value.

[3] Hörner (2002) considers a model of reputation (that is a model based on beliefs as to the type of a firm) in which price has a signalling role and so even under perfect competition, a price premium above the cost of production is maintained, thereby maintaining the incentives to exert effort and maintain reputation. As discussed at some length below, I suppose that price has no signalling role and under perfect competition price would be driven down to marginal cost.

In recent and related independent work, Kranton (2003) argues that competitive effects reduce the price premium in a model where consumers can use only current price in addition to past quality to make inferences on the current quality. Thus in Kranton's paper, the first of the two effects highlighted in the paragraphs above is elaborated, that is the role of competition in preventing an equilibrium where high quality is maintained. However, the second effect considered here, the potential for some degree of competition to help in sustaining high quality is precluded by an assumption that a firm that produces low quality in one period can be held down to zero profits in all future periods. Thus the non-monotonic effect of competition on the ability to commit to high quality—the focus of this chapter—does not arise.

5.2. Model

In order to consider the effect of the degree of competition a firm faces on its ability to commit to high effort, we introduce a simple model. We suppose that there is only one firm in the industry, Firm A, for which there is a commitment problem—in each period Firm A, chooses whether to produce high quality products at some cost or low quality costlessly. The firm has a commitment problem inasmuch as profits would be higher if it could credibly commit to high quality but in each period customers and rival firms, though they observe the firm's quality in previous periods, cannot observe the firm's quality in the current period until after all purchases have been made. It is assumed that there are n rival firms and that they can only produce high quality.[4]

In each period, all firms set quantities. Given these quantities and consumers' expectations of quality, prices are determined. The realization of the firm's quality in previous periods informs customers' and rivals' expectations of quality. In particular, we suppose that if the firm produces low quality then it is supposed that the firm would produce low quality forever—thus we restrict attention to "trigger strategy" equilibria. Moreover, quality realizations are commonly and publicly observed at the end of each period by all customers and all rival firms.

We assume that customers' and rivals' expectations are not affected by any firm's choice of quantity—that is an individual firm's choice of quantity plays no signalling role. Note, in particular, that this assumption precludes the possibility that firms collude in

---

[4]We show in Appendix E.2 that relaxing this assumption and instead supposing that all firms make quality decisions leads to qualitatively similar results.

this simple framework.[5] We restrict attention to a Markov perfect equilibrium in which the state can take two values corresponding to whether or not Firm A has ever produced low quality.

We work with a linear demand model with quality indices.[6] In this model all consumers (there are assumed to be a measure 1 of consumers) have the same utility function defined over the $n + 1$ goods in the industry (Firm A's and those of its n rivals) and of a separated outside good available at a fixed price of unity. The following utility function for a consumer in period t ensures that the consumer's demand schedule for each good takes a relatively simple linear form:

$$U = \sum_i (x_{it} - \frac{x_{it}^2}{v_{it}^2}) - 2\sigma \sum_i \sum_{k<i} \frac{x_i}{v_i} \frac{x_k}{v_k} + M, \tag{5.1}$$

where x_{it} is the quantity of good i consumed, v_{it} is its quality, $\sigma \in [0,1]$ measures the degree of substitutability between different firms' offerings and M denotes the consumption of the outside good which is priced at unity. Writing the consumer's income as Y, it follows that $M = Y - \sum_i p_i x_i$.

In particular, this model implies that the price in period t for firm i's product is given by:

$$p_{it} = 1 - \frac{2x_{it}}{u_{it}^2} - \frac{2\sigma}{u_{it}} \sum_{j \neq i} \frac{x_{jt}}{u_{jt}}. \tag{5.2}$$

Note that the price will depend not on the quality of the good and of rival goods, but rather on their anticipated quality in the period denoted by u_{it}. In particular, all firms but firm A always produce high quality and so $u_{it} = h$ for all $i \neq a$ and for all t.

Firm A can choose either to produce quality l at a cost 0 or quality h at a cost c per unit in each period. As discussed below it is socially more efficient and Firm A would prefer to commit to high quality when $h(1 - c) > l$.[7] All firms seek to maximise the discounted value of profits where the discount factor is given by r.

---

[5]Kranton (2003) also considers a similar assumption with respect to firms. However, in that model, though customers (as here) can use all past realisations of quality to make inferences about current quality, they can only use the *current* price. Here in contrast they can use neither current nor past prices to make inferences about current quality.

[6]See Appendix 2.2 of Sutton (1998) for further details concerning this model.

[7]See the last paragraph of Section 5.3.2 and the Appendix for proof of this claim.

Below, we discuss the relationship between the possibility that an equilibrium in which the strategic firm can commit to exert high effort and the different measures of competition, σ and n. First however, it is useful to consider as benchmarks, the static outcomes when quality is observed before purchase rather than after, as is assumed through most of the chapter.

5.3. Static benchmarks

Before exploring the model discussed above it is useful to calculate per-period profits in the case where Firm A can commit to high or low quality (or equivalently in the case where customers can observe quality before purchase and rivals can observe Firm A's quantity before deciding their own production).

5.3.1. Firm A produces high quality

Suppose first that Firm A produces high quality goods, then it produces a_h where a_h maximizes its profit, which are given by:

$$(1 - \frac{2x}{h^2} - \frac{2\sigma}{h}\sum_{j \neq a}\frac{x_j}{h} - c)x. \tag{5.3}$$

From the first order condition, it follows that $a_h = \frac{h^2}{4}(1 - c) - \frac{\sigma}{2}\sum_{j \neq i}x_j$.

Similarly, each of the rival firms i, in equilibrium, chooses x to maximise:

$(1 - \frac{2x}{h^2} - \frac{2\sigma}{h}\sum_{j \neq i,a}\frac{x_j}{h} - \frac{2\sigma a_h}{h^2} - c)x$

and so $x_i = \frac{h^2}{4}(1 - c) - \frac{\sigma}{2}\sum_{j \neq i,a}x_j - \frac{\sigma a_h}{2}$.

By symmetry $x_j = x_h = a_h$ for all j and so $x_h = a_h = \frac{h^2}{4}(1 - c) - \frac{\sigma}{2}nx_h$ and so

$$a_h = \frac{h^2(1 - c)}{2(2 + \sigma n)}. \tag{5.4}$$

It follows trivially that the profit in this case is given by :

$$\pi_h = \frac{1}{2}\frac{h^2(1 - c)^2}{(2 + \sigma n)^2}, \tag{5.5}$$

and the price is given by $p_h = \frac{1-c}{2+\sigma n} + c$.

Note that the quantity, price and profit vary with respect to the parameters of the model in an intuitive way. In particular, profits are increasing in h and decreasing in c,

σ and n. In particular, competition (as measured either by n the number of rivals in the industry or σ the degree of substitution between the firm and its rivals) decreases both the price and the quantity and thereby the profits.

5.3.2. Firm A produces low quality

Suppose that firm A produces low quality. Suppose that it produces a_l and that the rivals each produce x_{hl}. Then a_l maximizes

$$(1 - \frac{2x}{l^2} - \frac{2\sigma}{l}n\frac{x_{hl}}{h})x \tag{5.6}$$

so $a_l = \frac{l^2}{4} - \frac{\sigma l}{2h}nx_{hl}$ so long as this is positive, otherwise $x_l = 0$.
The quantity produced by a rival firm maximizes

$$(1 - \frac{2x}{h^2} - \frac{2\sigma}{h}((n-1)\frac{x_{hl}}{h} + \frac{a_l}{l}) - c)x. \tag{5.7}$$

Taking the first order condition and assuming symmetry yields

$$x_{hl} = \frac{h^2}{2(2 + \sigma(n-1))}(1 - c) - \frac{h\sigma}{l(2 + \sigma(n-1)}a_l \tag{5.8}$$

and substituting back into the earlier expression for a_l yields:

$$a_l = \frac{l}{2}\frac{2l + \sigma l(n-1) - \sigma h(1-c)n}{(2 + \sigma n)(2 - \sigma)}, \tag{5.9}$$

so long as $2l + \sigma l(n - 1) - \sigma h(1 - c)n > 0$ when this condition fails then $x_l = \pi_l = 0$ and note that this condition is more likely to hold the less competitive the environment (the smaller σ and n) since $h(1 - c) > l$. When this condition holds then $p_l = \frac{2l+l\sigma(n-1)-h\sigma n(1-c)}{l(2+\sigma n)(2-\sigma)}$ and

$$\pi_l = \frac{1}{2}(\frac{2l + l\sigma(n-1) - h\sigma n(1-c)}{(2 + \sigma n)(2 - \sigma)})^2. \tag{5.10}$$

Note in particular that in this case, the profit for the low quality producer (and both price and quantity) are increasing in l and c and decreasing in h, furthermore it can readily be verified that they are decreasing in n and in σ.[8]

---

[8]For example $\frac{d\pi_l}{d\sigma} = -\frac{2l+l\sigma(n-1)-h\sigma n(1-c)}{(2+\sigma n)(2-\sigma)}\frac{4n(h(1-c)-\sigma l)+\sigma^2 n^2(h(1-c)-l)+\sigma^2 nl}{(2+\sigma n)^2(2-\sigma)^2}$
where $2l + l\sigma(n - 1) - h\sigma n(1 - c) < 0$ and $h(1 - c) > l$ so $(h(1 - c) - \sigma l) > 0$ and so $\frac{d\pi_l}{d\sigma} < 0$

Similarly in the case where $2l + \sigma l(n-1) - \sigma h(1-c)n > 0$ then $\pi_{hl} = \frac{1}{2}(\frac{2h(1-c)-\sigma l}{(2+\sigma n)(2-\sigma)})^2$ as one might expect is increasing in h and decreasing in l, c and n but $\frac{d\pi_{hl}}{d\sigma}$ may be either positive or negative.[9] This result is perhaps not surprising—it is the effect of business stealing from a weaker rival which implies that greater substitution between the two might benefit the stronger competitor.

Note that in the case that $2l + \sigma l(n-1) - \sigma h(1-c)n < 0$, $\pi_h = \frac{1}{2}\frac{h^2(1-c)^2}{(2+\sigma n)^2} > 0 = \pi_l$. When $2l + \sigma l(n-1) - \sigma h(1-c)n > 0$ then $\pi_h > \pi_l$ if and only if $\frac{1}{2}\frac{h^2(1-c)^2}{(2+\sigma n)^2} > \frac{1}{2}(\frac{2l+l\sigma(n-1)-h\sigma n(1-c)}{(2+\sigma n)(2-\sigma)})^2$ or equivalently $(2+\sigma(n-1))(h(1-c)-l) > 0$ and so in this case $\pi_h > \pi_l$ if and only if $h(1-c) - l > 0$. In either case if $h(1-c) - l > 0$ then $\pi_h > \pi_l$ and so as earlier claimed if $h(1-c) - l > 0$ then if Firm A could credibly commit to high quality production then it would choose to do so.

5.4. Commitment

In this section we seek to examine how competition affects the feasibility of an equilibrium in which Firm A commits to high quality. We restrict attention to a Markov perfect equilibrium in which the state is a simple indicator of whether or not Firm A has ever produced low quality and in which quantity has no signalling role. We return to consider this restriction in the concluding section.

First note that in any subgame in which Firm A had previously produced low quality, there is no action which would change the state and so the unique equilibrium strategy in this subgame would be for Firm A to produce low quality in each period and for the static optimal quantities to be produced.

Suppose that Firm A has not produced a low quality good, if in equilibrium in this state it produces high quality, then since the quantities chosen by Firm A and its rivals would not change the state, it is clear that the quantities chosen in each period would be as in the static case in Section 5.3.1, where Firm A was committed to high effort. In this section, essentially we ask whether this is sustainable.[10] In order to do so first consider

---

[9]This is also true for $x_{hl} = \frac{1}{2}h\frac{2h(1-c)-\sigma l}{4+2\sigma(n-2)-\sigma^2(n-1)}$ though $p_{hl} = \frac{1}{h}\frac{2h(1-c)-\sigma l}{4+2\sigma(n-2)-\sigma^2(n-1)} + c$ has a different comparative static with respect to c.

Note that
$\frac{d\pi_{hl}}{d\sigma} = (2h(1-c) - l\sigma)\frac{l\sigma^2 n + 4hn(1-c)(1-\sigma) - 4(h(1-c)-l)}{(2+\sigma n)^3(2-\sigma)^3}$

For example when $h(1-c) = 0.9$, $\sigma = 0.1$, $n = 1$ and $l = 0.05$ then $\frac{d\pi_{hl}}{d\sigma} < 0$, but when $h(1-c) = 0.9$, $\sigma = 0.1$, $n = 1$ and $l = 0.1$ then $\frac{d\pi_{hl}}{d\sigma} < 0$ (and in both cases $2l + \sigma l(n-1) - \sigma h(1-c)n > 0$).

[10]Note that there always exists an equilibrium in which Firm A always produces low quality.

the short-term gain from deviating—that is of producing low quality when expected to produce high quality.

The quantity that Firm A produce to maximise profits when producing low quality when rivals and customers anticipate high quality is a_d where a_d maximizes:

$$(1 - \frac{2a}{h^2} - \frac{2\sigma}{h^2}nx_h)a. \tag{5.11}$$

Note that customers expect high quality and as do rival firms so that the price is $(1 - \frac{2a}{h^2} - \frac{2\sigma}{h^2}nx_h)$. From the first order condition and substituting for x_h, it can readily be shown that $x_d = \frac{h^2}{4}\frac{2+\sigma cn}{2+\sigma n}$, $p_d = \frac{1}{2}\frac{2+\sigma cn}{2+\sigma n}$ and profit is given by:

$$\pi_d = \frac{h^2}{8}(\frac{2+\sigma cn}{2+\sigma n})^2. \tag{5.12}$$

5.4.1. Condition to ensure no deviation

It follows that a necessary and sufficient condition which ensures that there is a Markov perfect equilibrium in which firm A produces high quality output is given by:

$$\frac{1}{1-r}\pi_h \geq \pi_d + \frac{r}{1-r}\pi_l \tag{5.13}$$

Equivalently $\Delta = \frac{r}{1-r}(\pi_h - \pi_l) - (\pi_d - \pi_h) > 0$.

We consider the comparative statics of Δ with respect to n and σ, the measures of competition, examining the cases where $\pi_l = 0$ and $\pi_l > 0$ separately. Note that the former case is more likely for high values of σ (when $\sigma > \frac{2l}{h(1-c)n-l(n-1)}$) and for high values of n (when $n > l\frac{2-\sigma}{\sigma(h(1-c)-l)}$). Furthermore as anticipated in the introduction, for large enough n, $\Delta = \frac{1}{1-r}\pi_h - \pi_d = \frac{1}{1-r}\frac{1}{2}\frac{h^2(1-c)^2}{(2+\sigma n)^2} - \frac{h^2}{8}(\frac{2+\sigma cn}{2+\sigma n})^2$ and as $n \to \infty$, $\Delta \to -\frac{h^2c^2}{8} < 0$, that is for large enough n the condition fails.

5.4.2. Comparative statics with respect to n and σ

Case I: Low quality produces nothing after deviation $2l+\sigma l(n-1)-\sigma h(1-c)n < 0$

In this case $\Delta = \frac{1}{1-r}\pi_h - \pi_d = \frac{1}{1-r}\frac{1}{2}\frac{h^2(1-c)^2}{(2+\sigma n)^2} - \frac{h^2}{8}(\frac{2+\sigma cn}{2+\sigma n})^2 = \frac{h^2}{8(2+\sigma n)^2(1-r)}(4(1-c)^2 - (1-r)(2+\sigma cn)^2)$

The effect of a change in the degree of substitution between the offerings of different firms is given by:

$$\frac{d\Delta}{d\sigma} = -2n\frac{h^2}{8(2+\sigma n)^3(1-r)}(4(1-c)^2 - (1-r)(2+\sigma cn)^2) - \frac{h^22cn(2+\sigma cn)}{8(2+\sigma n)^2}$$

or, equivalently, $\frac{d\Delta}{d\sigma} = -2n\Delta - \frac{h^2 2cn(2+\sigma cn)}{8(2+\sigma n)^2}$. It follows that if $\Delta > 0$ then $\frac{d\Delta}{d\sigma} < 0$ but if $\Delta < 0$ then $\frac{d\Delta}{d\sigma}$ may be either positive or negative.

Similar results obtain for the comparative statics in this case with respect to n.[11]

Case B: Low quality produces a positive quantity after deviation $2l+\sigma l(n-1) - \sigma h(1-c)n > 0$

In this case

$\Delta = \frac{1}{1-r}\pi_h - \frac{r}{1-r}\pi_l - \pi_d$

$\Delta = \frac{h^2(4(1-c)^2 - (1-r)(2+\sigma cn)^2)}{8(2+\sigma n)^2(1-r)} - \frac{r}{2(1-r)}\left(\frac{2l+l\sigma(n-1)-h\sigma n(1-c)}{(2+\sigma n)(2-\sigma)}\right)^2$

The effect of a change in the degree of substitution between the offerings of different firms is given by:

$$\frac{d\Delta}{d\sigma} = \frac{1}{2}h^2 n\frac{2(c-r)(1-c)+\sigma cn(1-r)(1-c)}{(2+\sigma n)^3(1-r)}$$
$$+ \frac{r}{(1-r)}n\frac{2l+l\sigma(n-1)-h\sigma n(1-c)}{(2+\sigma n)(2-\sigma)}\frac{4(h(1-c)-l\sigma)-\sigma^2(n-1)l+\sigma^2 nh(1-c)}{(2+\sigma n)^2(2-\sigma)^2}$$

The second term is always positive, however since the first term may be either positive or negative, the sign of this expression is ambiguous. Similar results apply with respect to n.[12]

Thus overall the effect of a change in the degree of competition, measured either by the number of rival of firms in the industry or as the degree of substitution between the products of different of different firms, has ambiguous effects on Δ, or equivalently on the possibility that a Markov Perfect equilibrium in which Firm A would produce high quality.

The figures below illustrate that an increase in the degree of competition can have ambiguous effects, holding all other parameters of the model constant. For example, Figure 5.1 illustrates that when $h = 0.81$, $l = 0.63$, $c = 0.1$, $n = 5$ and $r = 0.485$, then $\Delta > 0$ holds at $\sigma = 0$ but as σ increases, it fails, then holds and then fails again.

Figure 5.2 illustrates that the effect a change in the number of rivals may have an ambiguous effect on Firm A's ability to commit to producing high quality. The figure, which plots Δ against $\log_2 n$ with $h = 0.8$, $l = 0.77$, $c = 0.02$, $\sigma = 0.5$ and $r = 0.51$,

---

[11] $\frac{d\Delta}{dn} = -2\sigma\frac{h^2}{8(2+\sigma n)^3(1-r)}(4(1-c)^2 - (1-r)(2+\sigma cn)^2) - \frac{h^2 2\sigma c(2+\sigma cn)}{8(2+\sigma n)^2}$

$\frac{d\Delta}{dn} = -2\sigma\Delta - \frac{h^2 2\sigma c(2+\sigma cn)}{8(2+\sigma n)^2}$

and so if $\Delta > 0$ then $\frac{d\Delta}{dn} < 0$ but if $\Delta < 0$ then $\frac{d\Delta}{dn}$ may be either positive or negative.

[12] $\frac{d\Delta}{dn} = \frac{1}{2}\sigma h^2\frac{2(c-r)(1-c)+\sigma cn(1-r)(1-c)}{(2+\sigma n)^3(1-r)} + \frac{r\sigma}{(1-r)}\frac{2l+l\sigma(n-1)-h\sigma n(1-c)}{(2+\sigma n)(2-\sigma)}\frac{2h(1-c)-l\sigma}{(2+\sigma n)^2(2-\sigma)}$
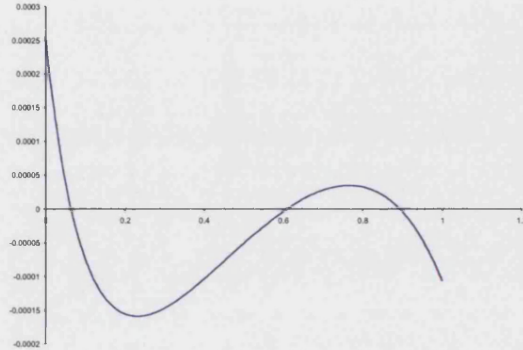
Figure 5.1. Delta against sigma at h=0.81, l=0.63, c=0.1, n=5 and r=0.485
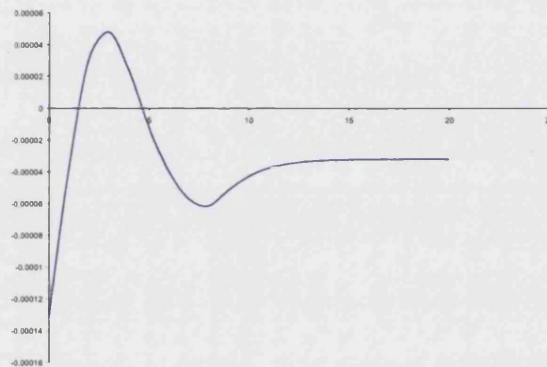
Figure 5.2. Delta against log(n) at h=0.8, l=0.77, c=0.02, sigma=0.5 and r=0.51

shows that $\Delta > 0$ when n is between 3 and 23, but if there are fewer than 3 or more than 23 rivals then the condition fails.

Suppose that there is an additional gain to producing high quality (such as pride in producing high quality or reputational spillovers to other products) then there is an equilibrium in which Firm A produces high quality so long as $\Delta > B$. As might be anticipated, given the results on $\frac{d\Delta}{d\sigma}$ in the case where $2l + \sigma l(n-1) - \sigma h(1-c)n < 0$, this might lead to even more peculiar relationships between the possibility of maintaining such an equilibrium and measures of competition. For example Figure 5.3 illustrates
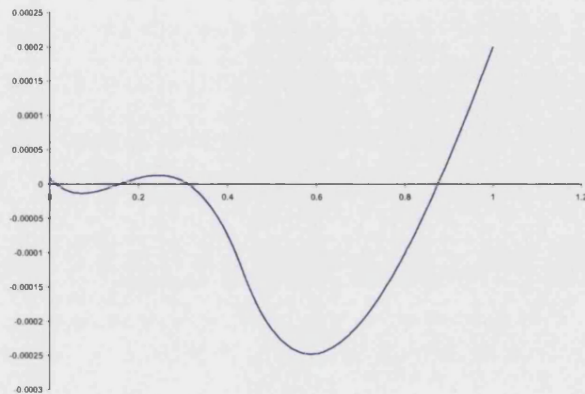
Figure 5.3. Delta-B against sigma at h=0.8, l=0.2, c=0.3, n=2, r=0.405 and B=0.01753

shows that how the condition $\Delta > B$ varies with σ for $h = 0.8$, $l = 0.2$, $c = 0.3$, $n = 3$, $r = 0.405$ and $B = 0.01753$.

5.5. Discussion

This chapter, and the examples in the results in the section above essentially make one simple point—that the degree of competition has ambiguous effects on a firm's ability to commit to high quality. Holding all other parameters of the model constant, an increase in the level of competition (measured either as the number of rival firms or the degree of substitution between different firms' products) can make the possibility of an equilibrium in which the firm produces high quality either more or less likely. This has a number of implications related to different strands of literature.

5.5.1. Competition and incentives

As Nickell (1996) states "Most people believe that competition is a good thing." However, empirical evidence that competition improves corporate performance is weak and theoretical results to date have been ambiguous, as discussed by Nickell (1996) and Shmidt (1997) for example. The formal theoretical work, beginning with Hart (1983), has primarily focussed on explicit contractual incentives and effects which are essentially static,

that is effects which are manifest in a single period of production of competition.[13] In addition, while much of this literature relies on better signals of managerial effort available in more competitive environments through comparative performance evaluation, this affect is absent from the model presented in this chapter, in which at the end of a period there is no uncertainty about what happened. A number of recent papers also consider how industry structure affects the returns to skill and effort and implications for wage inequality (Guadalupe (2003)), contracts (Cuñat and Guadalupe (2003a and b)) and organizational design (Harstad (2003)).

Recently, Raith (2003) argues that when the market structure is endogenous, a clear theoretical prediction emerges. Overall profits (including any sunk costs of entry) must be equal to zero and so in Raith's model greater competition (measured by a higher degree of substitution between different firms' offerings) implies larger output and so a greater incentive to reduce costs. Note that such an argument loses force in the model considered here, zero profits might be associated with either high or low quality and so the zero profit condition in a free entry version of the model presented in this noted would be insufficient to tie down output—in particular this observation suggests the possibility of many industry structures in equilibrium.

This chapter adds to the confusion in suggesting that competition has ambiguous effects on a firm's reputational incentives. Above, we focus on a single firm and for a given market structure. However, the intuition highlighted in the introduction and demonstrated above—that competition reduces the returns to being committed to high quality, the returns to a low quality producer and the costs of maintaining the expectation of high quality, but all at different rates—should be robust to relaxing both of these assumptions.[14]

Hörner (2002) presents a model of perfect competition in which an equilibrium where all firms exert high effort exists. An important feature in that model is that price can play a signalling role—that is firms set prices and customers make inferences on the quality of the firm's offering on the basis of the price set. In particular this assumption

---

[13]An exception is Meyer and Vickers (1997), discussed in Nickell (1996). Here the role that greater competition plays is to allow for comparative performance evaluation—essentially for better information about the performance of a manager. This can increase reputational incentives but the ratchet effect may imply that rather incentives are dulled.

[14]In the case where the market structure is endogenous. It is not clear that an increase in the degree of substitution would lead to a lower cost of maintaining reputation. Indeed such an increase in competition may lead to higher output and so a higher cost of maintaining reputation. But again there is no reason to suspect an overall monotonic effect on reputational incentives.

allows for a price premium above the cost even in a market with many firms whose products (when they all produce high quality) are indistinguishable. By contrast in this chapter, the only information that is assumed to guide customers as to the quality of the firm's offering is previous quality realization; in particular, in a perfectly competitive environment there can be no price premium.

5.5.2. Competitive strategy

A clear implication of the intuition presented in this chapter is that there may be occasions where a firm would be better off in a more competitive environment. By encouraging a firm to enter the industry or by making its own product more substitutable with rivals' products, a commitment to produce high quality could become credible and profits higher. Such intuition can easily be borne out through specific examples in the model.[15]

Other papers have noted that competition can help a firm's ability to commit to behave well, but in these papers, this is essentially a commitment to behave well in the future. For example, Farrell and Gallini (1988) and Shepard (1987) highlight that second-sourcing, or encouraging future competition constrains a monopolist's ability to charge a high price or produce at a low observed quality in the future and so might benefit the monopolist in increasing current sales to customers, who must incur product specific set-up costs. Dudey (1990) and Wernerfelt (1994) argue that firms might locate near each other as a means of committing to price at a reasonably low level and so encourage customers to pay the search costs required to visit the firms. In these models, therefore increasing competition increases a firm's commitment to future good behaviour and this has current benefits. In the model presented in this chapter, by contrast, the commitment problem is with respect to current unobservable behaviour and so whereas

---

[15]For the case of increasing profits through changing σ, it would be hard to interpret such a change when $n > 1$ (it is not obvious how could an individual firm affect the degree of substitution between two other firms. In the case $n = 1$ where changing the degree of substitution between its own and its rivals products is a more plausible modelling assumption, it is easy to construct examples in which Firm A could benefit from a higher σ. This is the case for example at $h = 0.3$, $l = 0.18$, $c = 0.1$, and σ around 0.54.

Encouraging another firm into the industry may be a plausible modelling assumption (for example in consulting this may be a choice to allow or encourage some partners to spin off or in general it may involve licensing or otherwise diffusing specialist knowledge). If this is possible, again a firm may benefit from doing so. For example, at $h = 0.8$, $l = 0.71$, $c = 0.05$, $\sigma = 0.6$, Firm A can commit to high quality and earn higher per-period profits when $n = 4$ than when $n = 3$ and it can make no such commitment.

in those papers more competition always helps to overcome the commitment problem, here as described above the effect is non-monotonic.

5.5.3. Caveats

The model in this chapter is deliberately made simple to make the point that competition has ambiguous effects on reputational incentives starkly. However, there are a number of issues which ought to be considered in refining this intuition in practical application.

In the model, we have focused on a single firm's incentives to produce high quality, taking as given that all other firms in the industry are committed to producing at high quality. It is certainly the case that as the degree of competition in the industry changes, the profits to holding a high reputation and the profits to holding a low reputation as well as the cost of maintaining reputation will all change and at different rates. Thus, it seems reasonable to suppose (and indeed it can be shown) that similar effects might apply if all firms had quality decisions to make in each period, [16] and if the number of firms in the industry was endogenously determined.

---

[16] In the spirit of Abreu (1988), to consider the viability of an equilibrium in which all firms produce high quality, one should consider a most severe feasible punishment continuation equilibrium for a firm that deviates. However, determining which continuation equilibria are feasible is not a trivial exercise (in particular it may not be feasible that all rival firms produce at high quality after one firm has produced at low quality).

Nevertheless, it can be shown in the static version of the model where all firms are committed either to low or high quality, that in an industry with $n + 1$ firms, profits for a firm committed either to low or high quality are non-decreasing in the number of firms committed to low quality. It then follows, that condition 5.13 defines a harshest (though possibly not feasible) continuation equilibrium and so can be thought of as a necessary condition for the existence of an equilibrium in which all firms produce high quality in a dynamic version of the model. Since a continuation where all firms produce low quality is always possible, a sufficient condition which ensures the existence of an equilibrium in which all firms produce high quality is given by:

$$\frac{1}{1-r}\pi_h \geq \pi_d + \frac{r}{1-r}\pi_{ll},$$

where $\pi_{ll} = \frac{1}{2}\frac{l^2}{(2+\sigma(n-1))^2}$ is the short-term profit that a firm in the industry earns when all the firms produce low quality and each optimises with respect to the quantity produced.

Examining this condition together with condition 5.13 allows us to identify parameters regions in which an equilibrium in which all firms produce high quality can exist and regions where such an equilibrium cannot exist (though there are some regions where the existence of such an equilibrium cannot be determined). However such an analysis can show that the existence of such an equilibrium is non-monotonic in the degree of competition.

This argument is elaborated in the Appendix.

It should be noted that collusion does not appear in the model and is left undiscussed. This is perhaps, the most widely discussed topic in formal dynamic models of industry outcomes with a fixed number of firms and so the omission may be a serious one. In part collusion does not arise in the model as a deliberate choice. Clearly the degree of competition in an industry affects firms' ability to collude; however the model presented here is intended to be as simple as possible to illustrate a simple point that competition has ambiguous effects on a firm's incentives to maintain reputation with respect to its customers, which I believe would be robust but perhaps more obscure to identify in a model that allowed for collusion between firms in the industry.

The modelling choice that prevents collusion from arising in the model is that firms and customers do not alter their strategies in the light of the past or current quantity decisions that firms make (only on the basis of the firm's quality decision). On the part of firms, as discussed above this assumption precludes any possibility of collusion. On the part of customers it suggests that customers base their assessment of a firm's current quality on characteristics of the industry as whole and on the quality of a firm's previous products. The current quantity decision of a firm may be hard for customers to observe or in some sense may be less salient to customers and, to me at least, the assumption that it does not affect customers' expectations of quality certainly seems plausible and worth investigating, as we have done in this chapter. Note however, that this is a fundamental assumption in the model. This assumption as it relates to rival firms may also be plausible.

References

[1] ABREU, D. (1988): "On the Theory of Infinitely Repeated Games with Discounting," *Econometrica*, **56**, 383-396.

[2] ABREU, D., D. PEARCE, and E. STACHETTI (1990): "Towards a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, **58**, 1041-1064.

[3] ALCHIAN A. A. and H. DEMSETZ (1972): "Production, Information Costs, and Economic Organization", *American Economic Review*, **62**, 777-795.

[4] ANDERSON, A. and L. SMITH (2002): "Assortative Matching, Reputation and the Beatles Break-Up," working paper.

[5] BAGWELL K. and G. RAMEY (1994): "Coordination Economies, Advertising, and Search Behavior in Retail Markets," *American Economic Review*, **84**, 498-517.

[6] BAKER, G., R. GIBBONS and K. J. MURPHY (1994): "Subjective Performance Measures in Optimal Incentive Contracts," *Quarterly Journal of Economics*, **109**, 1125-1156.

[7] BAKER, G., R. GIBBONS and K. J. MURPHY (2002): "Relational Contracts and the Theory of the Firm," *Quarterly Journal of Economics*, **117**, 39-84.

[8] BENABOU, R. and G. LAROQUE (1992): "Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility," *Quarterly Journal of Economics*, **107**, 921-958.

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

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

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

[12] BERNHEIM, B. D., and M. D. WHINSTON (1990): "Multimarket Contact and Collusive Behavior," *RAND Journal of Economics*, **21**, 1-26.

[13] BOOT, A. W. A. and T. T. MILBORN (2002): "Credit Ratings as Coordination Mechanisms," working paper, Tinbergen Institute DP 2002 - 058/2

[14] BRANDENBURGER, A. and B. POLAK (1996): "When Managers Cover Their Posteriors: Making the Decisions the Market Wants to See," *RAND Journal of Economics*, **27**, 523-541.

[15] BRETON, M., P. ST-AMOUR and D. VENCATACHELLUM (2003): "Dynamic production teams with strategic behavior", *Journal of Economic Dynamics and Control*, **27**, 875-905.

[16] BRETON, M., P. ST-AMOUR and D. VENCATACHELLUM (2002): "Inter- vs Intragenerational teams: A Young Worker's Perspective", working paper.

[17] CABRAL, L. (2000): "Stretching Firm and Brand Reputation" *RAND Journal of Economics*, **31**, 658-673.

[18] CARMICHAEL, H. L. (1988): "Incentives in Academics: Why is There Tenure?" *Journal of Political Economy*, **96**, 453-472.

[19] CHE, Y-K. and S-W. YOO (2001): "Optimal Incentives for Teams," *American Economic Review*, **91**, 525-541.

[20] CHOI, J. P. (1998), "Brand Extension as Informational Leverage," *Review of Economic Studies*, **65**, 655-669.

[21] CLARK, R. and I. HORSTMANN (2001): "Advertising and Coordination in Markets with Consumption Scale Effects," working paper.

[22] COTTERMAN J. D. (ed.) (2001): *Compensation Plans for Law Firms*, Altman Weil Inc., American Bar Association.

[23] CREMER, J. (1986): "Cooperation in Ongoing Organizations," *Quarterly Journal of Economics*, **101**, 33-55.

[24] CREMER, J. (1995): "Arm's Length Relationships," *Quarterly Journal of Economics*, **110**, 275-295.

[25] CRIPPS, M. W., G. J. MAILATH and L. SAMUELSON (2004) "Imperfect Monitoring and Impermanent Reputation," *Econometrica*, **72**, 407-432.

[26] CUNAT, V. and M. GUADALUPE (2003a): "Executive Compensation and Product Market Competition," working paper.

[27] CUNAT, V. and M. GUADALUPE (2003b): "How Does Product Market Competition Shape Incentive Contracts?" working paper.

[28] DARBY, M. and E. KARNI (1973): "Free Competition and the Optimal Amount of Fraud," *Journal of Law and Economics*, **16**, 67-88.

[29] DIAMOND, D. W. (1989): "Reputation Acquisition in Debt Markets," *Journal of Political Economy*, **97**, 828-862.

[30] DUDEY, M. (1990): "Competition by Choice: The Effect of Consumer Search on Firm Location Decision," *American Economic Review*, **80**, 1092-1104.

[31] FARRELL, J. and N. T. GALLINI (1988): "Second-Sourcing as a Commitment: Monopoly Incentives to Attract Competition," *Quarterly Journal of Economics*, **103**, 673-694.

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

[33] FARRELL, J. and S. SCOTCHMER (1988): "Partnerships," *Quarterly Journal of Economics*, **103**, 279-297.

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

[35] FELTOVICH, N., R. HARBAUGH and T. TO (2002) "Too Cool for School? Signaling and Countersignaling," *RAND Journal of Economics*, **33**, 630-649

[36] FISHMAN, A. and R. ROB (2003): "Is Bigger Better? Customer Base Expansion through Word of Mouth Reputation," working paper.

[37] FOMBURN, C. J. (1996) *Reputation: Realizing Value from the Corporate Image*, Harvard Business School Press.

[38] FUDENBERG, D. and D. K. LEVINE (1989): "Reputation and Equilibrium Selection in Games with a Patient Player," *Econometrica*, **57**, 759-778.

[39] FUDENBERG, D. and D. K. LEVINE (1992): "Maintaining a Reputation When Strategies Are Imperfectly Observed," *Review of Economic Studies*, **59**, 561-579.

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

[41] GALANTER, M. and T. PALAY (1991): *Tournament of Lawyers*, Chicago University Press.

[42] GARICANO, L. and J. SANTOS (2004): "Referrals," *American Economic Review*, forthcoming.

[43] GIBBONS, R. and K. J. MURPHY (1992): "Optimal Incentive Contracts in the presence of Career Concerns," *Journal of Political Economy*, **100**, 468-505.

[44] GILSON, R. and R. MNOOKIN (1985): "Sharing Among Human Capitalists: An Economic Enquiry into the Corporate Law Firm and How Partners Split Profits," *Stanford Law Review*, **37**, 313-392.

[45] GLAZER, A. (2002): "Allies as rivals: Internal and external rent seeking," *Journal of Economic Behavior and Organization*, **48**, 155-162.

[46] GLAZER, A. and B. SEGENDORFF (2001): "Reputation in team production," Working Paper Series in Economics and Finance, Working Paper No. 425, Stockholm School of Economics.

[47] GREEN, E. and R. PORTER (1984): "Non Cooperative Collusion Under Imperfect Price Information," *Econometrica*, **52**, 87-100.

[48] GREENWALD B. (1986), "Adverse Selection in the labour market," *Review of Economic Studies*, **53**, 325-247.

[49] GUADALUPE, M. (2003): "Does Product Market Competition Increase Wage Inequality," working paper.

[50] HARRIS, M. and Y. WEISS (1984): "Job Matching with Finite Horizon and Risk Aversion," *Journal of Political Economy*, **92**, 759-779.

[51] HARSTAD, B. (2003): "Organizations and Careers," working paper.

[52] HART, O. D. (1983): "The Market Mechanism as an Incentive Scheme," *Bell Journal of Economics*, **14**, 366-382.

[53] HERMALIN, B. (2001): "Economics and Corporate Culture," in *Handbook of Organizational Culture and Climate*, eds. S. Cartwright et al. Chichester, England: John Wiley & Sons Ltd.

[54] HOLMSTROM, B. (1982): "Moral Hazard in Teams," *Bell Journal of Economics*, **13**, 324-340.

[55] HOLMSTROM, B. (1999): "Managerial Incentive Problems—A Dynamic Perspective," *Review of Economic Studies*, **66**, 169-182; (originally appeared (1982) in *Essays in Honor of Professor Lars Wahlback*).

[56] HOLMSTROM, B. and R. MYERSON (1983): "Efficient and Durable Decision Rules with Incomplete Information," *Econometrica*, **51**, 1799-1819.

[57] HÖRNER, J. (2002): "Reputation and Competition," *American Economic Review*, **92**, 644-663.

[58] JEON, S. (1996): "Moral Hazard and Reputational Concerns in Teams: Implications for Organizational Choice," *International Journal of Industrial Organization*, **14**, 297-315.

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

[60] KAHN, C. and G. HUBERMAN (1988): "Two-Sided Uncertainty and Up-or-out Contracts," *Journal of Labor Economics*, **6**, 423-444.

[61] 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*.

[62] KLEIN, B. and K. B. LEFFLER (1981): "The Role of Market Forces in Assuring Contractual Performance," *Journal of Political Economy*, **89**, 615-641.

[63] KOSZEGI, B. and W. LI (2002): "Ambition and Talent," working paper.

[64] KRANTON, R. E. (2003): "Competition and the Incentive to Produce High Quality," *Economica*, **70**, 385-404.

[65] KREPS, D. (1990): "Corporate Culture and Economic Theory" in J. Alt and K. Shepsle (eds.), *Perspectives on Positive Political Economy*, (Cambridge: Cambridge University Press), 90-143.

[66] KREPS, D., P.R. MILGROM, D. J. ROBERTS and R. WILSON (1982): "Rational Cooperation In The Finitely Repeated Prisoner's Dilemma," *Journal of Economic Theory*, **27**, 245-252.

[67] KREPS, D., and R. WILSON (1982): "Reputation and Imperfect Information," *Journal of Economic Theory*, **27**, 253-279.

[68] LEVIN, J. (2001): "Career Concerns and Collective Reputation," working paper.

[69] LEVIN, J. and S. TADELIS (2002): "A Theory of Partnerships," working paper.

[70] MAILATH, G. J., and L. SAMUELSON (2001): "Who Wants a Good Reputation?" *Review of Economic Studies*, **68**, 415-441.

[71] MAILATH, G. J., and L. SAMUELSON (1998): "Your Reputation Is Who You're Not, Not Who You'd Like To Be," CARESS working paper 98-11, University of Pennsylvania.

[72] MEYER, M. A., T. E. OLSEN and G. TORSVIK (1996): "Limited intertemporal commitment and job design," *Journal of Economic Behavior and Organization*, **31**, 401-417.

[73] MEYER, M. A. and J. VICKERS (1997): "Performance Comparisons and Dynamic Incentives," *Journal of Political Economy*, **105**, 547-581.

[74] MILGROM, P. R. and D. J. ROBERTS (1982): "Predation, Reputation and Entry Deterrence," *Journal of Economic Theory*, **27**, 280-312.

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

[76] MILGROM, P. R. and D. J. ROBERTS (1992): *Economics, Organization & Management*, Prentice Hall.

[77] MORRISON, A. D. and W. J. WILHELM Jr. (2003): "Partnership Firms, Reputation and Human Capital," working paper.

[78] NELSON, R. L. (1988) *Partners with Power: The Social Transformation of the Large Law Firm*, University of California Press.

[79] NICKELL, S. J. (1996): "Competition and Corporate Performance," *Journal of Political Economy*, **104**, 724-746.

[80] O'FLAHERTY, B. and A. SIOW (1992): "On the Job Screening, Up or Out Rules and Firm Growth," *Canadian Journal of Economics*, **25**, 346-368.

[81] O'FLAHERTY, B. and A. SIOW (1995): "Up or Out Rules in the Market for Lawyers," *Journal of Labor Economics*, **13**, 709-735.

[82] OTTAVIANI, M. (1999): "Monopoly Pricing with Social Learning," working paper.

[83] PASTINE, I. and T. PASTINE (1999): "Coordination in Markets with Consumption Externalities: The Role of Advertising and Product Quality," working paper.

[84] PHELAN, C. (2001): "Public Trust and Government Betrayal," Federal Reserve Bank of Minneapolis Research Department Staff Report 253.

[85] PINAULT, L. (2001): *Consulting Demons: Inside the Unscrupulous World of Global Corporate Consulting*, HarperBusiness.

[86] PRENDERGAST, C. (1993): "The Role of Promotion in Inducing Specific Human Capital Acquisition" *Quarterly Journal of Economics*, **108**, 523-534.

[87] PRENDERGAST, C. (1999): "The Provision of Incentives in Firms" *Journal of Economic Literature*, **37**, 7-63.

[88] RAITH, M. (2003): "Competition, Risk and Managerial Incentives," *American Economic Review*, **93**, 1425-1436.

[89] REBITZER, J. B. and L. J. TAYLOR (2001): "When Knowledge is an Asset: Explaining the Organizational Structure of Large Law Firms," working paper.

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

[91] SCHELLING, T. C. (1960) *The Strategy of Conflict*, Harvard University Press.

[92] SEGENDORFF, B. (2000): "A signalling theory of scapegoats," Working Paper Series in Economics and Finance, Working Paper No. 406, Stockholm School of Economics.

[93] SHAPIRO, C. (1983): "Premiums for High Quality Products as Returns to Reputations," *Quarterly Journal of Economics*, **98**, 659-680.

[94] SHEPARD, A. (1987): "Licensing to Enhance Demand for New Technologies," *RAND Journal of Economics*, **18**, 360-368.

[95] SHMIDT, K. (1997): "Managerial Incentives and Product Market Competition," *Review of Economic Studies*, **64**, 191-213.

[96] SHERSTYUK K. (1998): "Efficiency in partnership structures," *Journal of Economic Behavior & Organization*, **36**, 331-346.

[97] SPULBER, D. (1999): *Market Microstructure: Intermediaries and the Theory of the Firm*, Cambridge University Press.

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

[99] SUTTON, J. (1998) *Technology and Market Structure: Theory and History*, MIT Press.

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

[101] TADELIS, S. (2002): "The Market for Reputations as an Incentive Mechanism," *Journal of Political Economy*, **92**, 854-882.

[102] TADELIS, S. (2003): "Firm Reputation with Hidden Information," *Economic Theory*, **21**, 635-651.

[103] TIROLE, J. (1996): "A Theory of Collective Reputations (with Applications to the Persistence of Corruption and to Firm Quality)," *Review of Economic Studies*, **63**, 1-22.

[104] TROTTER, M. H. (1997) *Profit and the Practice of Law: What's Happened to the Legal Profession*, The University of Georgia Press.

[105] TURNOR, R. (2001): "Partnership and Members' Agreements" in SACKER, T. (ed.) *A Practitioner's Guide to Partnership and LLP Law and Regulation*, City & Financial Publishing

[106] VAN DEN STEEN, E. (2001): "Organizational Beliefs and Managerial Vision," working paper.

[107] WALDMAN, M. (1984): "Job Assignments, Signalling and Efficiency," *RAND Journal of Economics*, **15**, 255-267.

[108] WERNERFELT, B. (1988), "Umbrella Branding as a Signal of Product Quality: An Example of Signalling by Posting a Bond," *RAND Journal of Economics*, **19**, 458-466.

[109] WERNERFELT, B. (1994): "Selling Formats for Search Goods," *Marketing Science*, **13**, 298-309.

APPENDIX A

Proofs for Chapter 2

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:

$$S(f(\lambda)) = \lambda g + (1-\lambda)b - c + \beta(\lambda g + (1-\lambda)b)\max\{0, f(\lambda^s)\} \qquad (A.1)$$
$$+\beta(1 - \lambda g - (1-\lambda)b)\max\{0, f(\lambda^f)\}.$$

where λ^s and λ^f are as defined in Equations (2.1) and (2.2) above.

We then proceed as follows: First we prove that it is true that $S : B[0,1] \longrightarrow B[0,1]$, that is we show S takes bounded, continuous real-valued functions with domain $[0,1]$ to bounded continuous real-valued functions with domain $[0,1]$. 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 (2.3), this requires that 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 $x \in [0,1]$ and 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.[1] Finally we prove that S takes increasing functions to increasing functions. That S preserves continuity and preserves the monotonically increasing property is sufficient for ensuring that $V^u(\lambda)$ is continuous and increasing.

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 $|x - y| < \delta$, $|(Sf)(x) - (Sf)(y)| < \varepsilon$. Consider first any $z \in [0,1]$, then after rearranging expressions, we can write

---

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

$$| (Sf)(x) - (Sf)(z) | = \qquad | x - z | (g - b) +$$
$$+ | \max \{0, f(x^s)\} - \max \{0, f(z^s)\} |$$
$$+ | \max \{0, f(x^f)\} - \max \{0, f(z^f)\} | \qquad . \quad (A.2)$$
$$+ | x - z | (g - b) | \max \{0, f(x^s)\} - \max \{0, f(z^f)\} |$$

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

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

$$| (Sf)(x) - (Sf)(y) | < \frac{\varepsilon}{4}(g - b) + \frac{\varepsilon}{4} + \frac{\varepsilon}{4} + \frac{\varepsilon}{4M}M < \varepsilon. \qquad (A.3)$$

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.

$$(S(f + a))(x) = \qquad xg + (1 - x)b - c$$
$$+ \beta(xg + (1 - x)b) \max \{0, f(\frac{xg}{xg+(1-x)b}) + a\} \qquad . \quad (A.4)$$
$$+ \beta(1 - xg - (1 - x)b) \max \{0, f(\frac{x(1-g)}{x(1-g)+(1-x)(1-b)}) + a\}$$

So

$$(S(f + a))(x) \leqslant \qquad \beta a + xg + (1 - x)b - c$$
$$+ \beta(xg + (1 - x)b) \max \{0, f(\frac{xg}{xg+(1-x)b})\} \qquad , \quad (A.5)$$
$$+ \beta(1 - xg - (1 - x)b) \max \{0, f(\frac{x(1-g)}{x(1-g)+(1-x)(1-b)})\}$$

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 $(Sf)(\lambda) - (Sf)(\mu)$. We can write

$$(Sf)(\lambda) - (Sf)(\mu) = (\lambda - \mu)(g - b)$$
$$+\beta(\mu g + (1 - \mu)b)(\max\{0, f(\lambda^s)\} - \max\{0, f(\mu^s)\})$$
$$+\beta(1 - \lambda g - (1 - \lambda)b)(\max\{0, f(\lambda^f)\} - \max\{0, f(\mu^f)\})$$
$$+\beta(\lambda - \mu)(g - b)(\max\{0, f(\lambda^s)\} - \max\{0, f(\mu^f)\})$$

$$(A.6)$$

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

$$(Sf)(\lambda) - (Sf)(\mu) \geq (\lambda - \mu)(g - b) > 0. \qquad (A.7)$$

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 (2.4) and (2.5), $\mu = \lambda$ 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 $\lambda < \lambda^*$.[2]

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

---

[2] 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.

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 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:

$$(S_\gamma f)(x) = \gamma[xg + (1-x)g] - c + \beta[b\max\{0, f(x^s)\} + (1-b)\max\{0, f(x^f)\}]. \quad \text{(A.8)}$$

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 $W_{\gamma_1}^b(\lambda) > W_{\gamma_2}^b(\lambda)$ where these are defined in a consistent way to the above notation, that is where $W_\gamma^b(\lambda)$ 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 $W_\gamma^b(\lambda) = 0$, and so since $W_{\gamma_1}^b(\lambda) > W_{\gamma_2}^b(\lambda)$ and these are increasing functions, $\lambda^*(\gamma_1) < \lambda^*(\gamma_2)$.[3] 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 4

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.

---

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

$$W^b(r, \lambda) = \lambda g + (1 - \lambda)b - c + \beta(\underline{\lambda}(r, \lambda)g + (1 - \underline{\lambda}(r, \lambda)b)\max\{0, W^b(r, \lambda^s)\}$$

$$+\beta(1 - \underline{\lambda}_t(r, \lambda)g - (1 - \underline{\lambda}_t(r, \lambda)b)\max\{0, W^b(r, \lambda^f)\}. \tag{A.9}$$

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 3, 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. Note μ_t depends on the beginning of period belief and on the belief about the seller's strategy, in particular on d the anticipated probability that a seller with a low signal would cease trading. 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.

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.[4] 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 5

Using the notation of the Proof of Proposition 4, 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

---

[4]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.

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 Lemma 6

We prove only (i), (ii) can be proven similarly.

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

$$\frac{\frac{\lambda g}{\lambda g+(1-\lambda)(1-\alpha)b}}{\frac{\lambda g}{\lambda g+(1-\lambda)(1-\alpha)b} + \frac{(1-\lambda)(1-\alpha)b}{\lambda g+(1-\lambda)(1-\alpha)b}(1-\alpha)b} = \frac{\lambda g}{\lambda g + (1-\lambda)(1-\alpha)^2 b^2}. \qquad (A.10)$$

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 is 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

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

side of the inequality is strictly less than

$$\frac{\lambda}{\lambda+(1-\lambda)(1-\alpha)}g + b - c + \sum_{n=1}^{\infty} \beta^n \left(\frac{\lambda g^n}{\lambda g^n+(1-\lambda)(1-\alpha)^{n+1}b^{n+1}}g + b - c \right)$$
$$< \frac{\lambda}{(1-\lambda)(1-\alpha)}g + b - c + \sum_{n=1}^{\infty} \beta^n \left(\frac{\lambda g^n}{(1-\lambda)(1-\alpha)^{n+1}b^{n+1}}g + b - c \right)$$
$$< \frac{\lambda}{(1-\lambda)(1-\alpha)}g + b - c + \frac{\lambda}{(1-\lambda)(1-\alpha)}g \frac{(1-\alpha)b}{(1-\alpha)b-\beta g} - (c-b)\frac{1}{1-\beta}. \text{ For } \lambda \text{ low enough this is}$$

negative which is the contradiction that we needed. ∎

Proof of Lemma 7

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

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

and

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

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

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

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

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 $(0, 1)$.[5] 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 hypothesized form. ∎

Proof of Proposition 8

We begin by introducing the following notation:

---

[5]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.

$$V_\alpha^g(\lambda; l, m) = g\mu(l, m) + (1 - \mu(l, m))b - c$$
$$+\beta[g\alpha \max\{0, V_\alpha^g(\mu(l, m)^s)\} + g(1 - \alpha)V_\alpha^g(\mu(l, m)^s) \quad \text{and}$$
$$+(1 - g)\alpha \max\{0, V_\alpha^g(\mu(l, m)^f)\} + (1 - g)(1 - \alpha)V_\alpha^g(\mu(l, m)^f)]$$
$$\text{(A.11)}$$

$$V_\alpha^b(\lambda; l, m) = g\mu(l, m) + (1 - \mu(l, m))b - c$$
$$+\beta[b\alpha \max\{0, V_\alpha^b(\mu(l, m)^s)\} + b(1 - \alpha)V_\alpha^b(\mu(l, m)^s) \quad \text{(A.12)}$$
$$+(1 - b)\alpha \max\{0, V_\alpha^b(\mu(l, m)^f)\} + (1 - b)(1 - \alpha)V_\alpha^b(\mu(l, m)^f)]$$

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

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

(i) $V_\alpha^g(l; l, m) = 0$, and (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 hypothesized 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 9

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 characterized in Proposition 2.

First, by Lemma 6, 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 $\bar{\alpha}$ with $l(\alpha) < \varepsilon$ for all $\alpha \in (\bar{\alpha}, 1)$. 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:

$$\frac{\lambda}{\lambda + (1 - \alpha)(1 - \lambda)}g + \frac{(1 - \alpha)(1 - \lambda)}{\lambda + (1 - \alpha)(1 - \lambda)}b - c - (1 - \alpha)\frac{c - b}{1 - \beta(1 - \alpha)} = 0. \quad (A.13)$$

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. ■

APPENDIX B

Additional results relating to Chapter 2

B.1. Finite horizon

B.1.1. Perfectly informed seller

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 recognized 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 summarized in the proposition below; the proof of which is based on backward induction and appears below. 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 B.1.3, which are uniquely robust to natural restrictions on beliefs.

Proposition 13. *Suppose that there are $N > 1$ trading opportunities and that the current reputation level is $\lambda \in (0, 1)$, then there exists an equilibrium 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 $\lambda g + (1 - \lambda)b - c \geq 0$.

Proof 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:

$$W_2^b(\lambda) = \lambda g + (1 - \lambda)b - c + \beta[bV_1^b(\lambda^s) + (1 - b)V_1^b(\lambda^f)] \qquad \text{(B.1)}$$

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 $W_N^b(\lambda)$:

$$W_{N+1}^b(\lambda) = \lambda g + (1 - \lambda)b - c + \beta[bV_N^b(\lambda^s) + (1 - b)V_N^b(\lambda^f)] \qquad \text{(B.2)}$$

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

$$
\begin{aligned}
W_{N+1}^b(\lambda_N) &> \lambda_N g + (1 - \lambda_N)b - c + \beta[bV_{N-1}^b(\lambda_N^s) + (1 - b)V_N^b(\lambda_N^f)] \quad \text{(B.3)} \\
&\geq \lambda_N g + (1 - \lambda_N)b - c + \beta[bV_{N-1}^b(\lambda_N^s) + (1 - b)V_{N-1}^b(\lambda_N^f)] \\
&= V_N^b(\lambda_N) = 0
\end{aligned}
$$

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 $V_N^b(\lambda)$:

$$
\begin{aligned}
V_{N+1}^b(\lambda) &= \lambda g + (1 - \lambda)b - c + \beta[bV_N^b(\lambda^S) + (1 - b)V_N^b(\lambda^f)] \quad \text{(B.4)} \\
&> \lambda g + (1 - \lambda)b - c + \beta[bV_{N-1}^b(\lambda^S) + (1 - b)V_N^b(\lambda^f)] \\
&\geq \lambda g + (1 - \lambda)b - c + \beta[bV_{N-1}^b(\lambda^S) + (1 - b)V_{N-1}^b(\lambda^f)] \\
&= V_N^b(\lambda)
\end{aligned}
$$

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. ∎

B.1.2. Imperfect private information

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 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 below.

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 summarize 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 characterizes 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 14. *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 $m \le n$ which has the following property: If $U_m(p) > 0$ and $m > 1$, then $V_m(\lambda, p) > 0$ for all $\lambda > 0$.*

Proof. 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 $(0, 1)^2$:

$V_m(p, p) \ge U_m(p)$;

$V_m(\lambda, p)$ *is non-decreasing in λ;*

$V_m(\lambda, p)$ *is non-decreasing in p; and,*

$V_m(\lambda, p)$ *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 $\lambda g + (1 - \lambda)b - c \ge 0$, 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 $\lambda g + (1 - \lambda)b - c + \beta E[V_{n-1}(\nu, q) \mid \lambda, p] \ge 0$ 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 $d_n(\lambda)$.[1]

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:

$$\mu g + (1 - \mu)b - c + \beta E[V_{n-1}(\nu, q) \mid \mu, p_n(\lambda)] = 0. \tag{B.5}$$

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 $d_n(\lambda)$.[2] 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 hypothesized 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.

---

[1] 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.

[2] 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.

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

$$\mu g + (1 - \mu)b - c + \beta E[V_{n-1}(\nu, q) \mid \mu, p_n] \leq 0 \tag{B.6}$$

but it also follows that:

$$\mu g + (1 - \mu)b - c + \beta E[V_{n-1}(\nu, q) \mid \mu, p_n] > p_n g + (1 - p_n)b - c + \beta E[V_{n-1}(\nu, q) \mid p_n, p_n] \geq 0 \tag{B.7}$$

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

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. $\square$

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 cannot 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 by backward induction.

B.1.3. Varying the degree to which the seller can behave strategically

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 15. *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.*

Proof. 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.[3] This, together with

---

[3]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.

Assumption 2, implies that an equilibrium must be of the hypothesized 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 hypothesized 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 $\lambda \in (0,1)$. $\square$

The result is thus proved 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 hypothesized form.

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

Proposition 16. *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.*

Proof. 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 would not 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 characterizes the lower bound, h_n, of the highest range.   □

-The intuition underlying this result again rests on the self-knowledge 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.

For general α, there is a multiplicity of equilibria of the type characterized in Proposition 16.[4] Note, however, that there is no multiplicity of equilibria of this type even with

---

[4]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:

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

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

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

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

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

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

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

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. $\overline{l}$ 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

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; though we argue that as α tends to 1, the set of feasible equilibria converges to a singleton.

Proposition 17. *As α tends to 1, the set of feasible equilibria of the hypothesized form (that is as described in Lemma 15) converge to a singleton. Moreover, this is the equilibrium characterized in Proposition 3.*[5]

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 focus Chapter 2 on the equilibria characterized in Sections 2.4 and 2.5.

B.2. Price as a signalling device

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 chapter of continuing to trade despite short-term losses. In this section, we relax this assumption and modify the model presented in Section 2.2 by supposing that if the seller chooses to trade, she makes a take-it-or-leave-it offer to one of the potential buyers.

---

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.

[5]The proof is similar to the proof of Proposition 9 and so is omitted.

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 characterized in Section 2.4. 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.[6]

It is worth noting that the equilibria described in Section 2.4 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

---

[6]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.

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.

B.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 characterize 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 characterized in Section 2.4 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 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 formalized 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.[7]

Definition 18. *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 below that working with this definition and proceeding by backward induction yields a unique equilibrium outcome, which is similar to that of Proposition 13 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 19. *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 characterized in Proposition 13, 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.[8] 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

---

[7]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).

[8]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 $Y_2^b(0.42) = 0$.

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 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.

Proof of Proposition 19

We make the following claims for the value for a good and bad seller of having a reputation λ with $n \leq N$ 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

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

$$Y_n^b(\lambda) = \lambda g + (1 - \lambda)b - c + \beta b Y_{n-1}^b(\lambda^s) + \beta(1 - b)Y_{n-1}^b(\lambda^f) \qquad \text{(B.9)}$$

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.

<u>n=1</u>

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 realizes 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 20. 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.[9] We consider each of these two possibilities separately.

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

$$p - c + \beta g Y_{m-1}^g(\lambda^s) + \beta(1 - g)Y_{m-1}^g(\lambda^f) \qquad (B.10)$$

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 $Y_{m-1}^g(\lambda^s)$ and $Y_{m-1}^g(\lambda^f)$, and the future beliefs must be λ^s and λ^f for consistency with the hypothesized equilibrium strategies that both good and bad sellers always sell. Furthermore, for this to be an equilibrium, it must be true that:[10]

$$p - c + \beta b Y_{m-1}^b(\lambda^s) + \beta(1 - b)Y_{m-1}^b(\lambda^f) \geq 0 \qquad (B.11)$$

This constraint ensures that it is rational for a bad seller to sell with probability one. Since $p - c + \beta g Y_{m-1}^g(\lambda^s) + \beta(1 - g)Y_{m-1}^g(\lambda^f)$ 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:[11]

$$\underline{\lambda}_m g + (1 - \underline{\lambda}_m)b + \beta b Y_{m-1}^b(\underline{\lambda}_m^s) + \beta(1 - b)Y_{m-1}^b(\underline{\lambda}_m^f) = 0 \qquad (B.12)$$

---

[9]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.

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

[11]Note we abuse notation and $\underline{\lambda}_m$ here is different from $\underline{\lambda}_t$ which appears in Section 2.5.

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

$$p - c + \beta g Y^g_{m-1}(\mu^s) + \beta(1 - g)Y^g_{m-1}(\mu^f) \qquad \text{(B.13)}$$

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

$$p - c + \beta b Y^b_{m-1}(\mu^s) + \beta(1 - b)Y^b_{m-1}(\mu^f) = 0 \qquad \text{(B.14)}$$

Consider the problem of maximizing the expression in (B.13) subject to the constraint (B.14) and $p \leqslant \mu g + (1 - \mu)b$. Substituting for p from (B.14), the problem can be recast as the maximization of

$$\beta g Y^g_{m-1}(\mu^s) + \beta(1 - g)Y^g_{m-1}(\mu^f) - \beta b Y^b_{m-1}(\mu^s) - \beta(1 - b)Y^b_{m-1}(\mu^f) \qquad \text{(B.15)}$$

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 (B.14) 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:

$$\mu_m g + (1 - \mu_m)b + \beta b Y^b_{m-1}(\mu^s_m) + \beta(1 - b)Y^b_{m-1}(\mu^f_m) \geq 0 \qquad \text{(B.16)}$$

Moreover since $\underline{\lambda}_m$ and $p = \underline{\lambda}_m g + (1 - \underline{\lambda}_m)b$ satisfy the constraint (B.14), it follows that:

$$\underline{Y}^g_m \geq \underline{\lambda}_m g + (1 - \underline{\lambda}_m)b - c + \beta g Y^g_{m-1}(\underline{\lambda}^s_m) + \beta(1 - g)Y^g_{m-1}(\underline{\lambda}^f_m) \qquad \text{(B.17)}$$

From the assumed properties of $Y_{m-1}^g(\lambda)$, $\lambda g + (1 - \lambda)b - c + \beta b Y_{m-1}^b(\lambda^s) + \beta(1 - b)Y_{m-1}^b(\lambda^f)$ is continuous and increasing in λ and so there exists a $\widehat{\lambda}_m \geq \underline{\lambda}_m$ such that $\underline{Y}_m^g = \widehat{\lambda}_m g + (1 - \widehat{\lambda}_m)b + \beta g Y_{m-1}^g(\widehat{\lambda}_m^s) + \beta(1 - g)Y_{m-1}^g(\widehat{\lambda}_m^f)$.[12]

Finally it is a simple exercise to show that there is a CPE with $\widehat{\lambda}_m$, p_m and $\underline{Y}_m^g$ 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 hypothesized inductive properties do hold. ■

Lemma 20. $Y_m^g(\lambda) - Y_{m-1}^g(\lambda) > 0$

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 $Z_m^g(\lambda)$, such that $Z_m^g(\lambda) - Y_{m-1}^g(\lambda) > 0$. Then by definition $Y_m^g(\lambda) \geq Z_m^g(\lambda)$ 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 2.4.2 and the two period version of Proposition 3, the equilibrium strategies in this proposition and the induced value which we here denote $Z_2^g(\lambda)$ satisfy the hypothesized properties.

Inductive step

Recall $\underline{\lambda}_m$ as defined in Equation (B.12), we consider two cases separately and $\widehat{\lambda}_m$ as defined in the paragraph below Equation (B.17). 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 $Z_m^g(\lambda) = \lambda g + (1 - \lambda)b - c + \beta g Y_{m-1}^g(\lambda^s) + \beta(1 - g)Y_{m-1}^g(\lambda^f)$ for $\lambda \geq \underline{\lambda}_m$ and $Z_m^g(\lambda) = \underline{\lambda}_m g + (1 - \underline{\lambda}_m)b - c + \beta g Y_{m-1}^g(\underline{\lambda}_m^s) + \beta(1 - g)Y_{m-1}^g(\underline{\lambda}_m^f)$ otherwise.

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

---

[12] 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 $\underline{Y}_m^g > \frac{(g-c)(1-\beta^{m+1})}{1-\beta}$.

$$
\begin{aligned}
Z_m^g(\lambda) - Y_{m-1}^g(\lambda) \quad &= \lambda g + (1-\lambda)b - c + \beta g Y_{m-1}^g(\lambda^s) + \beta(1-g)Y_{m-1}^g(\lambda^f) \\
&\quad -[\lambda g + (1-\lambda)b - c + \beta g Y_{m-2}^g(\lambda^s) + \beta(1-g)Y_{m-2}^g(\lambda^f)] \\
&= \beta g[Y_{m-1}^g(\lambda^s) - Y_{m-2}^g(\lambda^s)] + \beta(1-g)[Y_{m-1}^g(\lambda^f) - Y_{m-2}^g(\lambda^f)] \\
&\quad > 0
\end{aligned}
$$

(B.18)

where the inequality follows from the hypothesized inductive property.

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

$$
\begin{aligned}
Z_m^g(\lambda) - Y_{m-1}^g(\lambda) \quad &= \underline{\lambda}_m g + (1-\underline{\lambda}_m)b - c + \beta g Y_{m-1}^g(\underline{\lambda}_m^s) + \beta(1-g)Y_{m-1}^g(\underline{\lambda}_m^f) \\
&\quad -[\lambda g + (1-\lambda)b - c + \beta g Y_{m-2}^g(\lambda^s) + \beta(1-g)Y_{m-2}^g(\lambda^f)] \\
&\quad > \lambda g + (1-\lambda)b - c + \beta g Y_{m-1}^g(\lambda^s) + \beta(1-g)Y_{m-1}^g(\lambda^f) \\
&\quad -[\lambda g + (1-\lambda)b - c + \beta g Y_{m-2}^g(\lambda^s) + \beta(1-g)Y_{m-2}^g(\lambda^f)] \\
&\quad > 0
\end{aligned}
$$

(B.19)

where the first inequality follows since $Y_{m-1}^g(\lambda)$ is increasing in λ, and the second inequality follows in analogous way to the inequality in (B.18).

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

$$
\begin{aligned}
Z_m^g(\lambda) - Y_{m-1}^g(\lambda) \quad &= \underline{\lambda}_m g + (1-\underline{\lambda}_m)b - c + \beta g Y_{m-1}^g(\underline{\lambda}_m^s) + \beta(1-g)Y_{m-1}^g(\underline{\lambda}_m^f) \\
&\quad -[\widehat{\lambda}_{m-1}g + (1-\widehat{\lambda}_{m-1})b - c + \beta g Y_{m-2}^g(\widehat{\lambda}_{m-1}^s) + \beta(1-g)Y_{m-2}^g(\widehat{\lambda}_{m-1}^f)] \\
&\quad > \widehat{\lambda}_{m-1}g + (1-\widehat{\lambda}_{m-1})b - c + \beta g Y_{m-1}^g(\widehat{\lambda}_{m-1}^s) + \beta(1-g)Y_{m-1}^g(\widehat{\lambda}_{m-1}^f) \\
&\quad -[\widehat{\lambda}_{m-1}g + (1-\widehat{\lambda}_{m-1})b - c + \beta g Y_{m-2}^g(\widehat{\lambda}_{m-1}^s) + \beta(1-g)Y_{m-2}^g(\widehat{\lambda}_{m-1}^f)] \\
&\quad > 0
\end{aligned}
$$

(B.20)

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 $Z_m^g(\lambda) = \lambda g + (1-\lambda)b - c + \beta g Y_{m-1}^g(\lambda^s) + \beta(1-g)Y_{m-1}^g(\lambda^f)$ for $\lambda \ge \widehat{\lambda}_{m-1}$ and $Z_m^g(\lambda) = \widehat{\lambda}_{m-1}g + (1-\widehat{\lambda}_{m-1})b - c + \beta g Y_{m-1}^g(\widehat{\lambda}_{m-1}^s) + \beta(1-g)Y_{m-1}^g(\widehat{\lambda}_{m-1}^f)$ otherwise.

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

$$
\begin{aligned}
Z_m^g(\lambda) - Y_{m-1}^g(\lambda) &= \lambda g + (1 - \lambda)b - c + \beta g Y_{m-1}^g(\lambda^s) + \beta(1 - g)Y_{m-1}^g(\lambda^f) \quad \text{(B.21)} \\
&\quad - [\lambda g + (1 - \lambda)b - c + \beta g Y_{m-2}^g(\lambda^s) + \beta(1 - g)Y_{m-2}^g(\lambda^f)] \\
&= \beta g [Y_{m-1}^g(\lambda^s) - Y_{m-2}^g(\lambda^s)] + \beta(1 - g)[Y_{m-1}^g(\lambda^f) - Y_{m-2}^g(\lambda^f)] \\
&> 0
\end{aligned}
$$

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

$$
\begin{aligned}
Z_m^g(\lambda) - Y_{m-1}^g(\lambda) \quad &= 1 - \widehat{\lambda}_{m-1}g + (1 - \widehat{\lambda}_{m-1})b - c \\
&\quad + \beta g Y_{m-1}^g(\widehat{\lambda}_{m-1}^s) + \beta(1 - g)Y_{m-1}^g(\widehat{\lambda}_{m-1}^f) \\
&\quad - [\widehat{\lambda}_{m-1}g + (1 - \widehat{\lambda}_{m-1})b - c \quad\quad \text{(B.22)} \\
&\quad \beta g Y_{m-2}^g(\widehat{\lambda}_{m-1}^s) + \beta(1 - g)Y_{m-2}^g(\widehat{\lambda}_{m-1}^f)] \\
&\quad\quad > 0
\end{aligned}
$$

This covers all possibilities and concludes the proof.   ∎

APPENDIX C

Variations on the central model of Chapter 3

C.1. Agents know their own types at birth

In this section, we adapt the model presented in Section 4.3 and suppose that agents know their own types at birth. Analogous to the equilibrium described in there, we construct an equilibrium in which competent agents exert effort. Specifically, we distinguish between a competent and an inept agent. In equilibrium the following occurs:

In the first period of life, a competent agent either founds her own firm, working on her own in an unoccupied location or accepts a position as an employee if offered one at sufficiently attractive terms; in either case she exerts effort in the first period of her own life. In equilibrium in either case she will succeed (since it is assumed that a competent agent who exerts effort succeeds). If she had founded her own firm, then in the second period she hires a junior (offering a wage contract that pays w and offers the firm at P) and they work together with the founder exerting effort.

Alternatively she might begin life as an employee in an established firm with a (w, P) contract, so long as such a contract offers her in equilibrium at least as much lifetime earnings as founding her own firm. She exerts effort in the first period of life. In equilibrium, the firm produces two high quality outputs and she buys the firm at the specified price P, hires her own junior offering her junior the contract (w, P) and she works together with her junior and exerts effort.

An inept agent might begin life either as an employee or founding her own firm. In either case she exerts no effort, fails and poses as a new-born in the following period, accepting employment as an employee or founding her own firm.

It is clear that an inept agent, knowing that she is inept, would never exert effort (it is costly but otherwise does not affect outcomes) and would be willing to work as an employee so long as:

$$w \geq \mu \tag{C.1}$$

For the competent agent, following the strategy outlined above yields:

$$CV_f = \mu - c + (1 + \mu - c - w + \mu P). \qquad (C.2)$$

when starting life by founding her own firm, and

$$CV_e = w - c + (1 + \mu - c - w - P + \mu P). \qquad (C.3)$$

when starting as an employee. Note that in contrast to the expressions in (3.6) and (3.7) the competent agent, knowing that she is competent knows that she will succeed.

In equilibrium, a new born agent is willing to become a junior or equivalently $CV_e \geq CV_f$. In particular, this implies that:

$$w - \mu - P \geq 0. \qquad (C.4)$$

Given that there is a scarcity of junior slots available, hiring seniors will drive down the value that the contracts deliver to employees to the point where they are indifferent, in particular this implies:

$$w = \mu + P. \qquad (C.5)$$

The remaining deviations can be categorized into a number of separate groups (i) exert effort (incentive compatibility) (ii) hiring policy and (iii) buying the firm. Each group is considered in turn.

(i) exert effort (incentive compatibility)

For a second period agent who had been successful there is an incentive compatibility constraint—which is identical in both the case that she worked alone in the first period and the case when she is now a manager—specifically, this constraint is:

$$\mu P \geq c. \qquad (C.6)$$

In the first period in both cases, that the agent works alone or as a junior, it must be worthwhile to exert effort. The corresponding incentive compatibility conditions are the following:

$$1 + \mu - c - w + \mu P \geq c \qquad (C.7)$$

and

$$1 + \mu - c - w - P + \mu P \geq c \qquad (C.8)$$

Note that by (C.6), it follows that $P \geq 0$ and so (C.7) implies (C.8).

Trivially an inept agent does not exert effort

(ii) hiring policy

Without hiring in the second period of life, the agent could not commit to effort and so the best that she could do is pose as a new-born (and earn $\mu w + (1 - \mu)\mu$) and so the conditions that a second period agent who had success when working alone in the first period does indeed prefer to hire a junior is given by:

$$\mu + 1 - c - w + \mu P \geq \mu w + (1 - \mu)\mu \tag{C.9}$$

The same condition ensures that a second period agent who had been an employee and bought the firm, prefers to hire.

Suppose that an agent failed in the first period of life, then she must be inept and a single success would prove that her junior was competent and so a sufficient condition that would ensure that she does not hire is given by:

$$\mu w + (1 - \mu)\mu \geq -w + \mu P. \tag{C.10}$$

Finally, since customers hold the off equilibrium belief that any apparent new born attempting to hire must be a failed second period agent. The conditions that ensure no second period agent posing as a new born hires is also given by (C.10). A true new born agent, whether competent or inept, would not hire when this condition holds.

(iii) Buying the firm

An employee who had succeeded would indeed buy the firm so long as this generates more value than her alternative—posing as a new born. This is the case when:

$$1 + \mu - c - w - P + \mu P \geq \mu w + (1 - \mu)\mu \tag{C.11}$$

An employee who had failed is revealed as inept, and so rather than spending $P > 0$ to remain in this location, she would rather costlessly move to another location and pose as a new born agent where she would have a higher reputation. Note that (C.11) and (C.6) imply (C.9).

Thus sufficient conditions for the strategies to form an equilibrium are (C.5), (C.6), (C.7), (C.10), and (C.11). Substituting for $w = \mu + P$ from (C.5), and $\mu = \frac{\lambda}{2-\lambda}$, these inequalities reduce to $\min\{\frac{(2-\lambda)(1-c)-\lambda}{2(2-\lambda)}, \frac{(1-2c)(2-\lambda)}{2(1-\lambda)}\} \geq P \geq \frac{c(2-\lambda)}{\lambda}$, with $(c, \lambda) \in [0, 1] \times [0, 1]$. A sufficient condition is:

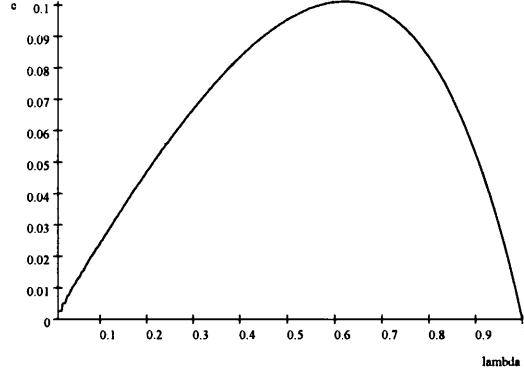$$(2 - \lambda)((1 + c)\lambda - 4c) - \lambda^2 \geq 0 \tag{C.12}$$

Figure C.1. Sufficient condition for equilibrium (know own type): c against lambda

This proves the following Proposition:

Proposition 21. *If (C.12) holds then there is a Perfect Bayesian Equilibrium in which competent agents exert effort in both periods of their careers.*

Condition (C.12) is illustrated in the Figure C.1 , where combinations of (c, λ) below the line satisfy the condition.

C.2. Case with $g < 1$, $b > 0$ and age observable

In this section, we alter the framework and suppose that a competent agent exerting effort succeeds with a probability $g < 1$ and that a competent agent who exerts no effort succeeds with probability $b > 0$, this is also the probability of success for an inept agent. We assume $g > b$. Further, it is convenient to assume that in addition to the options of working alone or working in a team, agents have the option of leaving the industry and claiming a (per-period) outside option of R.

Proposition 22. *There are parameter values and strategies which allow a Perfect Bayesian Equilibrium in which all agents exert effort in the first period of life, and agents who produce high quality when working alone in the first period or work in a firm, that produces two high quality products, as a junior exert effort in period 2 of their lives.*

In the first period of life an agent either founds her own firm, working on her own in an unoccupied location or as a junior for an agent currently in the second period of life; in either case she exerts effort in the first period of her own life. If she founds her

own firm and fails in the first period, then she leaves the industry and claims her outside option R. If she succeeds in the first period of life after founding her own firm, then in the second period she hires a junior (offering a wage contract that pays w_1 and w_2^m following a promotion) and they work together with the founder exerting effort. In the final ("retirement") period, she promotes a junior to manager (and pays the manager's fee w_2^m) if the firm as a whole produced two high quality units in the previous period and otherwise fires the junior.

Alternatively an agent might begin life as an employee in an established firm with a (w_1, w_2^m) contract. She exerts effort in the first period of life. If the firm as a whole produced one or no high quality outputs, she is fired and leaves the industry, earning her outside option R. If the firm produced two high quality outputs then she will be retained as manager and paid w_2^m, in which case she must hire her own junior offering her junior the contract (w_1, w_2^m); in this case she works together with her junior and exerts effort and at the end of the period gains control of the firm. In the final period, she promotes her junior to manager if the firm as a whole produced to high quality units in the pervious period and otherwise fires the junior.

The public and all agents believe that these are the equilibrium strategies and prices are set appropriately, with the industry capturing the full consumer surplus. Thus the price of the output of a new-born agent working alone is $p(\lambda) := \lambda g + (1 - \lambda) b$ since there is a probability λ that she is competent (and in equilibrium she is expected to exert effort and generate high quality output with probability g) and otherwise inept. To simplify notation, let the expected probability of success from a new-born agent when exerting effort be $n = p(\lambda)$; let $\lambda^s = \frac{\lambda g}{p}$ be the belief that an agent seen to have had a success is competent and let $s = p(\lambda^s)$ be the probability that an agent who had previously had a success has another success when exerting effort. With this notation, the behaviour along possible equilibrium paths, as described in the paragraph above, is illustrated in the figure below.

To complete the characterization of the equilibrium, one must also specify beliefs concerning off-equilibrium actions:

- if a senior and junior were both successful in the previous period and the senior did not promote the junior to manager but instead sought to retain her as a lone worker, then the un-promoted junior would have no concern for the future and so would not exert effort and the revenue generated would be b. We impose that in this case the senior would have to pay the un-promoted junior at least
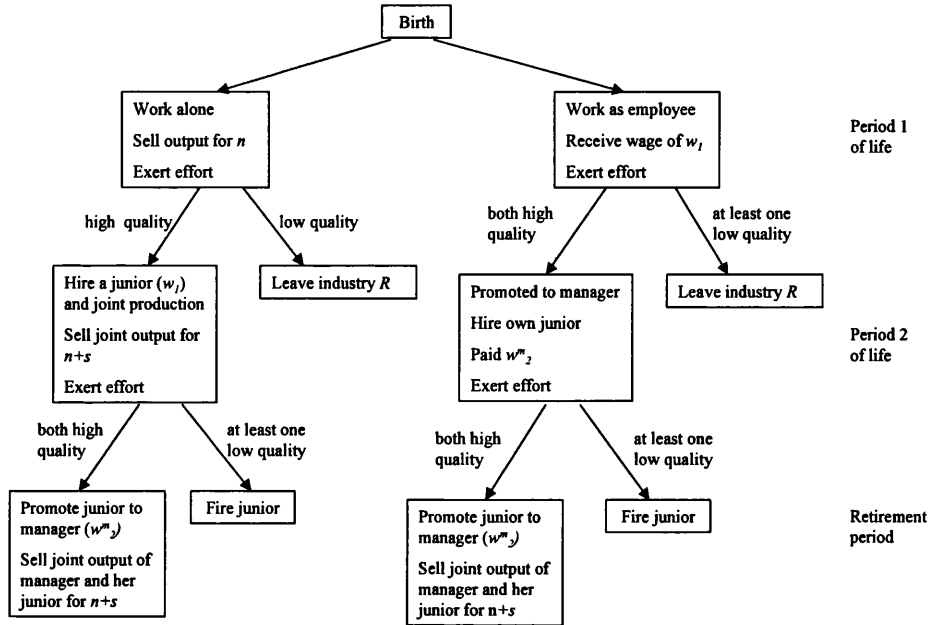
Figure C.2. Behaviour on possible equilibrium paths

her outside option R and characterize this condition explicitly below, calling it the "honest senior" condition (C.26);

- if the senior hires a junior as manager following failure, we assume in this section that the newly promoted manager will not exert effort;

- when an agent who had been successful does not hire, again there would be no possibility of retirement period revenues or wages and so the agent would not exert effort in the current period, as there are no future revenues for such effort to affect; thus this deviation is equivalent to the second period individual rationality conditions discussed below;

- when an agent who failed does hire, then we assume that the off-equilibrium beliefs are that this senior will not exert effort;

- if an agent does not offer the contract (w_1, w_2^m) then any potential juniors who will not have seen her history assume that she has previously failed; and finally,

- if an agent who is not expected to hire does so, she must offer the contract (w_1, w_2^m) but she is not expected to exert effort by customers and so does not raise the same revenues that a successful agent would; however, in this case

customers' updating on the basis of output will be different and specifically one high quality output may be sufficient to win the junior a promotion—such a deviation is considered explicitly below in Equation (C.24) and a new born agent deviating by seeking to hire in the first period is addressed by (C.25).

Note that the expected equilibrium value of beginning life by founding a firm (starting solo) is:

$$V_s = n - c + n[s + n - c - w_1 + sn(s + n - w_2^m)] + (1 - n)R \qquad (C.13)$$

This expression is derived as follows. In the first period the agent sells her service at a price n and exerts effort; if she fails, which occurs with probability $(1 - n)$, then in the second period she leaves the industry and so in the second period earns R. If she succeeds in the first period (which occurs with probability n) then it is known that she must be competent with probability λ^s, she hires a junior and sells their joint output at its expected value $s + n$, given that she exerts effort as does a new born agent, and she must also pay the junior w_1. Finally in the third period, if a senior and her junior succeeded in the second period, which occurs with probability sn, then again it must be that the junior is competent with probability λ^s and so when promoted to manager generates a revenue of $s + n$ (which is the expected value of the manager and her junior's output) but must be paid w_2^m. (We assume that the manager hires and pays her own junior but her wage is contingent on doing so). If a senior and junior fail in the second period, then in the third period the junior is fired, and the senior gains no revenue.

Similarly, the expected value for a new born agent who begins by working as a junior in joint production is given by the following expression.

$$V_j = w_1 - c + sn[w_2^m - c - w_1 + sn(s + n - w_2^m)] + (1 - sn)R \qquad (C.14)$$

The proof proceeds by construction; we characterize sufficient conditions for the parameters which ensure that the strategies described are indeed equilibrium strategies and verify that there are parameter values which satisfy all the relevant constraints.

First, we suppose that a junior refuses any contract other than (w_1, w_2^m). In addition it must be that the a new agent is willing to become a junior or equivalently $V_j \geq V_f$ and given the scarcity of such opportunities, the employer would be in a position to ensure that $V_j = V_s$

$$V_j = V_s . \qquad (C.15)$$

The remaining constraints fall into a number of categories. Specifically, these categories can be summarized under the headings (i) exit or remain in industry (individual rationality) (ii) exert effort (incentive compatibility) (iii) hiring policy and (iv) promotion policy.

The individual rationality constraints ensure that an agent prefers to stick to the strategy described to leaving the industry, either at the beginning of life or later when the strategy requires staying (for example following a high quality output when working alone in the first period). Individual rationality also requires that when the strategy requires that an agent leaves the industry, she prefers to do so than to remain within it. The first and second period incentive compatibility conditions ensure that an agent prefers to exert effort when the strategy requires her to do so.

The conditions on hiring policy are first that an agent hires in period 2 when in period 1 she had been in a "successful" firm, that is one that produced two high quality goods if she was there as a junior or where she produced one high quality good in a firm that she had founded on her own. In addition, we characterize sufficient conditions that ensure that an agent who had not been in a successful firm in period 1 does not hire in period 2, and that no agent hires in period 1.

Finally, on promotion policy the value of promoting a junior to manager in a successful firm must be greater than either the value of retaining the junior or firing her and a manager must prefer to fire a junior in an unsuccessful than retain or promote her.

Proof. The remaining deviations are categorized into the following groups (i) exit or remain in industry (individual rationality) (ii) exert effort (incentive compatibility) (iii) hiring policy and (iv) promotion policy. Each group is considered in turn and then we verify that there are indeed parameters at which all the requisite constraints are satisfied.

(i) exit or remain in industry (individual rationality)

It must be the case that pursuing these strategies from birth is preferred to exercising the outside option:

$$V_s \geq 2R \tag{C.16}$$

and $V_j \geq 2R$ or equivalently

$$w_1 - c + sn[w_2^m - c - w_1 + sn(s + n - w_2^m) - R] - R \geq 0 . \tag{C.17}$$

Note that (C.15) makes one of (C.17) and (C.16) redundant.

For a second period agent, the corresponding individual rationality constraints (the constraint that she wishes to remain in the industry and stick with the equilibrium

strategies) depend on whether the agent worked alone in a firm she founded in her first period of life or was successful as a junior in a team that produced two high quality outputs. The conditions are given respectively by:

$$s + n - c - w_1 + sn(s + n - w_2^m) \geq R , \tag{C.18}$$

$$w_2^m - c - w_1 + sn(s + n - w_2^m) \geq R . \tag{C.19}$$

Note that in (C.19) it is crucial that the agent cannot leave with her reputation intact—this follows from the assumption of informationally separate locations and that control of locations is determined by seniority.[1]

For an agent who failed when working alone or as a junior, she prefers to leave the industry so long as

$$R \geq b. \tag{C.20}$$

Note that an agent who failed would be unable to hire a junior (see (C.24) below) and so with only one period of working life remaining and no reputational concerns would be unable to commit to exerting effort and so would generate only b in revenue if she remained in the industry.

(ii) exert effort (incentive compatibility)

For a second period agent who had been successful there is an incentive compatibility constraint—which is identical in both the case that she worked alone in the first period and the case when she is now a manager—specifically, this constraint is as follows:

$$s + n - w_2^m - \frac{c}{sn} \geq 0 . \tag{C.21}$$

In the first period in both cases, that the agent work as a founder or as a junior, it must be worthwhile to exert effort. The corresponding incentive compatibility conditions are the following:

$$s + n - c - w_1 + sn(s + n - w_2^m) - R - \frac{c}{n} \geq 0 , \tag{C.22}$$

and

$$w_2^m - c - w_1 + sn(s + n - w_2^m) - R - \frac{c}{sn} \geq 0 . \tag{C.23}$$

---

[1]In the context of partnerships, restrictive covenants or bonding might be more plausible constraints which would allow a senior to claim at least some of the revenue attributable to a junior who succeeded, such mechanisms would operate in an informationally unified market.

Note that these two equations imply (C.18) and (C.19) respectively and since $-\frac{c}{n} \geq$ $-\frac{c}{sn}$ and by (C.21) $s + n \geq w_2^m$ it follows that (C.23) and (C.21) imply (C.22).

(iii) hiring policy

Without hiring in the second period of life, the agent could not commit to effort and so the conditions that a second period agent who had success when working alone in the first period does indeed prefer to hire a junior is given by:

$$s + n - c - w_1 + sn(s + n - w_2^m) \geq R$$

this is precisely (C.18) and by (C.20) which states $R > b$, this is the most profitable deviation (that is it is more profitable to leave the industry if not hiring rather than remaining in the industry and earning b). Similarly the corresponding condition that a manager prefers to hire is implied by (C.19) and (C.20).

Suppose that an agent failed in the first period of life, imposing the off-equilibrium belief that if she were to hire she would exert no effort, then a sufficient condition that would ensure that she does not hire is given by:

$$R \geq b + n - w_1 + (b + n)(s + n - w_2^m). \tag{C.24}$$

This condition is sufficient to ensure that such an agent does not hire but is not necessary. Given that the senior does not exert effort the first period revenue is $b + n - w_1$; if two successes ensue then the public believe that the junior is competent with probability λ^s and if observing one success this belief would be

$$\frac{\lambda g(1 - b) + \lambda(1 - g)b}{\lambda g(1 - b) + \lambda(1 - g)b + 2(1 - \lambda)b(1 - b)} < \lambda^s.$$

Consider a the value earned by an agent (who had previously failed) and hires in the second period of her life, when customers are over-optimistic in the sense that they believed that a junior who worked for such a senior was competent with probability λ^s when between them the senior and junior managed at least one success. This is the right hand side of (C.24) and since the condition states that even under these unrealistically favorable conditions an agent who had previously failed would still not want hire, this is a sufficient condition to ensure that such an agent does not hire. Note that in (C.24) implicitly it is assumed that customers in this off-equilibrium path believe that the senior would exert no effort. We have discussed this this off-equilibrium belief above.

Finally, the condition that a new born agent prefers not to hire is given by:

$$V_{joint} = V_{solo} \geq 2n - w_1 - c + n^2[s + n - c - w_1 + s + n - w_2^m + sn(s + n - w_2^m)] + (1 - n^2)R \ .$$
$$(C.25)$$

In this last expression we suppose that the new born agent exerts effort and that following one success the updated belief $\frac{1}{2}(\lambda^s + \frac{\lambda(1-g)}{\lambda(1-g)+(1-\lambda)(1-b)})$ is relatively small so that an agent with such a reputation would prefer to exercise her outside option.

(iv) Promotion policy

First, a senior must be honest, that is she must prefer to promote to manager a junior who had been successful. This conditions is given by:[2]

$$s + n - w_2^m \geq R \ .$$
$$(C.26)$$

In addition for the equilibrium characterized a senior in a firm with either one or two failures should not want to promote. This might be the case for example if the off-equilibrium beliefs in this case were that the promoted senior would exert no effort—then the relevant condition no promotion condition would be:

$$0 \geq -w_2^m + b + n$$
$$(C.27)$$

Verification

Thus sufficient conditions for an equilibrium are (C.15), (C.16), (C.20), (C.21), (C.23), (C.24), (C.25), (C.26) and (C.27). It can be easily verified that there are indeed parameter values and corresponding wages (w_1, w_2^m) for which these conditions are satisfied. For example they are satisfied when $\lambda = 0.45$, $g = 0.9$, $b = R = 0.1$, $c = 0.05$ and $w_1 = 0.565$ and $w_2^m = 1.124$. $\qquad\square$

C.3. A finite horizon up-or-out effort inducing equilibrium

First a three period career concern model is introduced in which reputational considerations lead to effort in the first period but not in the second of life for an agent working alone. We characterize restrictions which ensure that this is the case. Next it is shown that an agent who had been successful in the first period of life can commit to exerting effort in the second period of her life by hiring a new born agent of uncertain

---

[2]Note that if the senior attempted to retain the junior without promoting her, then the retained junior would raise no more than b in revenue which is less than her outside option by (C.20), which is the very least that she must be paid. Thus the senior would prefer to fire the junior rather than retain her without promotion.

quality, even though that agent will have no opportunity to hire her own agent and will only exert effort in the first period of her own life.

It is supposed that there are two types of agent. A good agent and a bad agent. A good agent differs from a bad agent in two ways. First she is more likely to be successful than a bad agent even when she exerts no effort and second, effort has a greater impact on her chances of generating success. Specifically suppose that a bad agent always fails whether exerting effort or not and a good agent succeeds with probability g with no effort and with probability $g + e$ when exerting effort. Exerting effort costs c. Initially the agent and labour market share the prior that the agent is good with probability λ.

Note that following any successes, all agents know that the agent must be a good type. Inferences following failure depend on whether or not it is believed that the agent exerted effort or not. Specifically let $\lambda^f(n)$ denote the belief that the agent is good given one failure when it is believed that the agent exerted no effort; $\lambda^{ff}(e, n)$ for the belief that the agent is good given two observed failures when it is believed that the agent exerted effort in the first period but not in the second; other beliefs are denoted similarly and can be written explicitly as follows

$$\lambda^f(n) = \frac{\lambda(1-g)}{\lambda(1-g)+1-\lambda} = \frac{\lambda(1-g)}{1-\lambda g}$$

$$\lambda^f(e) = \frac{\lambda(1-g-e)}{1-\lambda g-\lambda e}$$

$$\lambda^{ff}(n, n) = \frac{\lambda(1-g)^2}{\lambda(1-g)^2+1-\lambda}$$

$$\lambda^{ff}(e, n) = \lambda^{ff}(n, e) = \frac{\lambda(1-g)(1-g-e)}{\lambda(1-g)(1-g-e)+1-\lambda}$$

$$\lambda^{ff}(e, e) = \frac{\lambda(1-g-e)^2}{\lambda(1-g-e)^2+1-\lambda}$$

In addition, suppose there is an outside option R which can be invoked in any period.

First, we characterize conditions which ensure that when working alone an agent would exert effort in the first period but not in the second and would stay in the industry following a first period success but leave it otherwise.

First the condition that ensures that she would stay following success in the first period is:

$$g \geq R. \tag{C.28}$$

A condition which ensures that the she would leave the industry following two failures is that:

$$R \geq \lambda^{ff}(n, n)g. \tag{C.29}$$

Note that $\lambda^{ff}(n,n) > \lambda^{ff}(n,e) = \lambda^{ff}(e,n) \geq \lambda^{ff}(e,e)$ and so (C.29) is sufficient to ensure that the agent leaves the industry following two failures whatever the beliefs about her having exerted effort in the past.

Conditions which ensure that the agent leaves the industry following one failure and claims the outside option whether exerting effort or not are given respectively by:[3]

$$2R \geq \lambda^f(n)(g+e)(1+g) - c + (1 - \lambda^f(n)(g+e))R, \text{ and} \qquad \text{(C.30)}$$

$$2R \geq \lambda^f(n)g + \lambda^f(n)g^2 + (1 - \lambda^f(n)g)R. \qquad \text{(C.31)}$$

The condition which ensures that the agent prefers to stay in the industry and exert effort in the first period to leaving the industry immediately is

$$\lambda(g+e)(1+2g) - c + (1 - \lambda(g+e))2R \geq 3R. \qquad \text{(C.32)}$$

Finally the condition that states that she prefers to stay in the industry and exert effort in the first period to deviating and not exerting effort (even though effort is anticipated) is as follows:

$$\lambda(g+e)(1+2g) - c + (1 - \lambda(g+e))2R > \lambda(g+e) + 2\lambda g^2 + (1 - \lambda g)2R \qquad \text{(C.33)}$$

Thus inequalities (C.29)-(C.33) characterize parameters for which it is a PBE that an agent exerts effort in the first period of life and then stays in the industry following a success and leaves the industry following a failure.

Now suppose that the agent in the second period of her life has the opportunity to hire a new born agent though this agent cannot then go on to hire her own agent and show that this can be an equilibrium. The possible outcomes in this situation are summarized in the figure below.

An agent who works alone throughout life, supposing inequalities (C.29)-(C.33) hold, would earn

$$S = \lambda(g+e)(1+2g) - c + (1 - \lambda(g+e))2R. \qquad \text{(C.34)}$$

---

[3]Note that since $\lambda^f(n) > \lambda^f(e)$ and $g > R$ these conditions (which state that following one failure the agent leaves the industry given that she knows that she exerted no effort in the first period), it follows that the conditions also ensure that the agent leaves the industry following a failure when it she exerted effort in the first period.
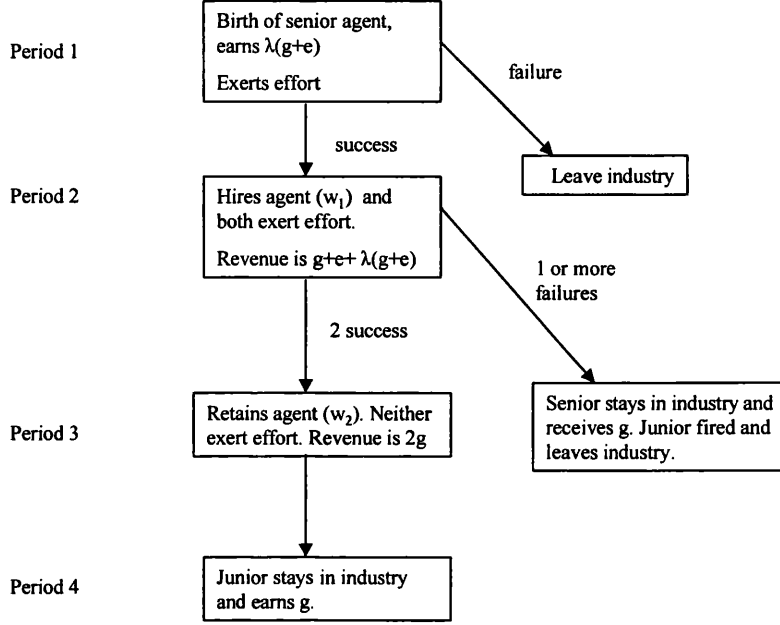
Figure C.3. Finite model: equilibrium path behaviour

Following a first period success the continuation value when working alone is $2g$. The continuation value if hiring with the contract (w_1, w_2) is instead

$H = 2g + e - c + \lambda(g + e) - w_1 + \lambda(g + e)^2(g + e - w_2)$. Thus the senior will hire in her second period following success so long as

$$2g + e - c + \lambda(g + e) - w_1 + \lambda(g + e)^2(g + e - w_2) \geq 2g \qquad (C.35)$$

She will not deviate by not exerting effort so long as the continuation value by hiring and exerting effort H is greater than the value by deviating and not exerting effort which is $2g + e + \lambda(g + e) - w_1 + \lambda(g + e)g(g + e - w_2)$. Thus the senior will not deviate by not exerting effort so long as this value is less than or equal to H, or equivalently:

$$e(g + e - w_2)\lambda(g + e) \geq c. \qquad (C.36)$$

She behaves honestly in retaining a junior following two success so long as:

$$g \geq w_2. \qquad (C.37)$$

Finally the junior agent must be willing to be hired and to exert effort, and stay following a success. This leads to three further conditions, specifically the value of being hired should be no less than the value to working alone:

$$w_1 - c + \lambda(g + e)^2(w_2 + g) + (1 - \lambda(g + e)^2)R \geq S, \qquad \text{(C.38)}$$

and the junior should be willing to exert effort in the first period of her life (corresponding to period 2 of the time path) so that the left hand side of (C.38) must be greater than or equal to $w_1 + \lambda(g + e)g(w_2 + g) + (1 - \lambda(g + e)g)R$, or equivalently

$$\lambda(g + e)e(w_2 + g - R) \geq c. \qquad \text{(C.39)}$$

The last condition is that following a success, the junior should prefer to stay in the firm to leaving the firm (and so also the industry) that is:

$$w_2 + g \geq 2R \qquad \text{(C.40)}$$

Finally, having characterized the relevant inequalities, it remains to show that there are indeed parameter values for which all the relevant inequalities hold. It can be easily verified that (C.29)-(C.33) and (C.35)-(C.40) are all satisfied, for example, when $R = 0.6$, $g = 0.5$, $e = 0.3$, $c = 0.01$, $\lambda = 0.8$, $w_1 = 1$, and $w_2 = 0.55$.

C.4. $g < 1$ and $b > 0$ age non observable

Again, as in Section 3.4, we suppose that agents do not know their own types at birth but, in contrast to the framework considered in Section 3.4, it is supposed that $g < 1$ and $b > 0$. We characterize conditions for an up-or-out effort inducing equilibrium in which both new born agents and successful agents in the second period of life exert effort. The equilibrium follows the structure characterized in Section 3.4 but in which agents who fail (either on their own or as part of a firm that suffered at least one failure) pose as new born agents, are willing to work as juniors when given the opportunity and exert no effort. For convenience, take $b = 0$.

The value of a new born who starts life working alone is given by:

$$HV_{solo} = \mu g - c + \lambda g(\mu g + g - c - w_1 + \mu g^2(g + \mu g - w_2^m)) + (1 - \lambda g)\mu g. \qquad \text{(C.41)}$$

Note that this expression is built up as follows. On observing an agent who appears to be new-born working alone, customers believe that she is a competent new born agent with probability μ, we characterize μ below but note that $\mu < \lambda$ since among the pool of those agents who appear to be new born there will be some second-period agents who are hiding their histories. The new born agent, who exerts effort and knows that she is new born expects success with probability λg and in this case she hires a another agent as a junior (note that this may be a failed agent pretending to be new born and so the probability that junior will be new born and competent is μg) and so can charge $\mu g + g$ for the joint service and in case of further success she promotes the junior, paying her the manager's fee and receiving the appropriate revenue.[4] Following a failure in the first period, the agent can pose as a new-born agent and so receive the revenue μg.

The value for a new born agent who works as an employee is given by the following expression:

$$HV_{joint} = w_1 - c + \lambda g^2 (w_2^m - c - w_1 + \mu g^2 (g + \mu g - w_2^m)) + (1 - \lambda g^2)\mu g. \quad \text{(C.42)}$$

As before we impose:

$$HV_{solo} = HV_{joint} \quad \text{(C.43)}$$

In addition to the conditions which are similar to those considered in Proposition 11. There will be new conditions to sustain an equilibrium, in particular, that a second period agent who is hiding her history prefers to work alone rather than be hired as a junior. Specifically, this condition is given by:

$$w_1 \geq \mu g \quad \text{(C.44)}$$

Say that there is a stable proportion of agents α in each generation who succeed in the first period of life hire new-borns is given by α. Then $\mu = \frac{\lambda}{2-\alpha}$ and $\alpha = \mu g(2-2\alpha) + \alpha \mu g^2$, so $\alpha = \frac{\lambda}{2-\alpha} g(2 - 2\alpha) + \alpha \frac{\lambda}{2-\alpha} g^2$. First note that

$$\left(\lambda^2 g^4 - 4\lambda^2 g^3 - 4\lambda g^2 + 4\lambda^2 g^2 + 4\right) = 4(1 - \lambda g^2) + 4\lambda^2 g^2(1 - g) + \lambda^2 g^4 > 0$$

Secondly note that there is a sensible solution only when $0 \leq \alpha \leq 1$.

---

[4]If there are two successes and the junior was a failed agent, she can not claim the promotion, but retires.

Write $f(\alpha) = \alpha\lambda g^2 + (2 - 2\alpha)\lambda g - \alpha(2 - \alpha)$. Then $f(0) = 2\lambda g > 0$ and $f(1) = \lambda g^2 - 1 < 0$

so there must be at least one solution for the quadratic equation $f(\alpha) = 0$ in the range, the question is, is there exactly one and which one?

Well consider

$$\lambda g(1 - \tfrac{g}{2}) + 1 + \tfrac{1}{2}\sqrt{(\lambda^2 g^4 - 4\lambda^2 g^3 - 4\lambda g^2 + 4\lambda^2 g^2 + 4)}$$

This is > 1 so the root in the range $(0,1)$ must be the negative root. That is the equation has exactly one root in the range $[0,1]$ and so

$$\alpha = 1 + \lambda g - \frac{1}{2}\lambda g^2 - \frac{1}{2}\sqrt{(4 - 4\lambda g^2 + 4\lambda^2 g^2 - 4\lambda^2 g^3 + \lambda^2 g^4)}.$$

Corresponding to Proposition 11, the remaining constraints can be categorized in groups as follows.

(i) work alone

Now there is no opportunity to leave the industry, but an agent can always choose to work alone in both periods. It must be the case that pursuing the hypothesized up-or-out equilibrium strategies from birth is preferred to working alone:

$$V_{solo} \geq 2\mu g \tag{C.45}$$

and $V_{joint} \geq 2\mu g$. Note that (C.43) makes this latter condition redundant.

(ii) exert effort (incentive compatibility)

For a second period agent who had been successful there is an incentive compatibility constraint—which is identical in both the case that she worked alone in the first period and the case when she is now a manager—specifically, this constraint is as follows:

$$g + \mu g - w_2^m - \frac{c}{\mu g^2} \geq 0 . \tag{C.46}$$

In the first period in both cases, that the agent works alone or as a junior, it must be worthwhile to exert effort. The corresponding incentive compatibility conditions are the following:

$$g - c - w_1 + \mu g^2(g + \mu g - w_2^m) - \frac{c}{\lambda g} \geq 0, \tag{C.47}$$

and

$$w_2^m - c - w_1 + \mu g^2(g + \mu g - w_2^m) - \mu g - \frac{c}{\lambda g^2} \geq 0 . \tag{C.48}$$

(iii) hiring policy

Without hiring in the second period of life, the agent could not commit to effort and so the conditions that a second period agent who had success when working alone in the first period does indeed prefer to hire a junior is given by:

$$g - c - w_1 + \mu g^2 (g + \mu g - w_2^m) \geq 0$$

which is implied by (C.47). Similarly the corresponding condition that a manager prefers to hire is implied by (C.48).

Suppose that an agent failed in the first period of life, imposing the off-equilibrium belief that if she were to hire she would exert no effort, then a sufficient condition that would ensure that she does not hire is given by:

$$0 \geq -w_1 + \mu g (g + \mu g - w_2^m). \tag{C.49}$$

Note that in this case, since $b = 0$, a single success demonstrates to customers that the junior must be competent.

Finally, the condition that a new born agent prefers not to hire is given by:

$$HV_{solo} \geq \lambda g + \mu g - w_1 - c + \lambda \mu g^2 [\mu g + g - c - w_1 + \mu g + g - w_2^m + \mu g^2 (g + \mu g - w_2^m)] + (1 - \lambda \mu g^2) \mu g \ . \tag{C.50}$$

In this last expression we suppose that the new born agent exerts effort and that following one success the updated belief is relatively small so that an agent with such a reputation would prefer to hide her history and pose as a new born.

(iv) **Promotion policy**

First, a senior honest, that is she must prefer to promote to manager a junior who had been successful. This conditions is given by:

$$g + \mu g - w_2^m \geq 0 \ . \tag{C.51}$$

In addition for the equilibrium characterized a senior in a firm with either one or two failures should not want to promote. This might be the case for example if the off-equilibrium beliefs in this case were that the promoted senior would exert no effort—then the relevant condition no promotion condition would be:

$$0 \geq -w_2^m + \mu g \tag{C.52}$$

Verification

It can readily be verified that all the relevant conditions are satisfied when $\lambda = g = 0.9$, $c = 0.05$, $w_1 = 0.74$ and $w_2^m = 1.454$.

APPENDIX D

Proofs for Chapter 4

Proposition 23. *Equation (4.4) has a unique solution in the range $[0, 1]$.*

Proof. Existence of a solution in the range $[0, 1]$

Consider $f(e) = \alpha + \frac{p(1-\alpha)}{p+r(e(1-\alpha-\beta)+\alpha)} - (\alpha - \frac{q\alpha}{q+r((1-e)(1-\alpha-\beta)+\beta)}) - e$.

$f(0) > 0$ and $f(1) < \alpha + (1-\alpha)\frac{p}{p+r(1-\beta)} - \alpha\frac{r\beta}{q+r\beta} - 1 < 0$, given that $f(e)$ is a continuos function, there exist $e^* \in (0, 1)$ such that $f(e^*) = 0$.

Thus Equation (4.4) has at least one solution in $(0, 1)$.

Uniqueness

First note that $f(e)$ is a convex function,

This follows since $\frac{d^2 f(e)}{de^2} = \frac{d^2 S}{de^2} - \frac{d^2 F}{de^2}$

Examining each term separately:

$\frac{dS}{de} = -(1-\alpha)\frac{pr(1-\alpha-\beta)}{[p+r(e(1-\alpha-\beta)+\alpha)]^2} < 0$ and $\frac{d^2 S}{de^2} = 2(1-\alpha)\frac{pr^2(1-\alpha-\beta)^2}{[p+r(e(1-\alpha-\beta)+\alpha)]^3} > 0$; and

$\frac{dF}{de} = -\alpha\frac{qr(1-\alpha-\beta)}{[q+r((1-e)(1-\alpha-\beta)+\beta)]^2} < 0$ and $\frac{d^2 F}{de^2} = -2\alpha\frac{q^2(1-\alpha-\beta)^2}{[q+r((1-e)(1-\alpha-\beta)+\beta)]^3} < 0$.

Thus overall $\frac{d^2 f(e)}{de^2} > 0$ and so $f(e)$ is a convex function.

Suppose that $f(e) = 0$ has more than one solution in $(0, 1)$. Let e^* be the first such root. Given that $f(0) > 0$, $f'(e^*) < 0$, but then $f(e)$ cannot have another root, e^{**}, in $(0, 1)$. Given the continuity of $f(e)$, the second root e^{**} must have a positive derivative $f'(e^{**}) > 0$, but this is not possible, since $f(e)$ is a convex function and this would imply that $f(e)$ is increasing for $e > e^{**}$, which contradicts $f(1) < 0$. $\qquad\square$

Lemma 24. $sign(\frac{de^*}{d\alpha}) = sign(-p\frac{p+r(1-\beta e^*)}{(p+r(e^*(1-\alpha-\beta)+\alpha))^2} + q\frac{q+r(1-e^*)(1-\beta)+r\beta}{(q+r((1-e^*)(1-\alpha-\beta)+\beta))^2})$

Proof. Using the implicit function theorem over $f(e^*) = 0$ (Equation (4.4), we obtain

$$\frac{\partial e^*}{\partial \alpha} = -\frac{\frac{\partial f(e^*)}{\partial \alpha}}{\frac{\partial f(e^*)}{\partial e^*}}$$

Given that for the previous proof, $\frac{\partial f(e^*)}{\partial e^*} < 0$, the sign of $\frac{de^*}{d\alpha}$ is equal to the sign of $\frac{\partial f(e^*)}{\partial \alpha}$. Taking the derivative over $f(e^*)$ with respect to α yields:

$$\frac{\partial f(e^*)}{\partial \alpha} = -\frac{p}{p+r(e^*(1-\alpha-\beta)+\alpha)} - \frac{pr(1-\alpha)(1-e^*)}{(p+r(e^*(1-\alpha-\beta)+\alpha))^2} + \frac{q}{q+r((1-e^*)(1-\alpha-\beta)+\beta)}$$
$$+\frac{qr\alpha(1-e^*)}{(q+r((1-e^*)(1-\alpha-\beta)+\beta))^2} - \frac{q\alpha(-r(1-e^*)-r(1-\alpha-\beta)\frac{de^*}{d\alpha})}{(q+r((1-e^*)(1-\alpha-\beta)+\beta))^2},$$

or equivalently:

$$sign(\tfrac{de^*}{d\alpha}) = sign(\tfrac{\partial f(e^*)}{\partial \alpha}) = sign(-p\frac{p+r(1-\beta e^*)}{(p+r(e^*(1-\alpha-\beta)+\alpha))^2} + q\frac{q+r(1-e^*)(1-\beta)+r\beta}{(q+r((1-e^*)(1-\alpha-\beta)+\beta))^2}).$$

$\square$

D.1. Proof of Proposition 12

Proof. Taking the derivative of e^* with respect to α from Equation (4.4) yields:

$$\frac{de^*}{d\alpha}(1 + \frac{pr(1-\alpha)(1-\alpha-\beta)}{(p+r(e^*(1-\alpha-\beta)+\alpha))^2} + \frac{(1-p-r)r\alpha(1-\alpha-\beta)}{(1-p-r+r((1-e^*)(1-\alpha-\beta)+\beta))^2}) =$$
$$-\frac{p}{p+r(e^*(1-\alpha-\beta)+\alpha)} - \frac{pr(1-\alpha)(1-e^*)}{(p+r(e^*(1-\alpha-\beta)+\alpha))^2}$$
$$+\frac{1-p-r}{1-p-r+r((1-e^*)(1-\alpha-\beta)+\beta)} + \frac{(1-p-r)r\alpha(1-e^*)}{(1-p-r+r((1-e^*)(1-\alpha-\beta)+\beta))^2}$$

Write $A = \frac{de^*}{d\alpha}(1 + \frac{pr(1-\alpha)(1-\alpha-\beta)}{(p+r(e^*(1-\alpha-\beta)+\alpha))^2} + \frac{(1-p-r)r\alpha(1-\alpha-\beta)}{(1-p-r+r((1-e^*)(1-\alpha-\beta)+\beta))^2})$ and

$$B = -\frac{p}{p+r(e^*(1-\alpha-\beta)+\alpha)} - \frac{pr(1-\alpha)(1-e^*)}{(p+r(e^*(1-\alpha-\beta)+\alpha))^2}$$
$$+\frac{1-p-r}{1-p-r+r((1-e^*)(1-\alpha-\beta)+\beta)} + \frac{(1-p-r)r\alpha(1-e^*)}{(1-p-r+r((1-e^*)(1-\alpha-\beta)+\beta))^2} .$$

Taking the total derivative with respect to p yields

$$\frac{dA}{dp} = \frac{d^2e^*}{d\alpha dp}(1 + \frac{pr(1-\alpha)(1-\alpha-\beta)}{(p+r(e^*(1-\alpha-\beta)+\alpha))^2} + \frac{(1-p-r)r\alpha(1-\alpha-\beta)}{(1-p-r+r((1-e^*)(1-\alpha-\beta)+\beta))^2})$$
$$+\frac{de^*}{d\alpha}(\frac{r(1-\alpha)(1-\alpha-\beta)}{(p+r(e^*(1-\alpha-\beta)+\alpha))^2} - 2\frac{pr(1-\alpha)(1-\alpha-\beta)}{(p+r(e^*(1-\alpha-\beta)+\alpha))^3}(1+r(1-\alpha-\beta)\tfrac{de^*}{dp}))$$
$$-\frac{(1-\alpha-\beta)}{(1-p-r+r((1-e^*)(1-\alpha-\beta)+\beta))^2} - 2\frac{(1-p-r)r\alpha(1-\alpha-\beta)}{(1-p-r+r((1-e^*)(1-\alpha-\beta)+\beta))^3}(1+r(1-\alpha-\beta)\tfrac{de^*}{dp}))$$

and

$$\frac{dB}{dp} = -\frac{1}{p+r(e^*(1-\alpha-\beta)+\alpha)} + \frac{p}{(p+r(e^*(1-\alpha-\beta)+\alpha))^2}(1+r(1-\alpha-\beta)\tfrac{de^*}{dp})$$
$$-\frac{r(1-\alpha)(1-e^*)}{(p+r(e^*(1-\alpha-\beta)+\alpha))^2} + 2\frac{pr(1-\alpha)(1-e^*)}{(p+r(e^*(1-\alpha-\beta)+\alpha))^3}(1+r(1-\alpha-\beta)\tfrac{de^*}{dp})$$
$$-\frac{1}{1-p-r+r((1-e^*)(1-\alpha-\beta)+\beta)} + \frac{1-p-r}{(1-p-r+r((1-e^*)(1-\alpha-\beta)+\beta))^2}(1+r(1-\alpha-\beta)\tfrac{de^*}{dp})$$
$$-\frac{r\alpha(1-e^*)}{(1-p-r+r((1-e^*)(1-\alpha-\beta)+\beta))^2} - 2\frac{(1-p-r)r\alpha(1-e^*)}{(1-p-r+r((1-e^*)(1-\alpha-\beta)+\beta))^3}(1+r(1-\alpha-\beta)\tfrac{de^*}{dp})$$

Now for r large enough then it must be the case that p and $1-p-r$ are small. Substituting in $p \approx 0$ and $1-p-r \approx 0$ and using Lemma 25 below[1] and supposing that $\frac{de^*}{dp}$ is bounded, which we verify in Lemma 26 below, then

$$\frac{dA}{dp} \approx \frac{d^2e^*}{d\alpha dp} + \frac{de^*}{d\alpha}(\frac{(1-\alpha)(1-\alpha-\beta)}{r(e^*(1-\alpha-\beta)+\alpha)^2} - \frac{(1-\alpha-\beta)}{(1-p-r+r((1-e^*)(1-\alpha-\beta)+\beta))^2}),$$

and

$$\frac{dB}{dp} \approx -\frac{1}{r(e^*(1-\alpha-\beta)+\alpha)} - \frac{(1-\alpha)(1-e^*)}{(e^*(1-\alpha-\beta)+\alpha)^2} - \frac{1}{r((1-e^*)(1-\alpha-\beta)+\beta)} - \frac{\alpha(1-e^*)}{r((1-e^*)(1-\alpha-\beta)+\beta))^2} .$$

Now consider $\frac{de^*}{d\alpha}$ as above at $p \approx 0$ and $q = 1-p-r \approx 0$ then $\frac{de^*}{d\alpha} \approx 0$.

---

[1]We have to be a little careful here because from Equation (4.4) it also follows that $e^* \approx 0$. Well so long as $\frac{p}{e^*} \to 0$ as $p \to 0$, which is Lemma 25, then everything here follows.

Now since $\frac{dA}{dp} = \frac{dB}{dp}$, it follows that

$$\frac{d^2e^*}{d\alpha dp} \approx -\frac{1}{r(e^*(1-\alpha-\beta)+\alpha)} - \frac{(1-\alpha)(1-e^*)}{(e^*(1-\alpha-\beta)+\alpha)^2} - \frac{1}{r((1-e^*)(1-\alpha-\beta)+\beta)} - \frac{\alpha(1-e^*)}{r((1-e^*)(1-\alpha-\beta)+\beta))^2} < 0$$

thus $\frac{d^2e^*}{d\alpha dp} < 0$. $\qquad\square$

In the proof given above, a little bit of care must be taken inasmuch as when $p \approx 0$ and $q = 1 - p - r \approx 0$ then from Equation (4.4) it also follows that $e^* \approx 0$. As long as $\frac{p}{e^*} \to 0$ as $p \to 0$, then the approximations in the proof above are correct. Thus the lemmas below completes the proof.

Lemma 25. $\frac{p}{e^*} \to 0$ as $p \to 0$.

Proof. This is trivially the case from Equation (4.4) in the case that $q > 0$.

Suppose that $q = 0$ then from Equation (4.4) $e^* = \frac{p(1-\alpha)}{p+r(e^*(1-\alpha-\beta)+\alpha)}$.

Writing $x = \frac{p}{e^*}$ and rearranging this is equation is equivalent to

$$xp + rp(1 - \alpha - \beta) + r\alpha x = x^2(1 - \alpha).$$

This equation has two solutions for x:

$$x_1 = \frac{1}{1-\alpha}\left(\frac{1}{2}p + \frac{1}{2}r\alpha + \frac{1}{2}\sqrt{4pr - 6pr\alpha - 4pr\beta + 4pr\alpha\beta + p^2 + 4pr\alpha^2 + r^2\alpha^2}\right),$$

$$x_2 = \frac{1}{1-\alpha}\left(\frac{1}{2}p + \frac{1}{2}r\alpha - \frac{1}{2}\sqrt{4pr - 6pr\alpha - 4pr\beta + 4pr\alpha\beta + p^2 + 4pr\alpha^2 + r^2\alpha^2}\right)$$

(though as earlier shown only one of these will be consistent with $e^* \in (0, 1)$).

Let $p \to 0$ then get that both $x_1 \to 0$ and $x_2 \to 0$. In particular, this implies that $x = \frac{p}{e^*} \to 0$ as $p \to 0$. $\qquad\square$

Lemma 26. $\frac{de^*}{dp}$ is bounded for r large.

Proof. Well substituting $1 - p - r = q$ and taking the derivative of Equation (4.4) with respect to p yields:

$$\frac{de^*}{dp} = \frac{(1-\alpha)}{p+r(e^*(1-\alpha-\beta)+\alpha)} - \frac{p(1-\alpha)}{(p+r(e^*(1-\alpha-\beta)+\alpha))^2}\left(1 + \frac{de^*}{dp}r(1 - \alpha - \beta)\right)$$
$$- \frac{\alpha}{1-p-r+r((1-e^*)(1-\alpha-\beta)+\beta)} + \frac{(1-p-r)\alpha}{(1-p-r+r((1-e^*)(1-\alpha-\beta)+\beta))^2}\left(1 + \frac{de^*}{dp}r(1 - \alpha - \beta)\right)$$

Rearranging this expression yields:

$$\frac{de^*}{dp}\left(1 + pr(1 - \alpha)\frac{1-\alpha-\beta}{(p+r(\alpha+e^*(1-\alpha-\beta)))^2} - r\alpha(1 - p - r)\frac{1-\alpha-\beta}{(1-p-r+r(\beta+(1-e^*)(1-\alpha-\beta)))^2}\right) =$$
$$\frac{1-\alpha}{p+r(\alpha+e^*(1-\alpha-\beta))} - p\frac{1-\alpha}{(p+r(\alpha+e^*(1-\alpha-\beta)))^2}$$
$$- \frac{\alpha}{1-p-r+r\beta+r(1-e^*)(1-\alpha-\beta)} + \alpha\frac{1-p-r}{(1-p-r+r(\beta+(1-e^*)(1-\alpha-\beta)))^2}$$

Well for r large, we substitute in $p \approx 0$ and $1 - p - r \approx 0$ to obtain:

$$\frac{de^*}{dp} \approx \frac{1-\alpha}{r\alpha+re^*(1-\alpha-\beta)} - \frac{\alpha}{r\beta+r(1-e^*)(1-\alpha-\beta)}$$

which is bounded. $\qquad\square$

Additional Material for Chapter 5

E.1. Proofs of results in the chapter

In this section, we show that if $h(1 - c) > l$ then utilitarian social welfare is higher when Firm A is committed to high quality production than when it produces low quality. When Firm A is committed to high quality then recall $a_h = x_h = \frac{h^2(1-c)}{2(2+\sigma n)}$ from Equation (5.4), substituting into the utility function (Equation (5.1)) yields

$$U_h = (n+1)\frac{h^2(1-c)}{2(2+\sigma n)}\left(1 - \frac{1}{h^2}\frac{h^2(1-c)}{2(2+\sigma n)}\right) - \frac{2\sigma}{h^2}\frac{(n+1)n}{2}\frac{h^2(1-c)}{2(2+\sigma n)}\frac{h^2(1-c)}{2(2+\sigma n)} + Y - (n+1)p_h\frac{h^2(1-c)}{2(2+\sigma n)}$$

Now for Firm A and each of the rival firms

$$\pi_h = (p_h - c)\frac{h^2(1-c)}{2(2+\sigma n)}$$

So that welfare when firm A is committed to high quality and there are n rivals is given by

$$W_h(n) = (n+1)\frac{h^2(1-c)}{2(2+\sigma n)}\left(1 - \frac{1}{h^2}\frac{h^2(1-c)}{2(2+\sigma n)}\right) - \frac{2\sigma}{h^2}\frac{(n+1)n}{2}\frac{h^2(1-c)}{2(2+\sigma n)}\frac{h^2(1-c)}{2(2+\sigma n)} + Y - (n+1)pc\frac{h^2(1-c)}{2(2+\sigma n)}$$

$$W_h(n) = \frac{h^2(1-c)^2}{2^2(2+\sigma n)^2}(n + 1)(3 + \sigma n) + Y$$

Now suppose that Firm A cannot commit to high quality but instead produces at low quality, there are two cases to consider here.

Case 1: $2l + \sigma l(n - 1) - \sigma h(1 - c)n < 0$

In this case Firm A produces nothing and so welfare is given by $W_l(n) = W_h(n - 1)$

Now $\frac{d}{dn}W_h(n) = \frac{h^2(1-c)^2}{4}\frac{2+\sigma n(1-\sigma)+4(1-\sigma)}{(2+\sigma n)^3} > 0$

and so in particular $W_h(n) > W_h(n - 1)$

and hence in this case $W_h(n) > W_l$

Case 2: $2l + \sigma l(n - 1) - \sigma h(1 - c)n < 0$

Then by Equation (5.9) $a_l = \frac{l}{2}\frac{2l+\sigma l(n-1)-\sigma h(1-c)n}{(2+\sigma n)(2-\sigma)}$ and by Equation (5.8) $x_{hl} = \frac{h}{2}\frac{2h(1-c)-\sigma l}{(2+\sigma n)(2-\sigma)}$

Industry profits are given by $n(p_{hl} - c)x_{hl} + p_l a_l$ and so in this case

$$W_l(n) = n\frac{h}{2}\frac{2h(1-c)-\sigma l}{(2+\sigma n)(2-\sigma)}\left(1-\frac{1}{h^2}\frac{h}{2}\frac{2h(1-c)-\sigma l}{(2+\sigma n)(2-\sigma)}\right)$$
$$-\frac{2\sigma}{h^2}\frac{(n-1)n}{2}\frac{h}{2}\frac{2h(1-c)-\sigma l}{(2+\sigma n)(2-\sigma)}\frac{h}{2}\frac{2h(1-c)-\sigma l}{(2+\sigma n)(2-\sigma)}-cn\frac{h}{2}\frac{2h(1-c)-\sigma l}{(2+\sigma n)(2-\sigma)}$$
$$+\frac{l}{2}\frac{2l+\sigma l(n-1)-\sigma h(1-c)n}{(2+\sigma n)(2-\sigma)}\left(1-\frac{1}{l^2}\frac{l}{2}\frac{2l+\sigma l(n-1)-\sigma h(1-c)n}{(2+\sigma n)(2-\sigma)}\right)$$
$$-n\frac{2\sigma}{hl}\frac{h}{2}\frac{2h(1-c)-\sigma l}{(2+\sigma n)(2-\sigma)}\frac{l}{2}\frac{2l+\sigma l(n-1)-\sigma h(1-c)n}{(2+\sigma n)(2-\sigma)}+Y$$

Equivalently
$$W_l(n) = n\frac{1}{4}\frac{2h(1-c)-\sigma l}{(2+\sigma n)(2-\sigma)}\frac{6h(1-c)+(2h(1-c)-\sigma l)\sigma(n-1)-3\sigma l}{(2+\sigma n)(2-\sigma)}$$
$$+\frac{1}{4}\frac{2l+\sigma l(n-1)-\sigma h(1-c)n}{(2+\sigma n)(2-\sigma)}\frac{6l+3\sigma l(n-1)-2\sigma^2 ln+\sigma hn(1-c)}{(2+\sigma n)(2-\sigma)}+Y.$$

Now consider $(W_h - W_l)$:
$$W_h - W_l = \frac{h^2(1-c)^2}{2^2(2+\sigma n)^2}(n+1)(3+\sigma n)-n\frac{1}{4}\frac{2h(1-c)-\sigma l}{(2+\sigma n)(2-\sigma)}\frac{6h(1-c)+(2h(1-c)-\sigma l)\sigma(n-1)-3\sigma l}{(2+\sigma n)(2-\sigma)}$$
$$-\frac{1}{4}\frac{2l+\sigma l(n-1)-\sigma h(1-c)n}{(2+\sigma n)(2-\sigma)}\frac{6l+3\sigma l(n-1)-2\sigma^2 ln+\sigma hn(1-c)}{(2+\sigma n)(2-\sigma)}$$

rearranging terms, it follows that:
$$(W_h-W_l)\frac{4(2-\sigma)^2(2+\sigma n)^2}{h(1-c)-l} = -7\sigma^2 ln-12\sigma l+3l\sigma^2 n^2-l\sigma^3 n^2+l\sigma^3 n+12\sigma ln+3\sigma^2 l+$$
$$12l+ch\sigma^2 n+3ch\sigma^2 n^2-\sigma^3 hn^2c-3h\sigma^2 n^2-h\sigma^2 n+\sigma^3 hn^2-4\sigma hn+4\sigma hnc-12\sigma h+$$
$$12\sigma ch+\sigma^3 nh-\sigma^3 nch+3\sigma^2 h-3\sigma^2 ch+12h-12ch$$

Equivalently
$$(W_h-W_l)\frac{4(2-\sigma)^2(2+\sigma n)^2}{h(1-c)-l} = -n\sigma h(1-c)(4+\sigma+3\sigma n-\sigma^2(n+1))+\sigma nl(12-7\sigma-$$

$$\sigma^2(n-1)+3\sigma n)+3(h(1-c)+l)(4(1-\sigma)+\sigma^2)$$

Rearranging terms again
$$(W_h-W_l)\frac{4(2-\sigma)^2(2+\sigma n)^2}{h(1-c)-l} = (3\sigma n-\sigma^2 n)(2l+\sigma(n-1)l-n\sigma h(1-c))-(2-\sigma)l(3\sigma n-$$
$$\sigma^2 n)-n\sigma h(1-c)(4+\sigma-\sigma^2)+\sigma nl(12-7\sigma+\sigma^2)+3(h(1-c)+l)(4(1-\sigma)+\sigma^2)$$

Note that $(3\sigma n-\sigma^2 n)>0$ and by assumption $(2l+\sigma(n-1)l-n\sigma h(1-c))>0$ and so
$$(W_h-W_l)\frac{4(2-\sigma)^2(2+\sigma n)^2}{h(1-c)-l} > -(2-\sigma)l(3\sigma n-\sigma^2 n)-n\sigma h(1-c)(4+\sigma-\sigma^2)+\sigma nl(12-$$
$$7\sigma+\sigma^2)+3(h(1-c)+l)(4(1-\sigma)+\sigma^2)$$

the right hand side of this inequality it equal to:
$$(2l+\sigma(n-1)l-n\sigma h(1-c))(4+\sigma-\sigma^2)-8l+2l\sigma+ln\sigma(2-\sigma)(1-\sigma)+2l\sigma^2+$$
$$l\sigma^2(1-\sigma)+3(h(1-c)+l)(4(1-\sigma)+\sigma^2)$$

and since $4+\sigma-\sigma^2>0$ is greater than or equal to
$$-8l+2l\sigma+ln\sigma(2-\sigma)(1-\sigma)+2l\sigma^2+l\sigma^2(1-\sigma)+3(h(1-c)+l)(4(1-\sigma)+\sigma^2)$$

and $n>1$ and supposing that since $h(1-c)>l$ then this in turn is greater than or equal to

$-8l + 2l\sigma + l\sigma(2-\sigma)(1-\sigma) + 2l\sigma^2 + l\sigma^2(1-\sigma) + 6l(4(1-\sigma) + 3\sigma^2) = 4l(4-\sigma)(1-\sigma) + 14l\sigma^2 > 0$

and thus in this case too when $h(1-c) > l > 0$, then $W_h > W_l$.

E.2. All firms make quality decisions

In this section, we modify the model in the main body of the text by assuming that all $n + 1$ firms in the industry are symmetric and must make a decision in each period whether to produce high or low quality (that is the n rivals of Firm A are no longer committed to high quality exogenously). The question we address below is whether an equilibrium exists in which all firms in the industry can credibly commit to high quality production.

E.2.1. Static benchmark: m firms low, $n + 1 - m$ high

First consider the static benchmark in which product quality is anticipated and can be observed prior to purchase and where of the $n + 1$ firms in the industry, m produce low quality and the rest produce high quality.

Suppose that in equilibrium each low quality producer's output is x_{lm} and each high quality producer's output is x_{hm}.

Then for a low quality producer, the x_{lm} is determined by choosing x to maximise:

$$(1 - \frac{2x}{l^2} - \frac{2\sigma}{hl}(n+1-m)x_{hm} - \frac{2\sigma}{l^2}(m-1)x_{lm})x$$

Similarly for a high quality producer, the problem is to choose x to maximise

$$(1 - \frac{2x}{h^2} - \frac{2\sigma}{h^2}(n-m)x_{hm} - \frac{2\sigma}{hl}mx_{lm} - c)x$$

So the first order conditions yield:

$x_{hm} = \frac{h^2(1-c)}{2(2+\sigma(n-m))} - \frac{mh\sigma}{l(2+\sigma(n-m))}x_{lm}$

and

$1 - \frac{4x_{lm}}{l^2} - \frac{2\sigma}{hl}(n-m+1)x_{hm} - \frac{2\sigma}{l^2}(m-1)x_{lm} = 0$

so long as $x_{lm} > 0$

in this latter case

$x_{lm} = \frac{l}{2}\frac{l(2+\sigma(n-m))-\sigma h(1-c)(n+1-m)}{4+2\sigma(n-1)-\sigma^2 n}$

Thus if $2l + \sigma l(n-m) - \sigma h(1-c)(n+1-m) > 0$ (a condition which is more likely to hold the larger is m)

then $x_{lm} = \frac{l}{2} \frac{l(2+\sigma(n-m-1))-\sigma h(1-c)(n-m)}{4+2\sigma(n-2)-\sigma^2(n-1)}$ and

$x_{hm} = \frac{h}{2} \frac{2h(1-c)+h\sigma(m-1)(1-c)-l\sigma m}{(2+\sigma(n+1)-\sigma)(2-\sigma)}$.

Then $\pi_{lm} = \frac{1}{2}(\frac{2l+\sigma l(n-m)-\sigma h(1-c)(n+1-m)}{4+2\sigma(n-1)-\sigma^2 n})^2$ which is increasing in m and $\pi_{hm} = \frac{1}{2}\frac{(2h(1-c)+h\sigma(m-1)(1-c)-l\sigma m)^2}{(2+\sigma(n+1)-\sigma)^2(2-\sigma)^2}$ which is increasing in m.

Otherwise, that is if $l(2+\sigma(n-m))-\sigma h(1-c)(n+1-m) < 0$, then $\pi_{lm} = 0$ and $\pi_{hm} = \frac{1}{2}\frac{h^2(1-c)^2}{(2+\sigma(n-m))^2}$ which is increasing in m.

Note in particular that consideration of these two cases implies that π_{lm} is increasing in m.[1]

E.2.2. Necessary and sufficient conditions

That π_{lm} is increasing in m suggests that for a firm producing low quality and anticipated to produce low quality the lowest profits are earned when all other firms produce high quality and so a most severe punishment for defection from a situation when all are producing high quality is that following a defection, all the rivals continue producing high quality. Thus a necessary condition for the existence of an equilibrium in which all firms in the industry produce high quality is (5.13).

However such a continuation may not be feasible, that is it may not be the case that a continuation in which one firm produces low quality and all others produce high can be sustained as an equilibrium. One continuation equilibrium that can always be sustained is that all firms produce low quality for ever. Specifically, we suppose that following a deviation by a firm, in the future all firms produce low quality outputs—it is clear that the most severe punishment that can be sustained in equilibrium is at least as severe as this one. Using the earlier static results it follows that a sufficient condition which ensures that there is an equilibrium in which all firms produce high quality output is given by, it follows that a sufficient condition for an equilibrium in which all firms produce high quality is given by:

---

[1] In addition and as one might anticipate π_{hm} is also increasing in m.
π_{hm} is increasing in the region $l(2+\sigma(n-m))-\sigma h(1-c)(n+1-m) > 0$ and in the region $l(2+\sigma(n-m))-\sigma h(1-c)(n+1-m) > 0$.
At $l(2+\sigma(n-m))-\sigma h(1-c)(n+1-m) = 0$ then $l = \frac{\sigma h(1-c)(n+1-m)}{2+\sigma(n-m)}$ and
$\frac{1}{2}\frac{(2h(1-c)+h\sigma(m-1)(1-c)-l\sigma m)^2}{(2+\sigma(n+1)-\sigma)^2(2-\sigma)^2} = \frac{1}{2}\frac{h^2(1-c)^2}{(2+\sigma(n-m))^2}$
so π_{hm} is continuous throughout and so in particular the two regions connect and so π_{hm} is increasing in m.

$$\frac{1}{1-r}\pi_h \geq \pi_d + \frac{r}{1-r}\pi_{ll},$$

where $\pi_{ll} = \frac{1}{2}\frac{l^2}{(2+\sigma n)^2}$ is the short-term profit that a firm in the industry earns when all the firms produce low quality and each optimizes with respect to the quantity produced.

Let $\Delta_{suff} = \frac{1}{1-r}\pi_h - \pi_d - \frac{r}{1-r}\pi_{ll}$. We proceed by considering the comparative statics of Δ_{suff} with respect to n and σ—the measures of competition.

$\Delta_{suff} = \frac{1}{1-r}\frac{1}{2}\frac{h^2(1-c)^2}{(2+\sigma n)^2} - \frac{h^2}{8}\left(\frac{2+\sigma cn}{2+\sigma n}\right)^2 - \frac{r}{1-r}\frac{1}{2}\frac{l^2}{(2+\sigma n)^2}$.

Equivalently $\Delta_{suff} = \frac{4h^2(1-c)^2-(1-r)h^2(2+\sigma cn)^2-4rl^2}{8(1-r)(2+\sigma n)^2}$.

The first derivative with respect to n is given by:

$\frac{d\Delta_{suff}}{dn} = -\frac{(1-r)h^22(2+\sigma cn)\sigma c}{8(1-r)2+\sigma n)^2} - 2\sigma\frac{4h^2(1-c)^2-(1-r)h^2(2+\sigma cn)^2-4rl^2}{8(1-r)(2+\sigma n)^3} = -\frac{(1-r)h^22(2+\sigma cn)\sigma c}{8(1-r)2+\sigma n)^2} - 2\sigma\Delta_{suff}$

or equivalently $\frac{d\Delta_{suff}}{dn} = -\frac{(1-r)h^22(2+\sigma cn)\sigma c}{8(1-r)2+\sigma n)^2} - 2\sigma\Delta_{suff}$.

Note in particular that if $\Delta_{nec} > 0$ then $\frac{d\Delta_{nec}}{dn} < 0$ but otherwise it is not clear how to sign this.

The first derivative with respect to σ is:

$\frac{d\Delta_{suff}}{d\sigma} = -\frac{(1-r)h^22(2+\sigma cn)cn}{8(1-r)(2+\sigma n)^2} - 2n\frac{4h^2(1-c)^2-(1-r)h^2(2+\sigma cn)^2-4rl^2}{8(1-r)(2+\sigma n)^3} =$

or equivalently $\frac{d\Delta_{suff}}{d\sigma} = -\frac{(1-r)h^22(2+\sigma cn)cn}{8(1-r)(2+\sigma n)^2} - 2n\Delta_{suff}$.

Again if $\Delta_{suff} > 0$ then $\frac{d\Delta_{suff}}{dn} < 0$ but otherwise the sign of $\frac{d\Delta_{nec}}{d\sigma}$ is ambiguous.

Thus increased competition either through an increase in n or an increase in σ cannot make it more likely that $\Delta_{suff} > 0$; however, if there is an additional benefit to non-deviation then the necessary condition $\Delta_{suff} + B > 0$ may have an ambiguous relationship with n and σ.

E.2.3. Existence of an equilibrium

When the sufficient condition holds, then an equilibrium in which all firms produce high quality can be sustained, however when it fails then the possibility that such an equilibrium can be sustained cannot be ruled out. However, existence of such an equilibrium can be ruled out when the necessity condition fails. Thus by examining the two conditions simultaneously, it is possible to show that the existence of such an equilibrium is not monotonic in the measures of the degree of competition in the industry. Specifically such an equilibrium exists when $\Delta_{suff} + B > 0$ but does not exist when $\Delta + B < 0$; in
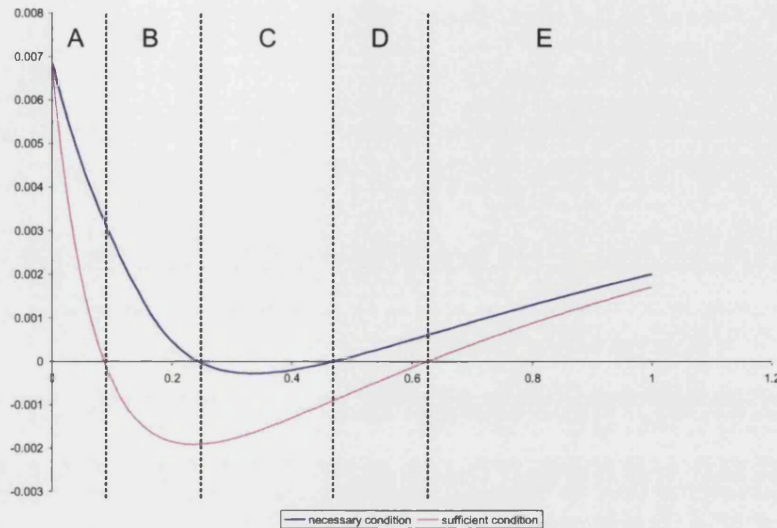
Figure E.1. Necessity and sufficiency conditions against sigma at h=0.8,
l=0.2, c=0.39, n=8, r=0.6 and B=0.02

the intermediate case (that is when $\Delta_{suff} < 0$ and $\Delta > 0$) the conditions are insufficient
to prove the existence or non-existence of such an equilibrium).[2]

In the light of earlier results and examples, it is perhaps unsurprising that the effect
of an increase in the level of competition (either as an increase in σ or in n) has an
ambiguous effect on the possibility that an equilibrium in which all firms can credibly
commit to produce high quality outputs. This is demonstrated by the Figures E.1 and
E.2. In both these figures an equilibrium in which all the firms in the industry can
credibly commit to high quality exists in the regions A and E but not in the region C.[3]

---

[2]It is trivial to show that $\Delta > \Delta_{suff}$ and indeed is a corollary of the result that π_{lm} is increasing in m.
[3]In these figures the line "necessary condition" represents $\Delta + B$ and the line "sufficient condition"
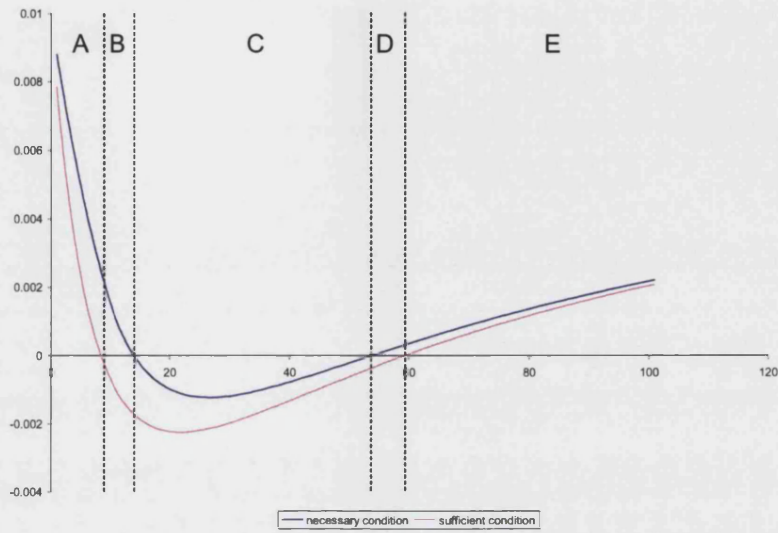$\Delta_{suff} + B$.

Figure E.2. Necessity and sufficiency conditions against n at h=0.9, l=0.2, c=0.3, sigma=0.1, r=0.5 and B=0.017