

Seller Reputation

Heski Bar-Isaac¹ and Steven Tadelis²

¹ *Department of Economics, New York University, 269 Mercer Street,
New York, NY 10003, USA, heski@nyu.edu*

² *Walter A Haas School of Business, University of California, Berkeley, CA
94720, USA, stadelis@haas.berkeley.edu*

Abstract

Seller reputation is an important asset because buyers often choose sellers on the basis of their reputation. This is particularly true when the quality of the good or service transacted is hard to measure and the parties cannot perfectly contract on the outcome of the transaction. As a consequence, the seller will be mindful of building and maintaining a good reputation through the information that buyers have about the seller, including previous transactions and the reports of other buyers.

We introduce a unifying framework that embeds a number of different approaches to seller reputation, incorporating both hidden information and hidden action. We use this framework to stress that the way in which consumers learn affects both behavior and outcomes. In particular, the extent to which information is generated and socially aggregated determines the efficiency of markets.

After reviewing these theoretical building blocks we discuss several applications and empirical concerns. We highlight that the environment in which a transaction is embedded can help determine whether the transaction will occur and how parties will behave. Institutions, ranging

from the design of online markets to norms in a community, can be understood as ensuring that concerns for reputation lead to more efficient outcomes. Similarly, the desire to affect consumer beliefs regarding the firm's incentives can help us understand strategic firm decisions that seem unrelated to the particular transactions they wish to promote.

We conclude by considering slightly different models of reputation that lie beyond the scope of our framework, briefly reviewing the somewhat sparse empirical literature, and highlighting and suggesting future directions for research.

1

Introduction

Most economic transactions are described by one party procuring goods or services from another party, either through monetary exchange, barter or the promise of future reciprocation. The most naive model of economic exchange assumes that the qualities and characteristics of these transactions are well known and understood by the parties involved, and that markets will clear to allocate the goods and services to those who value them the most. This approach — albeit useful in generating some building blocks of economic analysis — is often inadequate to describe many realistic situations of exchange. Common examples include mundane transactions in which a person buys a bottle of wine with unknown quality, a firm who hires an employee with unknown talent, or on a larger scale, a government who procures a weapons system with unknown properties in the battlefield.

In Akerlof's (1970) classic article *The Market for Lemons*, it is shown that this kind of uncertainty can hinder the operation of markets to the possible extreme of markets failing to operate despite obvious gains from trade. That is, in the face of inherent quality uncertainty market failures will prohibit efficient exchange. This uncertainty can

stem from two possible sources that are two central pillars of what is now referred to as “the economics of information.”

First, quality uncertainty may be a result of *unobserved primitives* that determine the quality of the good or service in the spirit of Akerlof’s “adverse selection.” As an example, consider a company that wishes to procure consulting services. If this skill of the employees in the consulting firm will determine the quality of the transaction then uncertainty over the employees skill may deter the company from paying a fee that would be adequate to engage the consulting firm. This is true even though both the company and the consulting firm would benefit from establishing the relationship.

Second, quality uncertainty may be a result of *unobserved actions* that determines the quality of the good or service, what is known now as “moral hazard.” Using the example above, if the company can gather information and screen the skills of different consulting firms but cannot monitor the consultants’ productive effort, then the company may fear that the hired consultants will shirk on their job, and the transaction may therefore be avoided altogether. Of course, both hidden information and hidden action might be present simultaneously.¹

Still, despite the existence of such problems in the marketplace, the exchange of extremely complex goods and services, with performance measures that are hard to describe or monitor, is commonplace. The question posed is then, what are the remedies that foster exchange in these hazardous market environments? One such remedy is the introduction of contingent contracts. Obviously, if the consultant’s effort can be contracted upon, then a contract of the form “you will be paid if you work adequately and you will not be paid if you shirk” will solve the moral hazard problem. Similarly, if skills can be later verified then a contract of the form “you will get a base pay commensurate with a low skill consultant, and a bonus if it turns out that you are a high skilled consultant” will offer the consultant in the consulting firm adequate reward, while shielding the company from overpaying for low skilled consultants.

¹ Interestingly, adding hidden information may either mitigate or exacerbate the problems created by hidden action. We explore some of these issues in detail later.

When such remedies are unavailable or prohibitively costly, then the information and beliefs that the buyers have about sellers will play a crucial role in determining whether a transaction takes place and the efficiency of trade. Such information and beliefs about the seller's skill and behavior, which we refer to as the seller's "reputation," are a consequence of many things. These include direct observations on past performance, experience with other sellers, reports from third parties, actions that the seller may undertake outside of the transaction, and numerous other factors for which we attempt to provide a taxonomy.

Specifically, focussing first on information conveyed by past transactions, we distinguish between three cases as follows.

Pure Hidden Information (Section 3): In this case a seller has no active control over the outcome of the transaction, but sellers vary in their innate ability, or "type." For example, a consultant may be more or less smart, affecting his ability to deliver; a chef may have talent in creating exciting recipes, or may not. As potential buyers observe the output produced by such sellers, they effectively learn over time about the seller's innate skill, and hence we also refer to this situation as "pure learning." In this case the seller's reputation is the buyer's belief about the seller's skill, or type.² As a somewhat amusing, yet concrete example, consider the movie *Gigli*. The quality of the movie is now fixed. Neither of the authors has seen this movie to date. In deciding not to watch it, we have beliefs about how much we would enjoy sitting through it based on reports of others (such as the reviewers and trusted friends) who have watched it and reported their (largely negative) experiences. Generally for movies (though possibly not for *Gigli*) some viewers would expect to enjoy it and others not, and so the movie's "reputation" can be seen as the probability with which an audience member would enjoy it.

Pure Hidden Action (Section 4): In this case there is only one "type" of seller, but this seller has active control over the outcome of the

²Including a discussion of learning in a review of reputation is a somewhat idiosyncratic choice, but while the formal literatures have largely been quite separate, in application at least there is an important and obvious connection between how a buyer learns and a seller's desire to affect what she learns.

transaction, and his actions within the transaction are not contractible. For example, a consultant may work hard or slack off; a chef/owner of a restaurant may use expensive fresh and high quality ingredients or instead choose to purchase cheap ones. The hidden action qualifier is that our buyer cannot observe the seller's behavior, and hence must try to infer what the seller will do. As a consequence, in this case, reputation is not about learning some underlying trait of the seller, but instead reflects the buyer's belief about the seller's equilibrium behavior.

Mixed Models (Section 5): This more realistic setting includes both hidden information and hidden action where sellers vary in their type and they can unobservably affect the outcome of the transaction. For example, a consultant's ability to successfully improve the company's performance depends on both his unobserved skill and the unobserved effort that he takes; the quality of a meal depends on both the chef's skill and the choices he makes about what kind of ingredients to purchase. In this case the "reputation" is the buyer's belief about the seller's type as well as the equilibrium behavior of the different types of sellers. Hence, in this more realistic setting, reputation includes both a belief about underlying traits and about anticipated equilibrium behavior.

The pure hidden-information approach rests on the idea that there is some underlying truth about sellers that is not manipulable, and buyers are using past performance to learn about this truth. Hence, the method of analysis incorporates learning by buyers as outcomes are revealed. In contrast, the pure moral hazard approach rests on the premise of repeated games with the idea that "what goes around, comes around." That is, if buyers and sellers interact time and time again, then inadequate behavior on part of the sellers can be punished by retaliatory behavior of the buyers: if some seller acts to deliberately cheat some buyer, then the buyers at large can stop interacting with this seller, assuming that information about his behavior is easily accessible. This "carrot and stick" approach is the basic model of a repeated game with complete information. In it, a seller sells repeatedly to one, or many buyers, and the outcome of each transaction is observed by all

the participants. The seller has short run incentives to shirk on quality, but will refrain from shirking if long run incentives to perform well are in place. These incentives are typically provided using the standard trigger (or bootstrap) equilibrium, often referred to as a “reputational equilibrium.” The mechanism at work is the celebrated “folk theorem.”

It is worth noting that the hidden information and mixed-model approaches still incorporate part of the underlying logic of “what goes around, comes around.” However, it is not so much based on the notion of a retaliatory punishment, but instead through beliefs that follow a rather appealing process. Following bad outcomes buyers will update their beliefs to accommodate the higher likelihood that poor quality is persistent, and thus will beware of sellers who provided poor quality in the past.

We will begin our review by exploring these three different models for what determines the outcome of a transaction. A few key lessons will emerge from these building blocks, which shed light on four factors that are key in determining the extent to which reputation acts effectively in leading to efficient trade. Specifically, these are: (a) the extent of uncertainty about the seller; (b) the rate at which buyers learn from outcomes, including, for example, the rate of information diffusion among buyers; (c) the seller’s discount factor, or value of future interactions; and (d) characteristics of demand that determine how sensitive buyers are to reputation.

We then turn in Sections 6–11 to applications and extensions of these building block ideas. First, we discuss how exogenous institutions — ranging from the existence of markets for reputations, the design of an electronic marketplace with built-in feedback mechanisms, to social norms in a community — affect how well reputational concerns work in achieving efficient market outcomes. In addition to institutions which are exogenous from the perspective of an individual buyer or seller, we discuss actions outside of transactions that the seller might take to influence his reputation or the way that his reputation evolves. Furthermore, the discussion highlights that when discussing “reputation” in the context of a particular application, a researcher must take care in thinking through precisely the question of “reputation for what.”

We conclude our review by raising concerns not adequately captured by our basic models, as well as the difficulty in applying some of the theoretical exercises to empirical scrutiny. Specifically, we briefly introduce and discuss a recent literature on “reputation for experts” where in contrast to the discussion above, even following a transaction, the outcome may be hard to classify as “good news” or “bad news,” and so a seller may be tempted to undertake actions to satisfy the buyer’s expectations regardless of the extent to which this is the right outcome from an efficiency perspective. We discuss implications of having multiple dimensions of skills or ability and multi-dimensional reputations. We also briefly highlight that while the bulk of the literature has focused on “vertical” (quality) aspects or the extent to which a seller is “good,” a seller may also take actions to affect “horizontal” reputation or buyers beliefs about the extent to which the seller will match their tastes.

Any survey will be incomplete, and this one reflects our overlapping interests in this topic, as well as our personal judgment on where the boundaries of this topic lie. In this regard, it is worth highlighting that there are a number of other recent reviews. In particular, Cripps (2006) provides a short but useful survey on building block models, Mailath and Samuelson (2006) provide a detailed treatment, and Macleod (2007) provides a complimentary survey that highlights a contractual interpretation of reputation and the interaction between formal contracts and reputation effects.

2

A Simple Framework

We begin by presenting a simple framework that will allow us to nest different approaches that have been used to model reputation, as we discussed earlier in the introduction. Consider a seller who over several periods can sell a good or service to a buyer, or to a sequence of buyers. In every period, the seller can engage in (at most) one transaction. The transaction can either yield a successful outcome, which is worth 1 to any buyer, or a failure, which is worth 0. Buyers can observe the outcome only after a transaction has taken place and cannot write contingent contracts based on the outcomes of the transactions. The likelihood of success may depend on the seller's effort as well as her ability.

Suppose that a seller can be either "good" or "bad." We index the seller's type by $\theta \in \{b, g\}$, where $0 \leq b \leq g \leq 1$ and let μ_t denote the buyers' belief that the seller is a good type in the end of the t th period of trade. This belief will surely play a role in determining a seller's "reputation." However, as we will see, other components will also have an influence.

The technology is such that the probability of success depends on the seller's effort and on his type. In particular, for any given transaction

the probability of success is given by $e\theta$, where $e \in [\underline{e}, 1]$ denotes the seller's effort, and $\underline{e} > 0$. Typically, we assume that it costs the seller $c(e)$ to exert effort e , where $c'(\underline{e}) = 0$ and this cost is increasing and convex ($c' > 0$ and $c'' > 0$). In order to ensure interior solutions assume that $c'(1)$ is sufficiently high. Finally, the buyer and the seller are risk neutral, and makes decisions in order to maximize the present value of lifetime earnings, when discounting the future at the rate δ .

We say that a transaction is executed if a buyer chooses to purchase the service provided by the seller at some price $p > 0$, and the seller chooses to supply the service at that price. In what follows, we will usually analyze the case in which the seller is a monopolist and the buyers are left with no rents. However, we will sometimes use a different approach. We will discuss each such case as it arises.

The timing of exchange will proceed as follows. There are $T > 1$ periods, and each period has the following sequence of events. First, a buyer is matched with the seller, and decides whether or not to purchase the service at the price p . If the buyer and seller choose to proceed with the transaction, then the buyer pays p to the seller. The seller then chooses how much effort to exert, and the service succeeds with probability $e\theta$. Finally, the buyer, and possibly other potential future buyers, observe whether the service was a success or not, and beliefs are updated accordingly. In the first period the buyers share a correct prior $\mu_0 > 0$ which is common knowledge.

Notice that the efficient level of effort that maximizes social surplus given a seller's type θ , which we denote as e_θ^* , is obtained by maximizing

$$\theta e - c(e),$$

which assuming an interior solution yields the optimality condition $\theta = c'(e_\theta^*)$. Since "higher" types are more productive (the multiplicative technology $e\theta$ exhibits a single crossing property) then $e_g^* > e_b^*$. Finally, to make things interesting we assume that $\underline{e} < e_b^*$ so that it is efficient for both types to exert some effort above the minimum level \underline{e} .

3

Pure Hidden Information (The Learning Approach)

In this section, we suppose that the seller cannot manipulate the probability of success, that is, there is no effort decision and we set $\underline{e} = 1$.¹ Thus, the probability of success depends only on the type of the seller, so that a type θ succeeds with probability θ . The beliefs that buyers have about the seller's type will, in this case, fully capture the seller's reputation. Over time, from observing the outcome of past transactions (either directly or perhaps indirectly from outcomes with other buyers), a potential buyer can learn the type of the seller.

Assume that a good seller has a strictly higher probability of success than a bad one so that $0 \leq b < g \leq 1$.² Also, assume that the outcome of any transaction is observed by all potential buyers. Following a successful outcome S in the first period, buyers will have a new updated belief that is calculated according to Bayes' rule. Specifically, we write,

$$\mu_1(\mu_0, S) = \frac{\mu_0 g}{\mu_0 g + (1 - \mu_0) b}.$$

¹Note that here we depart from the assumption that $c'(\underline{e}) = 0$, and instead assume that $c(1) = 0$ and $c(e) = \infty$ for all $e \neq 1$.

²If $b = 0$ and $g = 1$ then complete learning will occur after one period. Dynamics are more interesting when at least one of the two weak inequalities is strict.

Similarly, we can write an expression for the posterior following the observation of a failure F in the first period,

$$\mu_1(\mu_0, F) = \frac{\mu_0(1-g)}{\mu_0(1-g) + (1-\mu_0)(1-b)}. \quad (3.1)$$

It is intuitive, and it can easily be proved, that as buyers observes more and more outcomes, the reputation (or posterior belief) converges to the truth, in the sense that after sufficiently many observations, buyers will almost surely know whether they are facing a good or a bad type of seller. That is, as $t \rightarrow \infty$, μ_t will converge either to 0 (if the seller is a bad type) or to 1 (if the seller is a good type).

This intuitive “convergence to the truth” result depends on buyers transacting with the seller infinitely often. Interestingly, it may be the case that buyers will cease to interact with a seller before all the information is learned. If this occurs then buyers will cease to receive new information about the seller’s ability. To use our earlier example, such is the current reputation of *Gigli* that none of our acquaintances are likely to see it and we are unlikely to receive additional information about the extent to which it is worth viewing.

To see this imagine that the price p at which transactions can be executed is fixed (it is not necessary to set a fixed price but it simplifies things dramatically). Clearly, if $p > g$ then regardless of the seller’s reputation, buyers will not find it worthwhile to buy the service because the expected value of transacting is $E\theta = \mu g + (1-\mu)b \leq g$. Similarly, if $p < b$ then buyers will always purchase. The more interesting case occurs when $b < p < g$.

Suppose that buyers are short-lived and, as such have no reason to take actions for the purpose of generating information. In a single transaction, a buyer would buy whenever the reputation μ is high enough. Specifically, when $\mu g + (1-\mu)b > p$, or equivalently,

$$\mu \geq \underline{\mu} \equiv \frac{p-b}{g-b}.$$

Now imagine that $\mu_0 = \underline{\mu} + \varepsilon$, with ε small. That is, the likelihood of a good type is barely high enough to support trade. In the first period, trade will occur, but what will happen if this buyer experiences a failure? The new belief $\mu_1(\mu_0, F)$ will follow from (3.1) above, and for

ε sufficiently small we will have $\mu_1(\mu_0, F) < \underline{\mu}$ and no buyer will now be willing to buy at the price p . Thus, if $g < 1$, our seller who actually may be a good type that was unlucky, is “stuck” with a poor reputation that she does not deserve. Similarly even with a higher $\mu_0 > \underline{\mu}$, a good seller can suffer a run of failures and at some point in time the seller’s reputation can fall so low that no one buys and again the seller is stuck with an inappropriate low reputation.

The result that bad luck can drive a good seller out of the market can be sustained even when the price varies with the reputation. Bar-Isaac (2003) allows price to endogenously reflect buyers’ expected value for buying the good. Using our framework, if at period t the seller has a reputation of μ_t then the highest price buyers are willing to pay is $p = \mu_t g + (1 - \mu_t)b$, and assuming that the seller is a perfect monopolist, this price can be used as an endogenous price in equilibrium.³ In Bar-Isaac (2003), it is assumed that there is a marginal cost of production which lies between b and g so that trade with a good seller is efficient but trade with a bad seller is not. If due to a drop in reputation price falls below marginal costs, then sellers who trade will lose money on those transactions.

Notice that if the seller does not know her own type then again, a good seller may suffer a run of bad luck that drives her out of the market. Clearly, she will be as convinced as the buyers are about the likelihood that she is a good type, and will not see a reason to incur losses after her reputation drops sufficiently low to imply a price that is sufficiently below marginal cost.⁴ However, if the seller does know whether she is good or bad then buyers might learn from the seller’s decisions to continue trading rather than drop out of the market when price is sufficiently below marginal cost, and in the long run buyers will end up making the same decisions as if they knew the truth.⁵

³In Appendix B of Bar-Isaac (2004) similar qualitative results are shown when the buyer sets price.

⁴It is not enough that price falls just below marginal costs for the market to shut down because there is an option value from trading at a slight loss since reputations may subsequently improve.

⁵Pricing in the context of this model has also been discussed in the case of competing sellers who do not know their abilities. When there are competing sellers, then sellers set prices not only to extract surplus but also to induce further learning. Bergemann

Though the model is somewhat stylized, circumstances in which beliefs might get stuck with inefficient long run outcomes are much more general. Indeed one might ask under what circumstances a single long lived buyer who cares about the future might learn everything worth knowing. As suggested by the example above, when the buyer discounts the future then he may not learn everything — the classic reference, albeit not couched in terms of a buyer learning about the quality of a seller, is Rothschild (1974). The interested reader is also referred to Aghion et al. (1992) who introduce sufficient conditions for complete learning and discuss an interesting example where it fails. Further issues arise when there are a number of buyers interested in learning about a seller. In particular, buyers who can choose when to consume may free-ride and wait for the outcome of other buyers; alternatively, by experimenting and generating some new information they may stimulate others into further experimentation, as discussed in Bolton and Harris (1999).

In addition to reputations getting stuck in such a way that no new information is generated, there may also be difficulties in aggregating information efficiently. Specifically, suppose that there is a continuum of potential buyers, each of whom receives a private signal about the quality of a seller that is not shared with other buyers. Then, even though a buyer who was able to observe all the signals would know the type of the seller, without a mechanism to disseminate and centralize this information, some buyers may hold inappropriate beliefs. A natural means by which such information is spread in society is that potential buyers observe whether previous buyers who had the opportunity to buy chose to do so. However, following some histories, an individual buyer's private information might be overwhelmed by what he learns from predecessors' decisions, regardless of the information that he holds. In such circumstances, his decision is uninformative to future buyers and since they will face similar circumstances, no further

and Välimäki (1996, 2006) and Felli and Harris (1996) show that learning can be efficient, though of course this does not mean that buyers will eventually learn the truth. Returning to monopoly, but assuming that consumers have private signals about the outcomes and only learn from each others' purchase decisions (buy or not), the seller might choose prices in such a way as to induce information herding as discussed by Bose et al. (2006).

information is generated. This has been termed information-herding or an information cascade in a large literature drawing from Bikchandani et al. (1992) and Banerjee (1992).

The key insight to take away from this section — and one which is implicit in much of the literature discussed below — is the importance of the extent to which buyers keep getting *new* and *informative* signals which help them update beliefs about the seller.

4

Pure Hidden Action (The Repeated Game Approach)

In a model of “pure” hidden action (also referred to as moral hazard), buyers have no doubt as to the “type” of the seller and are fully aware of the seller’s ability. In the context of our framework this is achieved by setting $b = g$ (or taking $\mu_0 = 1$). Although the buyers have no doubts about the seller’s ability, they may have some concern as to the action that the seller undertakes. Recall that we assumed that $\underline{e} < e_g^*$ so that the seller requires some incentives to exert efficient effort. Buyers cannot observe the seller’s effort and must form some expectation about it. Furthermore, the seller may be tempted to choose sub-optimal effort since effort is costly. For example, a customer may be unsure of the quality of ingredients used by a chef in a restaurant, and the chef has control over this quality by deciding how much to spend on the ingredients.

To begin, suppose that the price of the good is fixed at p and that $\underline{e}g < p < e_g^*g$. Clearly, if there is only one opportunity for transacting, the unique outcome is that the buyer does not buy because he correctly anticipates that the seller would not exert effort beyond the minimum level \underline{e} . Thus, if we are to have some effort exerted that is above the minimal level \underline{e} , there must be some “carrot and stick” continuation

through which the seller is rewarded for good behavior and punished for poor behavior.

If there are potentially infinitely many opportunities for trade and the seller values the future sufficiently then there may be trade and a wide range of outcomes can be supported.¹ Underlying this new richer set of outcomes that is supported by “carrot and stick” style equilibria is the discipline imposed on the seller’s effort decision by the buyers’ future responses to various outcomes. This is the key insight in the seminal work of Klein and Leffler (1981).

As a specific (and familiar) example, we construct a “trigger strategy” or “bootstrap” equilibrium in which the buyers buy, and the seller exerts some effort level $\tilde{e} > \underline{e}$, so long as all transactions up to that date have yielded successes. Once a failure is realized, buyers stop buying and the seller will choose \underline{e} if a transaction is initiated. Let V denote the seller’s value of starting a period with the “reputation” intact, where he anticipates a transaction this period at price p , and expects to choose effort level \tilde{e} . This is defined recursively as $V = p - c(\tilde{e}) + \delta\tilde{e}gV$ or

$$V = \frac{p - c(\tilde{e})}{1 - \delta\tilde{e}g}.$$

In equilibrium the seller must be choosing effort optimally, that is to maximize expected utility $p - c(e) + \delta egV$ so that \tilde{e} solves

$$c'(\tilde{e}) = \delta gV.$$

Moreover, for buyers to be willing to purchase at the price p , equilibrium requires that $p \leq \tilde{e}g$.

A number of results arise immediately. First, assuming that there is an interior solution, it is easy to see that V is increasing in p , δ , and g . Furthermore, since $c'' > 0$, \tilde{e} is increasing in p, δ , and g . That is, the equilibrium effort is increasing in the discount factor δ , in the price of the good p and in g (which in this hidden action model can be thought of as parametrizing the efficacy of additional effort). The intuition for these results is simple and appealing. When g increases, the

¹It is well known and easy to see that if there is a well defined finite number of trade periods then the last period is just like the one-off transaction described above, and backward induction results in no trade in every period.

productivity of effort increases and hence \tilde{e} increases. When p increases or when δ increases, the value of the relationship increases and hence investing in it is more fruitful.

Several points are worth noting. First, the equilibrium described above is not the only equilibrium. There may be other equilibria that support similar average payoffs through different strategies, and other equilibria altogether. As always with repeated game models, there is an equilibrium in which no buyer buys anticipating that a seller exerts minimal effort, and the seller exerts minimal effort anticipating that no buyers will wish to buy in the future. Second, allowing the seller to extract all the surplus from transactions so that $p = \tilde{e}g$ does not lead to an efficient level of effort. For example, if δ is close to 0 then efforts are minimal. In contrast, when δ approaches 1 then many levels of effort can be sustained in equilibrium, even greater than the efficient level e_g^* if g is close enough to 1.

Finally, it is worth highlighting that for the seller to exert effort above the minimum level then $p > c(\tilde{e})$ is a necessary condition. The seller must earn a reputational premium if a reputational concern is to discipline her behavior. This is the condition that expresses the “carrot and stick” intuition. A question that naturally arises is where do these reputational premia arise from and why do numerous sellers not enter the industry in order to exploit them. Klein and Leffler (1981) argue that a firm may need to expend resources by launching an advertising campaign, building a marble lobby, or buying the good name of an existing business, in order to ensure that the equilibrium played is the reputational equilibrium rather the equilibrium in which no buyer ever buys and the seller never exerts effort. Thus, there must be some initial investment which the firm is able to recoup over time so that the outcome can be consistent with free entry into the industry.

Although this seems anecdotally reasonable, there is much to be tied down and there is considerable multiplicity of equilibria and no reason why one way of dissipating profits should be any better or worse than another (or for that matter there is no reason why buyers should be swayed by such rent dissipation). Though not overcoming the critique of multiple equilibria, Shapiro (1983) introduces an equilibrium in which firms choose the price and initially need to price below cost and exert

effort to eventually earn a reputation premium — thus keeping the initial dissipation phase within part of the larger repeated game and without having to assume an additional means of rent dissipation.²

There are numerous ways to extend this basic model and consider the robustness of these insights. Much of the literature has focussed on a discrete effort choice and folk theorem results that demonstrate that a wide range of outcomes (indeed any outcomes that yield average payoffs to each player higher than they can obtain from quitting the relationship) can be sustained as the discount factor δ approaches 1. There are a couple of excellent recent reviews on repeated games. For a shorter introduction see Kandori (2006), and for a more comprehensive treatment see Mailath and Samuelson (2006). This literature has made a distinction between perfect monitoring where outcomes perfectly reveal action, imperfect public monitoring where outcomes are correlated with actions and observed by all buyers and the seller equally, and imperfect private monitoring where the seller does not observe the outcomes that buyers do. The setup we offer above falls in the sphere of imperfect public monitoring.³

² Allen (1984) also considers maintaining reputational premia through the repeated game, rather than introducing additional factors. Perri and Rasmussen (2001) and Hörner (2002) also consider reputation and competition but in models with both adverse selection and moral hazard as discussed in Section 5.

³ If we set $b = g = 1$ and let effort be discrete with $e \in \{0, 1\}$, $c(0) = 0$ and $c(1) \in (0, 1)$, then this will correspond to perfect monitoring.

5

Mixed Models: Hidden Types and Actions

Both aspects discussed above, hidden information and hidden actions, might be present simultaneously, where buyers must make conjectures on both the type of the seller and on the seller's likely action. This can be captured by letting $b < g$ and $\underline{e} \leq e_b^*$, and it is convenient to distinguish between classes of models. The first class is where the seller does not know her own type and is often referred to as “signal-jamming” models. The second class is where she does know her type and is often referred to as “signaling” models. Much of the literature reserves the term “reputation” for signaling models that include both hidden action and hidden (private) information.¹ In both these types of models, however, the seller takes actions in order to affect the resolution of uncertainty. As discussed at greater length below, the influence that the seller's action has on the resulting uncertainty is key to disciplining the seller's behavior and introduces a “reputational” concern.

¹ See, for example, Fudenberg and Tirole (1991), Cripps (2006), and Mailath and Samuelson (2006).

5.1 Seller Does Not Know Her Type

Though perhaps this is not the traditional order of presentation when discussing models of reputation, it is convenient to start by introducing a variant of the signal-jamming model of Holmström (1999). In this model the seller does not know her own ability but is motivated to exert effort so that future buyers would be more likely to observe a success and attribute this to high ability, leading to a high level of future compensation. Specifically, since buyers value success at 1 and failure at 0, in this model the price equals the buyers' expectation of success (for example this would be the price that arises if there are many buyers who Bertrand compete for the seller's service).

Suppose that there are only two periods for trade. The buyers must anticipate that the seller will exert no effort above \underline{e} in the second and last period. However, the seller may exert effort in the first period. If the market expects the seller to exert effort \tilde{e} in the first period then using the prior μ_0 , the first period price will be

$$p_1 = \mu_0 \tilde{e}g + (1 - \mu_0) \tilde{e}b.$$

The second period price will depend on the posterior reputation. Following a success in the first period, the second period reputation μ_1^S will be calculated using Bayes rule as follows:

$$\mu_1^S = \frac{\mu_0 \tilde{e}g}{\mu_0 \tilde{e}g + (1 - \mu_0) \tilde{e}b} = \frac{\mu_0 g}{\mu_0 g + (1 - \mu_0) b}. \quad (5.1)$$

Since $e = \underline{e}$ is anticipated in the second period, if a success occurred in the first period then the second period price following a success will be

$$p_2^S = \underline{e}[\mu_1^S g + (1 - \mu_1^S) b].$$

Similarly, following a failure, the seller's reputation will be

$$\mu_1^F = \frac{\mu_0(1 - \tilde{e}g)}{\mu_0(1 - \tilde{e}g) + (1 - \mu_0)(1 - \tilde{e}b)}.$$

and the corresponding second period price is

$$p_2^F = \underline{e}[\mu_1^F g + (1 - \mu_1^F) b].$$

Recall that since the seller does not know her own ability she anticipates succeeding with probability $e[\mu_0g + (1 - \mu_0)b]$. Therefore, a seller chooses effort e in the first period in order to maximize her expected discounted present value of earnings,

$$p_1 + \delta e[\mu_0g + (1 - \mu_0)b]p_2^S + \delta(1 - e[\mu_0g + (1 - \mu_0)b])p_2^F - c(e). \quad (5.2)$$

After some rearranging this maximization yields the following FOC,

$$c'(e) = \delta[\mu_0g + (1 - \mu_0)b](\mu_1^S - \mu_1^F)(g - b)e, \quad (5.3)$$

and in equilibrium the solution to this equation must conform to the market's expectation \tilde{e} .

A number of features arise. First, there is a unique equilibrium outcome here in which $\tilde{e} > \underline{e}$. To see this, assume that expectations were such that $\tilde{e} = \underline{e}$. The analysis above shows that for any \tilde{e} , $\mu_1^S - \mu_1^F > 0$. Given that $c'(\underline{e}) = 0$, and $c''(\cdot) > 0$, (5.3) implies that the solution must have $\tilde{e} > \underline{e}$ because the right-hand side of (5.3) is positive, a contradiction. Second, similar to Section 4, equilibrium effort (or the strength of reputational incentives) increases with δ . Third, it can be shown that effort increases in $g - b$. One implication of this fact is that a mean-preserving spread in types (an increase in g and a decrease in b that leaves $\mu_0g + (1 - \mu_0)b$ constant) increases effort. This follows because first, outcomes are more informative and cause stronger updating, and second, positive updating is more valuable.² Finally, and related to the intuition on the informativeness of outcomes, consider the effect of μ_0 . If there is little uncertainty about the ability of the seller, e.g., μ_0 is close to 0 or 1, then $\mu_1^S \approx \mu_1^F \approx \mu_0$ and there is little effort (so long as neither b nor g are degenerate). Thus, effort requires sufficient uncertainty.

One can consider longer horizons, and, potentially, even an infinite horizon model. However, given the way our framework is set up, this would be a daunting task. To see this, notice that any effort in period t will affect the Bayes updating of buyers for all future periods. Thus, one needs to employ a modeling structure that would result

²We use a different model than the one used in Holmström's seminal paper. This corresponds to a decrease in the precision of the distribution of types h_η in Holmström's model.

in convenient recursive equations in order to characterize the solution, which is precisely what appears in Holmström (1999).³

It is worth discussing some of the features that are highlighted in Holmström's analysis. First, similar to the two period model, effort will diminish over time, but in a more "continuous" way. To see this, notice that in our two period model, there is a simple reason that effort declines: there is "nothing to prove" in the second period since there is no future. If, however, there are more than two periods, then in the second period, there is a reason to exert effort for the same reason that effort is exerted in the first period of our two period model. More generally, consider a T period model, and focus on a later period t' and an earlier period t , such that $1 < t < t' < T$. First, notice that in period t' , there are only $T - t'$ periods left, so the "bang for your buck" of effort in trying to influence the beliefs of buyers is smaller than that in the earlier period t since there are $T - t > T - t'$ periods left. A second effect that implies higher effort in earlier periods is that in period t' , buyers are better informed about the type of the agent (more learning had occurred), and as such, effort has less of an impact on the corresponding Bayes updating of reputation. As a result, when the model has more periods, effort is diminishing over time.

Interestingly when there are more than two periods it may be that in equilibrium effort in early periods is even higher than the first best efficient level of effort. The intuition is rather straightforward: by working harder the effect on reputation will last over a number of periods, and the combined value may be so high as to induce a high level of effort that exceeds first best.⁴

³First, the outcome is not dichotomous but rather continuous and can take on any real value. In particular, the outcome in period t is $y_t = \theta + e_t + \varepsilon_t$, where θ is the seller's ability, $e_t \geq 0$ is her effort and ε_t is a normally distributed noise with mean zero. By having θ also normally distributed with some mean m , the resulting Bayes updating of reputation takes on a very convenient recursive closed form representation. Notice that in Holmström's model, and in contrast to ours, effort acts as a substitute rather than a complement for ability (the two are additive), which also serves to simplify the analysis.

⁴Holmström's "career concerns" insights have been extended in several ways (Holmström's paper was originally published in a book in 1982, and hence the impact since). First, even when outcome-contingent contracts can be written, if reputation concerns already provide some incentives then the form or strength of incentive contracts should (and do) alter, as shown in Gibbons and Murphy (1992). Second, interesting applications of "signal jamming" have been used for industrial organization (e.g., Fudenberg and Tirole, 1986)

With infinitely many periods of Holmström’s model, the first effect described as the “bang for your buck” of effort does not arise since the seller always has an infinite future, but the second effect of diminishing uncertainty still operates since the type of seller is set at the beginning of time and as a result of perfect learning, effort converges to zero as time goes by. Holmström noted that for effort to stay bounded away from zero some uncertainty has to be present. One way to introduce this into the model would be to have the type of seller change over time — an idea that we will revisit later.

It is also worth contrasting the infinite horizon version of Holmström’s model to the repeated game models of reputation described in Section 4. Recall that with pure hidden actions there is a folk theorem that implies multiple equilibria. In Holmström’s model there is a unique equilibrium as described above.

As a final note in this section, we highlight that Holmström presents a second model in his seminal paper in which an agent can *costlessly* choose from a range of different projects. If the reward to holding a reputation is nonlinear in reputation then, depending on the reward function, the seller might have incentives to act in a manner that is more (or less) risk averse than might otherwise be warranted. In particular, suppose that below a certain reputation threshold the seller is simply never hired. This induces a convex return function, which in turn implies that relatively “conservative” or risk-averse choices will be made. Such behavior has been further explored theoretically (Holmström and Ricart-i-Costa, 1986) and empirically (e.g., Chevalier and Ellison (1999) and Hong et al. (2000)).

5.2 Seller Knows Her Type

When the seller does know her own type then in general different types may take different actions. Again, we analyze a two period model, where sellers will choose no effort above \underline{e} in the second period since any positive level of effort will not have future rewards. We suppose that buyers anticipate that a good type exerts effort \tilde{e}_g in the first period

and finance (e.g., Stein, 1989). More applications that are too many to cover here have been since developed.

whereas a bad type exerts effort \tilde{e}_b . As before, the first period price will be a function of these expectations as follows:

$$p_1 = \mu_0 \tilde{e}_g g + (1 - \mu_0) \tilde{e}_b b.$$

The second period price will depend on the posterior reputation. Unlike (5.1) where both types choose the same level of effort, here different types may choose different effort levels \tilde{e}_g and \tilde{e}_b , so that Bayes' rule leads to reputation updating as follows (following success and failure, respectively):

$$\begin{aligned} \mu_1^S &= \frac{\mu_0 g \tilde{e}_g}{\mu_0 g \tilde{e}_g + (1 - \mu_0) b \tilde{e}_b}, \quad \text{and} \\ \mu_1^F &= \frac{\mu_0 (1 - g \tilde{e}_g)}{\mu_0 (1 - g \tilde{e}_g) + (1 - \mu_0) (1 - b \tilde{e}_b)}. \end{aligned}$$

It follows trivially that second period prices are,

$$\begin{aligned} p_2^S &= \underline{e} [\mu_1^S g + (1 - \mu_1^S) b], \quad \text{and} \\ p_2^F &= \underline{e} [\mu_1^F g + (1 - \mu_1^F) b]. \end{aligned}$$

As a consequence, a type θ maximizes his expected utility given by

$$p_1 + \delta e \theta p_2^S + \delta (1 - e \theta) p_2^F - c(e),$$

which yields, after some rearranging,

$$c'(\tilde{e}_\theta) = \delta \theta (\mu_1^S - \mu_1^F) (g - b) \underline{e}.$$

The equilibrium is unique and has similar properties to those in the case of Section 5.1. Efforts (though now of both types) increase in the discount factor, δ , in the difference in ability of a good and bad type ($g - b$) and will fall to the minimal level if there is no uncertainty with regard to the seller's ability (that is $\mu_0 \in \{0, 1\}$). Recall that in our framework, the complementarity of effort and type yields a single-crossing condition, which implies that $\tilde{e}_g > \tilde{e}_b$.

In general, characterizing equilibrium for longer horizons of this model with known types is more difficult and especially so if there is a richer type space. Moreover, if the complementarity between type and effort is removed then multiple equilibria may arise to further

complicate the characterization of equilibria. This follows because without complementarities there is no structure that causes one type to prefer more or less effort than another, and the incentives are imposed only through off equilibrium beliefs. The literature has dealt with this problem primarily by simplifying the model to have only one type that has actions to choose from and has focused on a discrete effort choice that results in the need to solve for mixed strategy equilibria. We now discuss two specific forms that have been influential in the reputation literature.

5.2.1 The Imitation Approach

Consider the following special case of the stage-game of our framework where both types have equal skill, but where the good type is constrained to acting in the best possible manner whereas the bad type has discretion over how hard to work. In particular, let $g = b = 1$, $\underline{e}_g = 1$ and $\underline{e}_b = 0$. This implies that a good type will always succeed whereas a bad type choose the probability of success which will just equal her effort $e_b \in [0, 1]$.

This approach (with discrete effort $e_b \in \{0, 1\}$) was pioneered by the seminal “gang-of-four” papers, Kreps et al. (1982), Kreps and Wilson (1982), and Milgrom and Roberts (1982). The literature that followed has coined the good type who is committed to a certain action as a “commitment” type. Another common nomenclature introduced by Fudenberg and Levine (1989, 1992) is a “Stackelberg” type who is committed to the action that the strategic type would commit to if it were possible.⁵

Going back to our example, and considering two periods, the bad type would choose $e_b = 0$ in the second period. However, this would not be the case in the first period. Similar to the argument in Section 5.1, if buyers expect $\tilde{e}_b = 0$ in the first period then a success is an immediate indicator that they face a good type, which would imply that $p_2^S = 1$

⁵In this case the strategic type defined by $\underline{e}_b = 0$ would like to commit to e_b^* that solves $c'(e_b^*) = 1$, and following our assumption this would yield $e_b^* < 1$. If instead $e \in \{0, 1\}$ and $0 = c(0) < c(1) < 1$ then the strategic type would like to commit to $e = 1$ as in the gang-of-four model.

and $p_2^F = 0$. However, this is more than enough of an incentive to induce the bad type to choose some positive effort because there is a premium for a successful outcome, and since $c'(0) = 0$, there is no first order cost to choose a small positive level of effort.

Hence, there is a unique equilibrium $\tilde{e}_b > 0$ that is solved as follows. Since the good type always succeeds, it must be that following a failure $p_2^F = \mu_1^F = 0$. Following a success we have,

$$p_2^S = \mu_1^S = \frac{\mu_0}{\mu_0 + (1 - \mu_0)\tilde{e}_b},$$

and \tilde{e}_b will solve

$$c'(\tilde{e}_b) = \delta(p_2^S - p_2^F) = \frac{\delta\mu_0}{\mu_0 + (1 - \mu_0)\tilde{e}_b} > 0.$$

The comparative statics of this model are familiar: effort is increasing in the discount factor δ (the “carrot” is bigger) and it is also increasing in μ_0 . This version of the model has some other characteristics that have been expressed by Kreps et al. (1982) as well as other applications of this idea. In particular, as the number of periods grow, the future becomes increasingly important and the bad type loses more from a failure. As a consequence, if $c'(1)$ is finite, and if δ is high enough then in early periods the bad type will choose $\tilde{e}_b = 1$, thus fully imitating the good (commitment) type. Interestingly, fixing $c'(1) < \infty$ and δ close to 1, as the number of periods goes to infinity so does the number of periods in which the bad type fully imitates the good type.⁶ Notice also that in such an equilibrium with the bad type perfectly imitating the good type, there is no information generated in the periods where imitation occurs and buyers cannot learn anything about the seller’s type. Thus, uncertainty about the seller’s type can be sustained perpetually.

The result that with infinite or long enough horizons the strategic (bad) type will perfectly imitate the good type is special to the case of $b = g = 1$, that is, the case with perfect public monitoring, as shown by

⁶The idea of commitment types has been used by Fudenberg and Maskin (1986) to generate a “folk theorem.” Specifically, if one plays with the commitment type’s committed action, then one can support a whole host of behaviors by the strategic (bad) type. Fudenberg et al. (1994) extend these ideas to games with imperfect public monitoring.

Cripps et al. (2004). Consider the case of imperfect public monitoring with $b = g < 1$, $\underline{e}_g = 1$, and $e_b \in [0, 1]$. To see that full imitation in any period is impossible, suppose for contradiction that in some period t , $\tilde{e}_{bt} = 1$ so that the bad type is fully imitating the good type. In this case buyers learn nothing from the outcome of period t , which can be either a failure or a success since $b = g < 1$. As a consequence, future payments to the seller will not depend on the outcome of the period t transaction, and therefore there can be no incentives for the bad type to exert costly effort in period t , a contradiction to $\tilde{e}_{bt} = 1$. However, reputation effects will still play a role and effort can be arbitrarily close to full imitation in early periods of the game.⁷

5.2.2 The Separation Approach

In the subsection above, the bad type is strategic and imitates the good type, taking costly action to try to convince buyers that she is in fact the good type. A more recent literature, initiated by Mailath and Samuelson (2001), instead focusses on the opposite case, where the strategic type is the good type who takes actions to separate herself from the bad type. In our framework this can be represented by setting $b = 0$, $g = 1$, and $\underline{e} = 0$. Here a bad type is inept and will always fail whereas a good or competent type chooses the probability of success, which will just equal her effort $e_g \in [0, 1]$.

We begin by considering a two period example and as usual, in the second period the good type would choose $e_g = \underline{e} = 0$. However, now in contrast to the models described above this implies that there will also be no effort in the first period. In the second period, a seller known to be good will earn an identical fee to a seller known to be bad. There are no carrots and no sticks to discipline a seller and so the seller exerts no effort. A similar result, of course, applies with a longer fixed horizon. As in Section 4, like in any infinitely repeated game,

⁷Cripps et al. (2004) continue to show that despite the early reputation effects, eventually, in the very long run, reputation concerns are fully exhausted since the type will be almost surely learned. This result might at first glance appear to conflict with the “folk theorem” results of Fudenberg et al. (1994). However, notice that the folk theorem results are about average payoffs, while the Cripps et al. result is about behavior in the very long run, which when discounted back will have a negligible effect on the present value of lifetime earnings.

when there is no terminal period then a folk theorem result implies that practically any payoff can be sustained in equilibrium. However, Mailath and Samuelson (2001) show that when restricting to Markov Perfect strategies, the only equilibrium is one in which no effort is exerted in every period.

This might suggest that when reputation is a concern for separation rather than imitation then there is no scope for reputation to discipline the seller's behavior. This is not generally true. For example, if we modify the model slightly so that $\underline{e} > 0$ then a good type exerting minimal effort \underline{e} is better than a bad type who can never succeed, and as a result there is an incentive for a good type to prove herself to be good, even if after having done so she will continue to exert no effort above $\underline{e} > 0$ in the subsequent period.

To see this, note that since the bad type cannot succeed, we have $\mu_1^S = 1$. When buyers anticipate that a good type who chooses \tilde{e}_g in the first period fails with probability $1 - \tilde{e}_g$ (recall that now $g = 1$), Bayes' rule implies that

$$\mu_1^F = \frac{\mu_0(1 - \tilde{e}_g)}{\mu_0(1 - \tilde{e}_g) + (1 - \mu_0)}.$$

Since a good type will succeed with probability \underline{e} in the second period, it follows that first period equilibrium effort \tilde{e}_g solves

$$c'(\tilde{e}_g) = \delta(p_2^S - p_2^F) = \delta \left(1 - \frac{\mu_0(1 - \tilde{e}_g)}{\mu_0(1 - \tilde{e}_g) + (1 - \mu_0)} \right) \underline{e}.$$

Taking this model to many periods will result in an analysis that is reminiscent of Holmström (1999) model. In particular, outcomes must be informative about the type of the seller for them to provide incentives to exert effort. In an infinite horizon, it cannot be that a good type exerts substantive effort for many periods. As the uncertainty about her type diminishes due to the learning of buyers, so do incentives to exert effort. Notice also, that the more effort the seller exerts in earlier periods, the stronger is the Bayesian learning that buyers perform. This is in contrast to the imitation approach of Section 5.2.1 in which more effort in earlier periods slows down the learning of buyers, which was key in sustaining high effort for many periods in the imitation approach models.

If reputation concerns are to be sustained indefinitely then some mechanism must ensure that buyers never fully learn the type of the seller. One way to achieve this, as explained in Mailath and Samuelson (2001), is by adding uncertainty that is constantly replenished. This can occur if a seller's type might exogenously change in an unobserved way so that in every period the type of seller may revert from good to bad or bad to good (similar setups have been used by Holmström (1999), Benabou and Laroque (1992), and Phelan (2006)). In application, for example, the environment might change and a seller's skill might become obsolete. An alternative to changing types, which is being explored in a very recent stream of papers is endowing buyers with finite memory (Monte, 2007; Ekmekci and Wilson, 2007) or with some positive costs of observing history (Liu, 2007) as a means of precluding buyers from learning fully and so perpetually sustaining the uncertainty that ensures that reputational incentives do not disappear. Bar-Isaac (2007), explores the institution of partnership and team production as an endogenous means of reintroducing type uncertainty and sustaining reputation incentives. Finally, there is a recent literature on name trading (in particular Tadelis (2002) and Mailath and Samuelson (2001)) that speaks to this issue and is outlined in the next section.

6

Trading Reputations

Most studies of reputation captured by the simple framework developed earlier consider one seller with a finite or infinite horizon of activity. This approach very much views reputation as attached to a certain identity, or individual. It is clear, however, that reputations are often associated with entities such as firms, teams, and organizations in which individuals may be replaced with other individuals, thus creating a wedge between the identities of the actual sellers and the entity that carries the reputation in the market.

A first attempt to capture this idea in the spirit of hidden action models of Section 4, was made by Kreps (1990). Consider the simple case of $b = g = 1$ and $e \in \{0, 1\}$ with $0 = c(0) < c(1) < 1$. In any finitely repeated version of the game there will be no trade because of the obvious “unraveling” reason. There can be trade (or in Kreps’ term “trust”) with an infinite sequence of this game. This is true whether there is a single buyer, or a sequence of short-lived buyers who can observe the history of trades (or indeed just the outcome of the previous trade) by using a trigger-strategy equilibrium as discussed in Section 4.

Now, what happens if our long-lived seller is replaced by a sequence of short lived sellers? A first guess would be that short-lived sellers will

have no incentive to perform well. Kreps, however, demonstrates that reputation can become a tradeable asset that provides incentives even when buyers and sellers live for only one period. The argument is simple and appealing: a “name” will be chosen for the firm by the first seller. Each subsequent seller will be trusted by the buyer of that period if and only if trust was never abused in the past, *and* if the current seller acquired the firm’s “name” from his predecessor. Sellers’ strategies will be to buy the “name” and honor trust if and only if trust was never abused. Thus, a seller will be able to sell his own good name if and only if he himself honors the trust of the client. If the loss from not being able to sell a name outweighs the benefits from abusive behavior, then sellers will have incentives to honor trust, and cooperative behavior is sustained in equilibrium. Cremer (1986) presents a related model, where rather than honoring trust in order to sell a good name, young agents honor trust in order to sustain a “culture” from which they will benefit in the twilight of their careers.

Like the pure hidden action models, the approach suggested by Kreps, and the related work of Cremer, has some difficulties. As Kreps writes, “The reputation construction is decidedly fragile: If reputation works only because it works, then it could fall apart without much difficulty. In real life, these risks will appear as substantial costs of undertaking transactions in this way” (p. 111). Furthermore, the bootstrap equilibrium in the repeated game approach fails to illuminate the process by which names become valuable after good performance, a process that is well documented in reality. Hence, it seems that some natural dynamics of reputation building is missing from this approach. Two related papers, Tadelis (1999) and Mailath and Samuelson (2001) try to illuminate the way in which reputations are built, and more importantly, why they may have value as a tradeable asset that is separated from the identity of the agents who are behind the firm’s activities.

In Tadelis (1999) sellers are either good or bad, and within our framework he assumes that $e = 1$, and $0 = b < g < 1$.¹ Departing from most of the reputation literature, sellers live for two periods but a

¹ Assuming that $b = 0$ is not necessary, but it is very convenient.

new cohort of sellers enter the economy every period, creating an overlapping-generations demography of sellers, while a new cohort of buyers who live for one period enter every period. The basic model has only two periods in total: in the first there are “old” sellers and “young sellers,” the latter who will continue to be active in the second period as “old,” and the former who will retire. Then there is a new “young” cohort who will be active in the second period, after which the economy ends. Clients are plentiful, and hence each seller is able to charge the expected benefit, or $p_t = \Pr\{\text{success}|t\} = \mu_t g + (1 - \mu_t)b = \mu_t g$.

The twist is that sellers are represented by names, and that sellers who are active in the economy in the second period can either choose a name or buy a name from a seller who wishes to sell her name (e.g., sellers from the initial old generation who will retire.) A name is just associated with the history of success or failure that has been generated by it. Hence, μ_t becomes a function of what clients believe about what types of sellers are buying names and what types of sellers are continuing with their names.

Tadelis (1999) offers two main results. The first is that if clients cannot observe whether or not trading of names occurs, then trade of names *must* occur in *any* equilibrium. The reasoning is simple: if clients were to believe that no names are traded in equilibrium, then successful names in the second period must have been generated by good type sellers who succeeded in the first period and then continued with their successful name. But then, any new seller would rather have an old successful name that she can get from one of the successful retirees, and would be willing to pay a positive amount for it. As a result, an active market for names (which are the vehicles for reputation) must be part of any equilibrium.

The second main result requires more than two periods² and shows that there cannot be an equilibrium in which good sellers can use the market for successful names to fully separate themselves from the bad sellers, which is somewhat counter to the intuition that is generated by pure hidden action models. For example, in Kreps’s model,

²Tadelis (2003) extends the 3 period, two-type model of Tadelis (1999) to include a continuum of types and infinite horizons. Given that sellers are active for two-periods, the results can be derived without relying on clients observing infinite histories.

sellers will buy a good name only if they intend to maintain it, which in a hidden information model may suggest that good types value good names more than bad types because it is easier for them to maintain the name. Tadelis calls this the “Reputation Maintenance Effect.” Another effect, however, is present: it is easier for good types to build their own name. As a result, good types value an existing good name less than bad types who cannot build a name for themselves. Tadelis calls this the “Reputation Start-up Effect.” It is shown that if only good types buy names then market expectations cause the Start-Up effect to overcome the Maintenance effect, which in turn causes bad types to value names more than good types do. The intuition is that if it is very likely that a good type buys successful names, then clients cannot update too strongly against such a name if it then fails, since good types do sometimes fail. This in turn implies that bad types will not be penalized too much from buying such a name and failing, and as a consequence, are willing to pay more than a good type are because their outside option of not buying a name is rather bleak.

Mailath and Samuelson (2001) offer a repeated game model in the spirit of the separation approach discussed in Section 5.2.2 that shows that both the start-up and maintenance reputational effects discussed carry over to their framework, which combines both hidden action and hidden information.³ Tadelis (2002) introduces a mixed model that includes hidden action into the basic hidden information set-up of Tadelis (1999) and shows that the two main results described earlier go through, and a third result emerges. Namely, the career concerns of young sellers from future increases in their payments offer the same quantitative effect that name-selling concerns have on old sellers. The intuition is that the price in the market for names is endogenous and tightly tied to the price differential in the product/service market that successful histories gain over unsuccessful ones.

³They do show that bad types are likely to value a very good reputation more than good types. Their partial equilibrium analysis requires an exogenous assumption that good types have a better outside option than bad types, which arises endogenously in Tadelis (1999, 2002).

When name-trading is unobservable then the identity of the seller is separate from the entity to which reputation is attached.⁴ Similar effects can arise when the identity of a seller, or his contribution to performance, is at least partially observable, as occurs in teams. We discuss such effects, in the context of our broader discussion of the effect of organization design in Section 9.3. A closely related idea that different identities are linked through some common entity arises in the context of co-branding a number of products and is discussed in Section 9.2. Further, transfers of names which are fully observable can still play an informative role, when one type of seller is more likely to buy a good name than another, as endogenously arises in Hakenes and Peitz (2007), Deb (2007), and Wang (2007). Finally, Marvel and Ye (2008) offer an extension of Tadelis' model to trademark sales and introduce endogenous and costly entry to explore the welfare effects of allowing trademark sales.

⁴In Tadelis's papers, the first result stating that names must be traded in every equilibrium relies on a sufficient amount of non-observability. With complete observability name trading is an equilibrium, but there are also equilibria with no trade of names. The no-separation result carries over even with complete observability.

7

Exogenous Factors Influencing Reputation

The building blocks of Sections 3, 4, and 5 and the intuitions that arise have been painted in a fairly broad brush. The repeated games literature introduced in Section 4 has mainly focused on folk theorems and allows for a multiplicity of outcomes. The mixed models of Section 5 can rely delicately on the assumptions made on the possible types that the sellers might take,¹ and the early literature focussed on possibility results and how type uncertainty can lead to reputation effects with a finite horizon. For applications and empirical work, however, there is considerable interest in identifying factors that influence the dynamics of reputation and the comparative statics of reputation effects: the circumstances in which concerns for reputation are more pronounced or more muted. These factors rely on the fundamental forces identified and discussed above, such as the discount factor, the extent of uncertainty about types and, in general, the relative size of the “carrot and stick”

¹ Abreu and Sethi (2003) introduce an evolutionary model to endogenize commitment types in a symmetric model of reputation in bargaining. However, papers generally take the possible types available as a primitive of the model. In practice, and in empirical work, these are primitives that can be hard to identify.

for different behavior and outcomes. Nevertheless, it is worth showing how these fundamental forces play out in application.

Much of the discussion here is closely related to our discussion on empirical work on reputation that appears in Section 11, since empirical work naturally requires sources of exogenous variation. Broadly speaking, one might imagine two classes of factors that might introduce exogenous variation. First, there are factors that influence the returns to having a good or bad reputation. These form the focus of the discussion in this section, and in particular we highlight life-cycle effects, the size of the seller's business, and the competitive environment that the firm faces. In addition reputation effects depend on the current reputation, a theme that is picked up in particular in some recent work on reputation in online trading environments, and, as discussed widely in Macleod (2007), on the efficacy and nature of formal contracting. Second, there are factors that affect the ability of buyers to observe and gather information, which in turn will have implications on the value of building and maintaining a reputation.

7.1 Life-cycle and Dynamic Effects

The models described in Section 5 share a number of features. Recall that reputational incentives to exert effort decline with time for two reasons. First, there is less "bang for the buck" in investing in reputation when the number of future periods remaining to enjoy the reputational benefits diminish. Second, since uncertainty diminishes over time buyers have less to learn and current seller actions have a smaller effect on future beliefs. In particular, this suggests that sellers who are older and more established, and/or are closer to retirement should have diminished incentives to exert costly effort. That the reputation incentives of sellers might differ through a seller's lifetime yields an observable source of variation (age, experience) that has been empirically examined, as we discuss in Section 11, particularly, in the work on career concerns of managers and analysts. However, the insights in these building block models are relatively simple inasmuch as the models focus squarely on either an imitation incentive or a separation

incentive, and the rewards that a seller anticipates are linear in the market's expectation of the seller's performance.

Diamond (1989) explores lenders (equivalent to our sellers) whose reputation will determine whether they get loans and at what interest rates, and offers a model that includes both incentives for separation and for imitation. First, there are three potential types of sellers (a "bad" commitment type, a "good" commitment type, and a strategic type) and so potentially both types of reputation concerns are present. For sufficiently low reputations lenders would simply not extend loans, and so, a borrower's value is nonlinear in expected performance. As a consequence, when initially starting with a low reputation, a strategic borrower will undertake a risky (equivalent to our low-effort) project. If successful, after which reputation is built as a consequence of luck, the seller may undertake a safe (equivalent to our high-effort) project in order to protect her enhanced good reputation. However, in late periods with nothing much to lose, the seller takes her preferred risky project. Thus the incentives to take a safe rather than risky project can change non-monotonically over time, in contrast, for example, to the simpler models discussed in Section 5.

Similarly, Benabou and Laroque (1992) present a model in which the reputation concerns of a seller (or in their application an informed trader such as a guru, manager or journalist) can fluctuate over time. They present a rich model where the seller makes costless cheap-talk pronouncements which might reveal information (and thus this is a model of reputation for expertise, as described in Section 10.1) and can trade an asset whose price might be affected by her pronouncement and for which she has some private information. The informed insider trades-off a current benefit from manipulating the market to her advantage, with the reputational cost which would prevent her from exploiting other traders in the future. Benabou and Laroque (1992) show that for low reputations, the insider seeks to build up her reputation by reporting her information relatively honestly, but if her reputation is sufficiently high she prefers to milk it by lying. Since, this is in effect a model of imperfect monitoring, this can lead to cycles and phases where there is relatively more investment in reputation, or where the insider milks her reputation; however, in the long run

the type of the agent can be statistically determined, as described in the last paragraph of Section 5.2.1, and so eventually the truth will emerge and reputational concerns will disappear. If, however, the type of the insider can change exogenously, then cycles of “investment in reputation” and “milking reputation” continue indefinitely. Gale and Rosenthal (1994) also offer an analysis where firms build up their reputation only to later “milk” it down when exogenous shocks to the environment occur.

7.2 Competition

Much of the (particularly early) literature focuses on a monopolist seller. However, the nature and extent of competition can influence the value of having a good reputation and the strength of reputation concerns. In effect, competition shapes both the size of the “carrot” and of the “stick” of market discipline that result from changes in reputation.

We have already mentioned in Section 4 that Klein and Leffler (1981) and others have suggested an early phase of rent dissipation that ensures that a reputational premium can be sustained even if there is free entry into a market.² Hörner (2002) suggests that competitive discipline ensures that firms continue exerting effort as any failures will drive them out of the market. There are rewards to developing or sustaining reputation since buyers make inferences on a seller’s ability from the price offered so that a reputational premium is sustained.

When buyers make no inferences from prices and when the industry structure is exogenously fixed, competition can erode this reputational premium as shown in Kranton (2003) so that too much competition diminishes reputation effects. Some competition, however, can help: for example, a monopolist who does not face the threat of competition, might do little to develop a reputation, confident that it would lose no buyers (even though it must charge a lower price if losing a reputation for providing quality). This discipline of competition, coupled with its erosion of the reputational price premia, can lead competition

²The initial rent dissipation exactly offsets the value generated once the reputation is established and so firms are indifferent between entering the industry and staying out.

to have non-monotonic effects on the strength of reputational concerns, as outlined in Bar-Isaac (2005).

More generally, the nature of competition might affect the rewards to reputation. Consider the function that maps buyers' beliefs about a seller to what they would pay for her service. This function clearly depends on the nature of competition. Under Bertrand competition with no differentiation (beyond quality) and no capacity constraints, superstar effects in the spirit of Rosen (1981) arise. Only the best seller would earn profits, leading to a much more convex reward-to-reputation function than the case when competition is not as fierce (for example when competition is Cournot or buyers have loyalties to sellers independent of reputation) as discussed in Harstad (2007). Similarly, heterogeneous demand, arising from different jobs or employers with different productivities, can lead to a convex return function, as discussed by Casas-Arce (2005), Martinez (2008), and Kovrijnykh (2007). The convexity of the return function affects the strength of reputation concerns and, in particular, means that the reputation concern depends on the current reputation.

A further effect that depends on the industry structure is that when a number of firms are involved in an industry, a price premium can be sustained through collusive agreements or equilibria. Dana and Fong (2007) explain that an oligopolist might be more concerned to sustain quality than a monopolist, since the long-run cost of shirking on quality is higher because firms expect both an adverse impact on buyers' expectations about future quality and an adverse impact on cooperative pricing.

7.3 Size of the Firm

A small literature has highlighted that the size of the firm might affect the rate at which potential buyers learn about the behavior of the seller. In particular, Rob and Fishman (2005) introduce an infinitely repeated hidden action model where in each period many new buyers arrive and sellers can service many customers simultaneously (but in each period must devote identical effort in their dealings with each buyer). Every new buyer speaks to an old buyer and learns his experience before

deciding on which firm to approach. It follows that when a firm is bigger, it has a larger buyer base, and so, new buyers (who are relatively more likely to meet an old buyer of a large seller) are more likely to hear about successful or failed transactions of a large seller than a small seller. Since such information is disseminated more rapidly, the carrot and sticks disciplining a large seller are relatively high. The reputation concerns (and induced efforts of a large firm) will be greater. Rob and Fishman (2005) then build on this idea to consider firm dynamics. Loosely, and as seems anecdotally plausible, once firms “make it” and are large enough, they are likely to “stick around” due to the stronger reputational concerns.³

³The work of Cai and Obara (2006) is somewhat related in motivation. In their model all buyers observe the outcomes of all transactions and a seller can take different actions with respect to different buyers in the same period but there is imperfect monitoring (so that $b = g < 1$). Their idea is that although imperfect monitoring means that the outcome of any single transaction may be a success or failure, there is much less uncertainty about the distribution of success and failures in a very large number of outcomes. This suggests that it is easier to sustain an equilibrium in which the seller exerts effort in all transactions and all periods when the firm is larger. Cai and Obara (2006) assume that in each period the firm receives some shock that affects outcomes in all transactions in a given period. This ensures that even if a seller is servicing infinitely many buyers, there is still some residual uncertainty about the overall distribution of outcomes. As a result, the seller is tempted to cheat a few of her customers and the optimal firm size for sustaining a high effort equilibrium, may be finite.

8

Reputation and Institutions

Seller reputations develop in the context of broader markets or societies. Institutional features play an important role in creating and sustaining reputational concerns. In particular, social and market institutions affect the rate at which information is disseminated and used by buyers, as well as the efficacy of rewards and punishments. We begin by reviewing the effect of social institutions on reputation and then turn to a specific application of a designed environment — e-commerce — where features of market design affect the efficacy of reputation and efficiency of trade.

8.1 Social Institutions

Information is at the heart of all of our models of reputation. Buyers must observe outcomes if they are to update beliefs, or understand that they are to punish inappropriate behavior, and hence buyers' information about a seller and her previous behavior play an important role. To this point we have implicitly been assuming that all buyers observe the full history of past transactions.¹ In practice, however, this

¹An exception is in our brief discussion of the social aggregation of information and the possibility of information herding at the end of Section 3.

is a strong assumption to make and it is of considerable interest to understand how reputational incentives are maintained depending on the nature and extent of information that is available to the community as a whole.

Much of the theoretical literature on community enforcement has focussed on the repeated game or pure hidden action approach to reputation. Moreover, it has considered trust or cooperation more broadly, making no distinction between the buyer and seller side of the transaction. Indeed this literature has typically focussed on the prisoners' dilemma and has either investigated folk theorem results, or explored what outcomes can be sustained for sufficiently high discount factors. Hence, it is not so much about seller reputation, but is close enough in methodology to warrant some discussion in this survey. There are certainly contexts where reputation and trust are important for both sides of a transaction; for example, one side of a transaction may be concerned about receiving timely payment and the other side concerned about receiving timely delivery of quality goods.

One branch of this literature (Kandori, 1992; Ellison, 1994; Harrington, 1995) has focused on the extreme case in which parties to a transaction observe only the outcomes of the transactions in which they have been directly involved. Even with this minimal amount of information diffusion, cooperation can be maintained as long as the community contains a finite number of members and the discount factor is sufficiently high (i.e., there is a folk theorem). Loosely, bad behavior by any agent causes his transacting partner to behave badly and similarly his future partners will behave badly, and then their future partners and so on. In this way, such bad behavior contagiously spreads through the population. Since the population is finite, the original misdeed will eventually cause distrust through the entire community and come back to haunt the original perpetrator.

Another branch of the literature (Kandori, 1992; Okuno-Fujiwara and Postlewaite, 1995) allows for infinitely many members in a community and assumes that each member is labeled with a status or reputation which is observable to the party with whom he is transacting. The way that the reputation or status changes depends not only on the actions and outcomes in a transaction but also on the

transacting partner's status (in particular this can allow a member to gain credit rather than suffer a drop in reputation for acting badly and thereby punishing a member with a bad reputation).

More generally, the development, efficiency and stability of institutions is a vast and burgeoning literature and a field in itself encompassing both theoretical analyses, as well as applications in historical and contemporary studies. It is well beyond the scope of this review to address this field, but interested readers are referred to the works of North (1990), Dixit (2004), and Greif (2006). Articles more closely related to reputation and community enforcement include Greif (1989), Greif et al. (1994), Milgrom et al. (1990) and Richman (2006).

8.2 Electronic Marketplaces: A Designed/Shaped Institution

One of the earliest and well known applications of reputation in online markets is eBay's "Feedback Rating," which records feedback that sellers (and buyers) receive from their trade partners in previously completed transactions (and now there are many others, e.g., "Amazon Marketplace.") While, the literature on Social Institutions described above in Section 8.1 can be viewed primarily as a positive exercise in describing social phenomena, the new electronic marketplaces have emerged and led to a greater focus on normative questions precisely because these electronic marketplaces are very much designed. An important question is, therefore, how should they be designed optimally *vis-à-vis* reputation tracking in order to facilitate trade?

While, reputation, trust and limits on information can act as impediments to trade in a bricks-and-mortar environment, in electronic markets where buyers and sellers can trade anonymously, these problems might seem even more severe. The market designers (such as eBay, Yahoo! Auctions, Amazon Marketplace and others) can make decisions that affect the information available for the parties considering trade to help mitigate these problems. This can range from allowing sellers to volunteer more information (whether in text, pictures, links to external sites and so on), allowing people to submit quantitative and qualitative feedback information, constraining the kinds of information that they

can leave (for example, one could consider a 3 point scale, a 5 point scale, a continuum, feedback on different aspects of the transaction and so on), summarizing the information collected in particular ways (prioritizing the most recent information, censoring information, prioritizing feedback by value of transaction or the “reputation” of the person leaving the feedback) and constraining who can leave information (only those who have transacted, or potentially other recommenders).

This field is, to some extent, is still in its infancy and the literature to date, which is well summarized by Dellarocas (2006) and Dellarocas et al. (2007), has applied the logic of the models outlined in Sections 4 and 5 above to demonstrate the role that reputation mechanisms can play in disciplining seller behavior in an online environment. The literature also provides opportunities to empirically explore such effects as discussed later in Section 11. Some features that are particularly prominent in the electronic marketplace applications are worth describing further.

First, reputation in electronic exchanges tends to rely on buyers’ reports about the seller. A central concern, is eliciting such feedback from buyers, and doing so in a way that is useful (for example avoiding “shills” who fallaciously build up a great reputation) and perhaps consistent (do all buyers have the same understanding of what constitutes a “good” or “bad” transaction). For example, Avery et al. (1999) consider paying buyers to provide information, and Miller et al. (2005) propose a mechanisms for eliciting honest feedback in environments with pure adverse selection, whereby those leaving the feedback are rewarded on the basis of future reporters’ feedback. More generally, one might consider the reputation of those leaving the feedback, and rating the raters. In doing so, however, one must be wary of perverse incentives to leave positive feedback in order to avoid retaliation that itself penalizes raters (see, in particular, Klein et al. (2006) and Dellarocas et al. (2007)).

Second, in a number of papers summarized in Dellarocas (2006), Dellarocas and others consider the granularity of a rating system, the frequency with which it is updated, and the extent to which older outcomes should be discounted. One feature that arises is that there may be a benefit in a system which “forgets” sufficiently old outcomes.

The intuition is in line with the discussion in Section 5, curtailing feedback can perpetually maintain uncertainty about the seller's type and so prevent the asymptotic learning that will eventually wipe out reputation incentives. Furthermore, in online environments one could, in principle at least, use additional information beyond feedback to construct ratings. For example, one could use characteristics of parties in past transactions or characteristics (such as price) of past transactions.

9

Strategic Manipulation of Reputational Environments

To this point, we have focussed primarily on the role of reputation in influencing a seller's behavior within a transaction. We have argued that reputation, or career concerns, might lead managers to work harder, or sellers of goods and services to provide higher quality. In this section we highlight that reputation might also affect a seller's actions outside of a given transaction either to affect buyers' beliefs before the transaction occurs, to affect what buyers observe following a transaction, or the way in which buyers interpret observations.

9.1 Signaling

In addition to actions that the seller could take to affect the outcome of transactions with buyers, as described in Sections 4 and 5, the seller may be able to take other actions that do not affect the quality of the good, but influence the buyers' beliefs and their propensities to buy. The classic Spence (1973) model of education as a signal is relevant. Consider a niche consultant specializing in regulatory advice. This person might choose to obtain a PhD in economics, not only to acquire some useful skills and training, but also to demonstrate through

a costly signal that he possesses the willingness and capacity to devote himself to tedious and careful work of a particular sort.¹

Another example for costly signaling can be the hiring decisions of firms when there is uncertainty about the firm's product's inherent quality, but not about its workers. For example, Hollywood movie-makers clearly take into account in choosing the cast that potential customers will be influenced by this choice in deciding whether or not to see the movie; similarly, academic departments and professional services might display an intention that they are changing strategy with a high profile hire. In general such costly signals are meaningful inasmuch as they have more value for one type of firm than for 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 for either a "good" or "bad" firm, but the benefits differ, creating an appropriate single-crossing condition that make signaling both credible and valuable.

Typically this arises because customers may buy repeatedly if they have a good experience. 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), but they will only buy again if the firm performed well. This accounts for the single-crossing condition necessary for informative signaling: attracting customers with an expensive signal will only pay off if they return for future profitable transactions.² Note that this signaling value 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.³

Perhaps, the clearest example of seller actions outside of the transaction to affect buyer beliefs is costly advertising. For example,

¹ In general, signaling models can allow for multiple equilibria. A long (and somewhat unsatisfying) literature has developed in order to offer refinements and selection of appropriate equilibria. See, for example, Section 5 of van Damme (1992) and Kreps and Sobel (1994) for reviews of this literature.

² 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 not to see it.

³ The notion that wages reflect not only productive capacity but also the opportunity for others to establish reputation has been recently explored in a somewhat different context by Anderson and Smith (2006) and Almeida Costa and Vasconcelos (2007).

Milgrom and Roberts (1986) discuss the use of advertising and low pricing as money-burning signals when customers might buy in each of two periods. By pricing low, or spending on advertising, the firm loses profits in the current periods, but a high quality firm would be confident that it will be able to charge higher prices to customers who were encouraged to try the good in the first period and discovered it to be high quality. Thus the appropriate single crossing property is satisfied and the signal is meaningful: a high quality seller gains from the costly signal of advertising or low-pricing, whereas a low quality seller does not. An excellent and comprehensive review of signaling models of advertising is contained in Bagwell (2007a).

9.2 Linking Transactions

Firms can also affect buyers' expectations and the speed and manner in which buyers update these expectations by, in effect, informationally linking different transactions. In particular, firms can choose to market different products under the same brand name (the literature has variously termed this reputation-stretching, brand stretching and umbrella branding). Wernerfelt (1988) concisely summarizes the underlying intuition 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 buyers to pool their experience with the two products to infer the quality of both.

In effect, by linking different transactions, the choice to make this link can act as signal. Moreover, it causes buyers to have more frequent interactions with the brand's seller, leading to faster learning and an effectively shorter discount factor. Different papers have focussed on these different aspects (and sometimes on a combination of these factors).

For example, Choi (1998) presents an infinitely repeated game where a firm chooses whether or not to umbrella brand a new good in each period; the firm can "behave well" by extending its name only to good new products, or "cheat" by releasing a bad new product under

the brand name. The promise of future returns can discipline the firm to extend its name only to good products, as in Section 4: in effect the firm develops a reputation for extending its brand name only to good products.

Wernerfelt (1988) presents a signaling model in which the appropriate cost of signaling (which ensures its credibility) arises as buyers have only partial information about the original product. After experiencing the new good (and thereby potentially gaining new information), there is an additional trading opportunity for the original good. Cabral (2000) extends this hidden type approach to umbrella branding, assuming that the two products are necessarily of the same quality and builds on some of the ideas from Tadelis (1999). Cabral highlights that following the consumption of the first good in the first period, there is a direct reputation effect, where the experience informs expectations of future consumption of this good; consumption of the second good informs expectations under umbrella branding, since consumers understand that these goods are of the same quality — a reputation feedback effect; finally, since in equilibrium the firm does not choose to umbrella brand with probability one, there is information in the observation of an umbrella brand that entails a signaling effect. Miklós-Thal (2006) allows the firm to choose the extent of correlation of the quality of the two goods. Andersson (2002) and Cabral (2008) consider firms that choose quality, as well choosing price, and who must decide whether or not to umbrella brand in infinitely repeated environments. The idea of linking different transactions between buyers and the firm owes much to the analysis of Bernheim and Whinston (1990), developed in the context of collusion. Dana and Spier (2007) consider linking transactions through bundling rather than co-branding and assume that buyers observe outcomes on their own purchases rather than realizations for all buyers in the market (imperfect monitoring). Finally, Hakenes and Peitz (2008a) also consider hidden action but in a finite horizon model where the hidden action problem is of a somewhat different nature: a firm has to decide on a level of investment as a fixed cost in product quality that determines quality in all periods of trade.⁴

⁴Other papers on umbrella branding include Aaker and Keller (1990), Montgomery and Wernerfelt (1992), Peppall and Richards (2002), and Hakenes and Peitz (2008b).

The phenomenon has also been empirically explored in the marketing literature by Erdem (1998), who uses panel data for two oral hygiene products, toothpaste and toothbrush, and by Balachander and Ghose (2003) who use scanner data for yoghurt and detergents. Both find some support for the idea that umbrella branding plays a role in conveying information. Sullivan (1990) considers two event studies showing that defects in the Audi 5000 reduced demand for other Audi models, while a new model design by Jaguar boosted demand for older models. There is additional evidence in Sappington and Wernerfelt (1985) who argue that umbrella branding may reduce uncertainty about a new product's attributes (and thus relates to "horizontal reputation" as discussed in Section 10.2), which increases value if buyers are risk averse.

While most of the discussion above concerns a single firm with multiple goods choosing whether or not to market these goods under the same brand name, one can easily imagine linking transactions between different firms. For example, a well-reputed firm may co-brand or directly certify less-established firm. Such co-branding has been explored, for example, in Garella and Peitz (2007) and Choi and Jeon (2007). Using the reputation of an intermediary to certify uncertain quality of senders is implicit even in static models of certification (such as Lizzeri (1999)), while dynamics and the reputational concerns of intermediaries have been explicitly considered in Chu and Chu (1994), Biglaiser (1993), and Biglaiser and Friedman (1994). In these papers the "future" comes sooner for the intermediary, and as a result reputational effects might discipline the intermediary despite not being as effective for individual sellers; interestingly, increasing returns to scales suggest limits to competition in market for intermediation.

9.3 Organizational Form and Personnel Policies

Organizational structures and policies can also affect buyers' beliefs and the ways in which they are updated. Indeed, choosing whether or not to admit an employee can function in much the same way as whether or not to stretch a brand name to a new product. For example, just as the Choi (1998) model considers a firm seeking to develop a reputation

for using its brand name only on good products, one could analogously think of a firm's role as certifying the quality of its employees. The firm maintains a reputation for hiring only good employees, and should it ever lose this reputation (which it may be tempted to by employing cheaper, low quality employees) then it loses the option value of future profitable hires. In this setup, the firm acts as an intermediary between customers and the labor market, maintaining a reputation for selecting only good employees.⁵ Though not much discussed in the formal literature, this certification role seems of considerable practical importance in professional services industries.⁶

Highlighting the role that organizations can play in sustaining good outcomes in the face of short-run temptations to cut corners was introduced by Cremer (1986). Cremer considers an infinitely repeated environment with overlapping generations of employees (or sellers), where buyers worry about the effort of these employees as discussed in Section 4 (hidden action). Using the idea of profit-sharing across generations of employees, or partners, Cremer shows that the members of the group can "bond" their behavior to be good, and buyers' correctly anticipate this in equilibrium. Specifically, Cremer characterizes an equilibrium in which old-aged partners receive the share of their younger partners work only if they themselves performed well when they were younger. This perpetuates a norm of hard work among the younger partners, from which they will benefit when older.

The alternative problem of hidden information, rather than hidden action, can also be addressed using profit sharing mechanisms, as demonstrated by Levin and Tadelis (2005). They argue that economies of scale or other complementarities encourage firms to include many

⁵ For a thorough discussion of different aspects of the firm acting as an intermediary between customers and suppliers see Spulber (1999).

⁶ In addition, it is worth highlighting that while our discussion and the literature on reputation is concerned with "vertical" aspects; "horizontal" concerns may also play a role. A firm may wish to sustain a reputation for hiring a certain "horizontal" type or style of worker. Similarly, firms may constantly take actions in order to build a consistent "horizontal" reputation. For example, in the consulting industry of the 1960s, McKenna (2006, p. 154) argues that although "To outsiders competition among these elite management companies . . . emphasized the striking homogeneity among the elite firms more than any immediate differences," in fact there was considerable variety and the elite firms "emphasized their work in specific sectors of the economy to differentiate their services."

members, and that partners who are committed to share their earnings will fear the dilution of their partnership shares from admitting worse partners who bring in less value, hence ensuring that their partnerships are small and consist of relatively high quality members compared to the corporate form. They show that only when buyers have imperfect information about quality (using a hidden information model) will this commitment pay off in terms of higher profits and higher quality of service. Their explanation squares well with the observed patterns that profit sharing is prevalent in service sectors that are plagued by hidden information and where labor is the key component in production, whereas other sectors rarely have profit sharing mechanisms.

Morrison and Wilhelm (2004) and Bar-Isaac (2007) consider teams composed of senior partners and juniors. This structure of internal promotion, coupled with a production process that obscures individual contributions, can lead to equilibria in which members of the firm exert effort. In particular, in Bar-Isaac (2007), since it is hard to know whether the senior or junior is responsible for the failure of a project, a senior who exerts effort gives the junior an opportunity to establish herself. This can motivate the senior to exert effort in order to sell the firm to the junior at a higher price. The junior is motivated by the prize of partnership. Thus, juniors are motivated by concern for their own reputations and seniors by the reputation of the firm that they own. As described in Section 5, type uncertainty is crucial for reputation incentives. The institution of partnership and joint production introduces uncertainty which sustains reputational concerns. In Morrison and Wilhelm (2004) the hidden action problem is with respect to seniors' training and development of their employees. Again the senior's concern is to work to ensure that the junior will find it valuable to buy into the partnership.

Another potential role that organizational structure can play is to affect the speed with which buyers and (current or potential) employers can learn about the ability of employees. As highlighted above, the rate at which information on performance is generated and diffused will affect the strength of reputational concerns. This can explain the choice to staff production teams as combinations of younger and older workers as discussed in Meyer (1994) and Jeon (1996), the delegation of

power, as in Ortega (2003) and Blanes-i-Vidal (2007), and promotion policies as in Koch and Peyrache (2005). In an application that focuses on firm policies affecting priors, rather than rate and nature of learning, Bar-Isaac and Ganuza (2008) highlight that recruitment and training policies affect priors about the ability of current workers and so also the strength of their reputational concerns.

To be sure, when one considers an organization that is composed of many key employees, the reputation of the organization as a whole will be tightly connected to the individual performances of its members, and hence tied to individual reputations. Tirole (1996) offers a model of group reputation that is tied to the reputation of its members. The model is an overlapping generations model of group members, and it incorporates both hidden information and hidden action of these members. Tirole assumes that there is imperfect monitoring of individual behavior, but that group affiliation is perfectly observable, as is group outcome as a whole. Hence, the group aggregates many noisy signals to one perfect signal of the aggregate behavior. Imperfect observability of individual behavior thus underlies the phenomenon of collective reputation, and each member's welfare and incentives are thus affected by the group's reputation. As a result, the behavior of new members of a group depends on the past behavior of their elders.

10

Beyond Reputations for Quality

The literature on the economics of reputation is vast and it is impossible for any review to be fully comprehensive either about all the relevant theoretical considerations or empirical applications.¹ In this section, we review a number of theoretical ideas of interest that lie beyond the scope of what can be captured by the framework we have outlined in Section 2.

10.1 Reputation for Expertise

We have thus far highlighted factors that influence a seller's reputation to take costly actions that will improve the outcomes of transactions, generally perceived as offering high quality goods and services. However, there is an extensive literature that has addressed a seller's concern to establish a reputation not so much for delivering a high quality good or service but rather for having a certain expertise.²

¹ For example a search for “reputation and economics” on Google Scholar yields almost 70,000 matches and “seller reputation” close to 30,000.

² Papers in this literature include Scharfstein and Stein (1990), Harrington (1993), Zwiebel (1995), Prendergast and Stole (1996), Ottaviani and Sørensen (2006a,b,c), Levy (2004, 2007), and Prat (2005). Another concern, somewhat related inasmuch as the concern is not

For example, when the seller is selling advice, there may be no additional direct cost to giving some kind of advice rather than another. Advice may simply be cheap talk.³ In these models, buyers have no doubts about *which* action the seller has taken, but instead have trouble in determining whether it is the *appropriate* action.⁴

The knowledge or expertise that the seller holds can affect the quality of her advice. However, when such a seller is in an ongoing relationship with buyers, she may tailor her advice to influence her reputation as being more knowledgeable or expert. For example, if there is a particular advice that buyers anticipate is relatively more likely to be given by a more expert agent, then, even if the seller has conflicting knowledge, she might provide that piece of advice.⁵

Similarly, some advice or some choices, might be more likely to reveal expertise than others and so if returns to reputation are nonlinear (or if they are linear but the seller is not risk neutral) the seller has incentives that might be more revealing (or more concealing) of her expertise. This insight appears in the second model presented in Holmström (1999) which considers a manager's incentives to take more or less risky actions. Similar effects have been explored both theoretically and empirically, in particular in the portfolio investment decisions of fund managers (see for example, Chevalier and Ellison (1999)).

Experts models need not rely on the actions or advice being costless. For example, in Ely and Välimäki (2003) and the more general model of Ely et al. (2008), there are costs to different actions.

to undertake costly actions but rather "appropriate" actions is a concern to demonstrate similar preferences to consumer rather than differential expertise. See, in particular, Morris (2001), Maskin and Tirole (2001) and Benabou and Laroque (1992)."

³ Crawford and Sobel (1982) introduces a workhorse model of cheap talk which has received an extensive workout in this literature on reputation for expertise. An early and important paper in this respect is Sobel (1985).

⁴ In practice, both concerns may operate simultaneously. Bar-Isaac (2008) develops a model where an agent exerts effort (an unambiguously preferred action) to generate information which would help decide on appropriate course of action. Since the expertise concern can lead to particular choices which look "smart," this can lead an agent to disregard his information and so diminish his incentives for information-gathering. However if "smart" actions are relatively revealing of ability, the expertise concern in leading the agent to take more revealing actions can boost information-gathering incentives.

⁵ In particular, this will be the case if she is not certain of her conflicting knowledge and/or if the consequences of the advice are noisy or will take a long time to be revealed.

Consider the clever example offered by Ely and Välimäki in which a car mechanic has the option to fix a faulty car by either performing a tune-up, which helps minor problems only, or by performing a more expensive maintenance that fixes any problem. If some mechanics are crooks, then these will always perform the more expensive service even when it is not needed. How will an honest mechanic distinguish himself? By performing tune-ups even when these are not sufficient to solve the problem. The reason is that buyers have an expectation over which different actions the seller takes over time, and reputation concerns will cause a seller to try to match this anticipated frequency. Since the environment is stochastic (in terms of which problems arise and the solutions that best fit them), trying to match an expected distribution can be inappropriate, and the resulting behavior of the expert can be less efficient than in the absence of reputation concerns.

In this work, buyers' expectation over what behaviors are appropriate are, largely, exogenously given. However, much of the literature has been concerned with endogenously determining what buyers would perceive as "good behavior" particularly in the context of competition between sellers. When there are many different sellers, then each seller's behavior or performance is assessed in comparison to others. This can lead both to herding or conformism, as described for example by Scharfstein and Stein (1990) and Zwiebel (1995), but can also lead to anti-herding effects since a seller may wish to distinguish herself, as in Levy (2004).

On a final note about experts models, they sometimes yield counter-intuitive insights. As argued above, "bad reputation" effects can arise — the reputational concern to conform to some expectation of buyer may lead a seller to discard useful information or even take costly actions that she knows to be wrong.⁶ Prat (2005) shows that transparency may be detrimental; he distinguishes between transparency of outcomes, which lead to faster learning about expertise and so are beneficial, and transparency of the seller's actions or decisions, which might boost a seller's incentives to conform to some expected behavior that may be inappropriate for the particular transaction.

⁶See also Brandenburger and Polak (1996) for a related idea.

10.2 Multiple Dimensions of Interest

Following the vast majority of the literature, we have focused on a single dimension for all objects of interest: a seller's type, his reputation, and the attributes of outcomes and how it is perceived by buyers. In many cases, however, reputation is multi-dimensional. There may be many channels a firm might try to use in order to convince its buyers about its "type," and its type might itself be a list of multi-dimensional attributes. For example, a management consulting firm may have a reputation for how broad or narrow is the scope of projects it handles, as well as a reputation for how good a job it does.

With regard to multiple dimensions for observations, it is easy to imagine a possibility of multiple equilibria by focusing on the possible dimensions over which reputation is relevant. If buyers expect a seller to devote its effort to a particular aspect of the job, this belief can be self-fulfilling (see Dewatripont et al. (1999a,b)). With regard to multiple dimensions of type, a firm may seek to establish a reputation in its industry as a tough competitor, however being tough in turn may depend on its cost structure and its patience (Bagwell, 2007b). A worker may be doing well because of her innate ability or her ambition (Közsegi and Li, 2008) — though these may have similar effects at the start of her career; potential employers may wish to distinguish between these as the employee nears retirement when ambition plays less of a role. In these examples, multi-dimensional abilities are summarized in unidimensional outcomes which might make the inference problem difficult. However, observations in different periods of a lifetime, for example, may allow observers to distinguish or make separate inferences on the constituent elements of type.

Of course, it is easy to imagine situations in which there are many dimensions both with respect to observations and type. In particular, multi-dimensional reputations can also have implications when there are distinct observations on different dimensions of type. For example, Bar-Isaac and Hörner (2008) suggest that since forcing an agent to specialize will lead the agent to undertake that task more often, it will enhance reputation incentives with respect to that specialized task.

The additional effort induced might override efficiency losses associated with undertaking an *ex-post* inefficient task.

Throughout this survey, we have typically considered only a single buyer or at least a single type of buyer; however, in many environments of interest, not only are there many buyers, but also buyers might vary in their tastes. Many of the intuitions developed above in the context of vertical or “quality” aspects apply to the case of horizontal or “match” concerns. The extent of uncertainty, rate of learning, and frequency of transaction can play an important role in determining firm strategies in this case as well. For example, consulting, law and other professional firms, differ not only in “quality” but also in their “styles” and the type of work they undertake and continually take actions to communicate and maintain reputations for particular style. Elite law firms in London, for example, vary from the more aggressive to the more traditional and consensual in their cultures. Their reputations and desire to sustain them affect both the kinds of clients and employees that they attract and retain. Similarly this “horizontal” view can help explain marketing and branding decisions. Umbrella branding can provide informational cues and lead to faster learning and improve efficiency (for example through risk reduction). Anand and Shachar (2004, 2005), Byzalov and Shachar (2004), and Sappington and Wernerfelt (1985) look at this role of risk reduction in umbrella branding and present empirical evidence in support of this role.

11

Empirical Studies

A great deal of discussion of “reputation” is somewhat imprecise and this has been reflected in casual empirical studies that attempt to value the reputation of a firm, or brand value by, for example, comparing stock market valuations of firms with the book value of physical assets. Increasingly though, the marketing literatures and others have attempted to value brands using structural models, and attempting to carefully control for extraneous effects and common factors.¹

Any attempt to empirically test the implications of more formal reputation models faces a number of challenges. The empirical challenge is to understand the strength with which “reputation” can overcome information problems, both with regard to imperfectly observed characteristics and imperfectly observed behavior, and the ability of reputation concerns to affect a seller’s behavior. Ideally, a researcher would know both the characteristics and behavior of the seller, and the beliefs and information available to buyers. In practice, this is all but impossible.

¹ See, for example, the discussion on empirical work on umbrella branding at the end of Section 9.2.

In particular, it is hard to find data where the researcher might know more than buyers do,² and, in many markets, buyers obtain information from numerous sources and social networks, and are heterogeneous in these sources in a way that may be difficult for a researcher to track. Finally, in many markets where we anticipate that reputation plays a key role, such as professional services, customers may have heterogeneous preferences with respect to what they may consider a “good” outcome and there may be many dimensions along which a transaction could be considered to have succeeded or failed. At a minimum, the researcher must observe *ex-post* outcomes that measure an important aspect of the transaction over which reputation will play a role.

There are settings in which some of these concerns are muted. In particular, in online trading environments for standardized goods, it seems reasonable to believe that the “quality” of the transaction merely reflects whether what was promised is actually delivered and that this happens in a timely manner (for example the 1/10 oz. 5 dollar gold coin of 2002 vintage). Furthermore, in a large online trading environment, it is reasonable to imagine that the only information that a buyer has about a seller is the information that appears online, so that the researcher can perfectly reconstruct the buyer’s information set. Finally, in such environments, it is even possible to perfectly observe characteristics of the seller through field experiments where the researcher can pose as a seller.

In the light of these observations, it is perhaps not surprising, that the last few years have witnessed a significant number of interesting papers that seek to explore reputation in the context of online trading environments. We do not offer to review this work but suggest the excellent summaries in Dellarocas (2006) and Bajari and Hortaçsu (2004). Broadly speaking, Dellarocas (2006) suggests that the principal conclusions of this literature is that feedback in such online environments affect prices and probability of sales, though the precise effects are subtle; the impact of feedback on prices and likelihood of sales is greater for relatively risky and expensive products; and negative feedback, especially recent negative feedback, are particularly influential.

²Though note that the researcher working with historical data might easily observe outcomes which are in the future from the perspective of buyers.

Using evidence from online markets, Cabral and Hortaçsu (2006) apply a “mixed model” where both imitation and separation incentives arise, in the spirit of Diamond’s (1989) model discussed in Section 7.1. Although, the evidence in support of a particular model of reputation in this environment is inconclusive, Cabral and Hortaçsu (2006) conclude in stating that: “Regardless of which theoretical model best explains the data, an important conclusion of our paper is that eBay’s reputation system gives way to noticeable strategic responses from both buyers and sellers. That is, the mechanism has ‘bite’” (p. 40). Further they highlight evidence that the behavior of sellers, as well as of buyers — the sole focus of much of the literature in both online and offline markets³ — is affected by reputation and feedback mechanisms.

Empirical work in the offline bricks-and-mortars world is somewhat more sparse. A few influential papers have built on the Holmström (1999) career concerns framework discussed in Section 5.1, and exploited variation in sellers’ incentives through their careers, as suggested in Section 7.1. Chevalier and Ellison (1999) consider the incentives of mutual fund managers: an industry where one would imagine that buyers have a fairly clear sense of a good or bad outcome *ex-post* (the fund’s rate of return) and thereby gain information on the manager. Performance, however, is somewhat more subtle inasmuch as managers are compared to their peers, reflecting that reputation here is for “expertise” and can lead to herding as suggested in Section 10.1. Incentives arise, primarily, from the prospect of termination and promotion, which imply reputational returns that are nonlinear in performance. Chevalier and Ellison show that “termination” is more performance-sensitive for younger managers, giving them stronger incentives to “herd.” Hong et al. (2000) show similar incentives and similar results in the security analysts industry. Hong and Kubik (2003) focus on the reputation incentives, rather than the analysts’ responses to these incentives. They highlight that not only might the strength of the career concerns incentives differ over a career, but that the nature of these incentives change, and that “reputation”

³For offline markets, see for example Borenstein and Zimmerman (1988) and Hubbard (2002).

might be multi-dimensional. In particular, incentives (and associated promotions) appear to be associated not only with the accuracy of forecasts (especially, early in a security analyst's career) but also for optimism later in one's career (especially, for analysts covering stocks underwritten by their brokerage houses).

Jin and Leslie (2003, 2008) exploit a natural experiment to show reputational effects and their influence on the behavior of both buyers and sellers. In particular, in December 1997 the Los Angeles County government passed an ordinance requiring restaurants to publicly display hygiene grade cards resulting from Department of Health Services hygiene inspections. Since such inspections had long been conducted, Jin and Leslie were in the unusual and enviable position of having data on quality that buyers would have liked to have had before the measure was introduced, as well as measures of outcomes before and after the introduction of posted hygiene cards. Using this unique opportunity, Jin and Leslie (2003) document that buyers responded to this information and that restaurants responded and improved hygiene. The effects were dramatic, and as a consequence of changes in both buyer and restaurant behavior, the hygiene grade cards caused a significant decrease in the number of people admitted to hospitals with food-borne illnesses. In their 2007 paper, Jin and Leslie highlight that, even before the introduction of grade cards, we should anticipate that reputation concerns should discipline and affect restaurants' behavior; however, these concerns vary by restaurant. In particular, incentives are indeed found to be stronger for chain-affiliation restaurants, or restaurants that rely relatively more on repeat business.

12

Future Directions

The greater focus on theory than empirical work in this review, to a large extent reflects the authors' knowledge and research interests, but also captures the state of progress in this field. While some broad intuitions are robust (in particular, something must be at risk in order to motivate a seller, the seller must care about the future, buyers must observe at least some kind of signals of past performance), the vast number of models and approaches adopted for different application suggest that there are subtleties. In trying to derive intuitions in application, the fine details can matter.

As we highlighted, in Section 5.2, different assumptions about the possible types of the seller can lead to different results and, in application, the range of possible types need not be obvious. For example, it is far from clear whether a lawyer (or for that matter an assistant professor in economics) is primarily motivated to prove herself to be brilliant rather than competent (as in the imitation approach of Section 5.2.1) or competent rather than inept (as in Section 5.2.2); however, as we have seen, one assumption can allow reputation concerns to last perpetually, and, in the other case, reputation concerns die out more quickly. Indeed, there seems to be more work to be done in pinning

down the determinants of reputational concerns and the actions that interested parties take to influence these. This brief discussion suggests that progress here would be made by focussing on particular applications.¹ More nuanced notions of reputation, for example, allowing for a variety of buyers with different interests or beliefs and perhaps a multi-dimensional view of the reputation of a seller could provide interesting avenues. In addition, and as discussed in Section 11, empirical work requires some measure of outcomes, separate from measures of the sellers reputation, which can be hard to find (though there have been a number of interesting studies that have been able to do so).

Looking forward, beyond progress on the applied and empirical direction, there are a number of exciting areas of current research which might have implications for work on reputation. First, and as we have briefly noted in Section 5.2.2, the implications of bounded rationality and, in particular, finite memory are being explored.² Somewhat more speculative and with less active research to date, it may be interesting to consider how the structure and design of the way that information spreads affects reputational incentives. For example, a wide sociology literature (here a classic reference is Wasserman and Faust (1995)) looks carefully at the structure of social interactions using the tools of network analysis. A growing literature in economics is addressing similar issues.³ Natural questions that arise, for example, are how an agent's position in society (or location in a network) affects her reputational concerns and behavior, as in Lippert and Spagnolo (2007), or how organizations and hierarchies might be designed in order that the communication of information leads agents to have stronger or more appropriate reputational concerns from the perspective of the designer.⁴

¹On a personal note, one of the authors of this piece has for a long time learned much (though not yet enough to fully satisfy him) by focussing attention on the role of reputation in the context of professional services. See Bar-Isaac (2004).

²In this context, in addition to the papers mentioned above. Phelan and Skrzypacz (2006) consider a repeated hidden action rather than a mixed model and restrict attention to strategies that can be represented by finite automata. One might consider the sophistication of an agent (either a consumer or seller) as represented by the number of states of an automaton.

³A good starting point for the interested reader is Jackson (2006).

⁴See, for example, Li (2007) for motivation and a first approach on this question.

Last but not least, as more transactions seem to be moving to the electronic marketplace, it seems natural that there will be more need to understand how reputations can play a role in supporting anonymous transactions across buyers and sellers. The past few years have offered just the beginning in answering his important question.

Acknowledgments

We thank Adam Brandenburger, Luis Cabral, Theodosios Dimopoulos, Joe Harrington, Martin Peitz, and Giancarlo Spagnolo for their comments and discussions. Victor Bennett provided excellent research assistance. Steven Tadelis is grateful to the National Science Foundation and to Yahoo! Research for financial support.

References

- Aaker, D. and K. Keller (1990), 'Buyer evaluations of brand extensions'. *Journal of Marketing* **54**, 27–41.
- Abreu, D. and R. Sethi (2003), 'Evolutionary stability in a reputational model of bargaining'. *Games and Economic Behavior* **44**, 195–216.
- Aghion, P., P. Bolton, C. Harris, and B. Jullien (1992), 'Optimal learning by experimentation'. *Review of Economic Studies* **58**, 621–654.
- Akerlof, G. A. (1970), 'The market for "lemons": Quality uncertainty and the market mechanism'. *Quarterly Journal of Economics* **84**, 488–500.
- Allen, F. (1984), 'Reputation and product quality'. *RAND Journal of Economics* **15**, 311–327.
- Almeida Costa, L. and L. Vasconcelos (2007), 'Share the fame or share the blame? The reputational implications of partnerships'. Mimeo, NOVA.
- Anand, B. and R. Shachar (2004), 'Brands as beacons: A new source of loyalty to multiproduct firms'. *Journal of Marketing Research* **41**, 135–150.
- Anand, B. and R. Shachar (2005), 'Advertising the matchmaker'. Mimeo, Harvard Business School.

- Anderson, A. and L. Smith (2006), 'Assortative matching and reputation'. Cowless Foundation Discussion Paper No. 1553.
- Andersson, F. (2002), 'Pooling reputations'. *International Journal of Industrial Organization* **20**, 715–730.
- Avery, C., P. Resnick, and R. Zeckhauser (1999), 'The market for evaluations'. *American Economic Review* **89**, 564–584.
- Bagwell, K. (2007a), 'The economic analysis of advertising'. In: M. Armstrong and R. Porter (eds.): *Handbook of Industrial Organization*, Vol. 3, Chapter 8. North Holland, pp. 1701–1844.
- Bagwell, K. (2007b), 'Signaling and entry deterrence: A multi-dimensional analysis'. *The RAND Journal of Economics* **38**(3), 670–697.
- Bajari, P. and A. Hortaçsu (2004), 'Economic insight from internet auctions'. *Journal of Economic Literature* **42**, 457–486.
- Balachander, S. and S. Ghose (2003), 'Reciprocal spillover effects: A strategic benefit of brand extensions'. *Journal of Marketing* **67**, 4–13.
- Banerjee, A. V. (1992), 'A simple model of herd behavior'. *Quarterly Journal of Economics* **107**, 797–817.
- Bar-Isaac, H. (2003), 'Reputation and survival: Learning in a dynamic signalling model'. *Review of Economic Studies* **70**, 231–251.
- Bar-Isaac, H. (2004), 'Reputation and professional services: Survival, teams and incentives'. Ph.D Thesis, University of London.
- Bar-Isaac, H. (2005), 'Imperfect competition and reputational commitment'. *Economics Letters* **89**, 167–173.
- Bar-Isaac, H. (2007), 'Something to prove: Choosing teamwork to create reputational incentives for individuals'. *RAND Journal of Economics* **38**, 495–511.
- Bar-Isaac, H. (2008), 'Transparency, career concerns, and incentives for acquiring expertise'. Working Paper, NYU.
- Bar-Isaac, H. and J. Ganuza (2008), 'Recruitment, training and career concerns'. *Journal of Economics and Management Strategy*, forthcoming.
- Bar-Isaac, H. and J. Hörner (2008), 'Reputation and specialization'. In progress, NYU.

- Benabou, R. and G. Laroque (1992), 'Using privileged information to manipulate markets: Insiders, gurus, and credibility'. *Quarterly Journal of Economics* **107**, 921–958.
- Bergemann, D. and J. Välimäki (1996), 'Learning and strategic pricing'. *Econometrica* **64**, 1125–1149.
- Bergemann, D. and J. Välimäki (2006), 'Dynamic price competition'. *Journal of Economic Theory* **126**, 232–263.
- Bernheim, B. D. and M. D. Whinston (1990), 'Multimarket contact and collusive behavior'. *Rand Journal of Economics* **21**, 1–26.
- Biglaiser, G. (1993), 'Middlemen as experts'. *Rand Journal of Economics* **24**, 212–223.
- Biglaiser, G. and J. W. Friedman (1994), 'Middlemen as guarantors of quality'. *International Journal of Industrial Organization* **12**, 509–533.
- Bikchandani, S., D. Hirschleifer, and I. Welch (1992), 'A theory of fads, fashion, custom, and cultural change as informational cascades'. *The Journal of Political Economy* **100**, 992–1026.
- Blanes-i-Vidal, J. (2007), 'Delegation of decision rights and the winner's curse'. *Economics Letters* **94**, 163–169.
- Bolton, P. and C. Harris (1999), 'Strategic Experimentation'. *Econometrica* **67**, 349–374.
- Borenstein, S. and M. B. Zimmerman (1988), 'Market incentives for safe commercial airline operation'. *American Economic Review* **78**, 913–935.
- Bose, S., G. Orosel, M. Ottaviani, and L. Vesterlund (2006), 'Dynamic monopoly pricing and herding'. *RAND Journal of Economics* **37**, 910–928.
- 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.
- Byzalov, D. and R. Shachar (2004), 'The risk reduction role of advertising'. *Quantitative Marketing and Economics* **2**, 283–320.
- Cabral, L. (2000), 'Stretching firm and brand reputation'. *Rand Journal of Economics* **31**, 658–673.
- Cabral, L. (2008), 'Umbrella branding with imperfect observability and moral hazard'. Mimeo, NYU.

- Cabral, L. and A. Hortaçsu (2006), 'The dynamics of seller reputation: Theory and evidence from eBay'. Mimeo, NYU.
- Cai, H. and I. Obara (2006), 'Firm reputation and horizontal integration'. Mimeo, UCLA.
- Casas-Arce, P. (2005), 'Career tournaments'. Working Paper, Universitat Pompeu Fabra.
- Chevalier, J. and G. Ellison (1999), 'Career concerns of mutual fund managers'. *Quarterly Journal of Economics* **114**, 389–432.
- Choi, J. P. (1998), 'Brand extension as informational leverage'. *Review of Economic Studies* **65**, 655–669.
- Choi, J. P. and D. S. Jeon (2007), 'A leverage theory of reputation building with co-branding: Complementarities in reputation building'. Working Paper, Michigan State University.
- Chu, W. and W. Chu (1994), 'Signaling quality by selling through a reputable retailer: An example of renting the reputation of another agent'. *Marketing Science* **13**, 177–189.
- Crawford, V. and J. Sobel (1982), 'Strategic information transmission'. *Econometrica* **50**, 1431–1451.
- Cremer, J. (1986), 'Cooperation in ongoing organizations'. *Quarterly Journal of Economics* **101**, 33–50.
- Cripps, M. W. (2006), 'Reputation'. *New Palgrave Dictionary of Economics*, second edition.
- Cripps, M. W., G. J. Mailath, and L. Samuelson (2004), 'Imperfect monitoring and impermanent reputation'. *Econometrica* **72**(2), 407–432.
- Dana, Jr., J. D. and Y.-F. Fong (2007), 'Product quality, reputation, and tacit collusion'. Mimeo, Kellogg School of Management.
- Dana, Jr., J. D. and K. Spier (2007), 'Bundling and product reputation'. Mimeo, Kellogg School of Management.
- Deb, J. (2007), 'Observability and sorting in a market for names'. Mimeo, Kellogg School of Management.
- Dellarocas, C. (2006), 'Reputation mechanisms'. In: T. Hendershott (ed.): *Handbook on Information Systems and Economics*. Elsevier Publishing, pp. 629–660.
- Dellarocas, C., F. Dini, and G. Spagnolo (2007), 'Designing reputation mechanisms'. In: N. Dimitri, G. Piga, and G. Spagnolo (eds.):

- Handbook of Procurement*, Chapter 18. Cambridge University Press, pp. 446–482.
- Dewatripont, M., I. Jewitt, and J. Tirole (1999a), ‘The economics of career concerns, Part I: Comparing information structures’. *Review of Economic Studies* **66**, 183–198.
- Dewatripont, M., I. Jewitt, and J. Tirole (1999b), ‘The economics of career concerns, Part II: Application to missions and accountability of government agencies’. *Review of Economic Studies* **66**, 99–217.
- Diamond, D. W. (1989), ‘Reputation acquisition in debt markets’. *Journal of Political Economy* **97**, 828–862.
- Dixit, A. (2004), *Lawlessness and Economics: Alternative Modes of Governance*. Princeton University Press.
- Ekmekci, M. and A. Wilson (2007), ‘A note on reputation effects with finite memory’. Mimeo, NYU.
- Ellison, G. (1994), ‘Cooperation in the prisoner’s dilemma with anonymous random matching’. *Review of Economic Studies* **61**, 567–588.
- Ely, J., D. Fudenberg, and D. K. Levine (2008), ‘When is reputation bad?’. *Games and Economic Behavior* **63**, 498–526.
- Ely, J. and J. Välimäki (2003), ‘Bad reputation’. *Quarterly Journal of Economics* **118**, 785–814.
- Erdem, T. (1998), ‘An empirical analysis of umbrella branding’. *Journal of Marketing Research* **65**, 655–669.
- Felli, L. and C. Harris (1996), ‘Job matching, learning and firm-specific human capital’. *Journal of Political Economy* **104**, 838–868.
- Fudenberg, D. and D. K. Levine (1989), ‘Reputation and equilibrium selection in games with a patient player’. *Econometrica* **57**, 759–778.
- Fudenberg, D. and D. K. Levine (1992), ‘Maintaining a reputation when strategies are imperfectly observed’. *Review of Economic Studies* **59**, 561–579.
- Fudenberg, D., D. K. Levine, and E. Maskin (1994), ‘The folk theorem in repeated games with imperfect public information’. *Econometrica* **62**, 997–1039.
- Fudenberg, D. and E. Maskin (1986), ‘Folk theorem for repeated games with discounting or with incomplete information’. *Econometrica* **54**, 533–554.

- Fudenberg, D. and J. Tirole (1986), 'A "signal-jamming" theory of predation'. *The RAND Journal of Economics* **17**, 366–376.
- Fudenberg, D. and J. Tirole (1991), *Game Theory*. The MIT Press.
- Gale, D. and R. Rosenthal (1994), 'Price and quality cycles for experience goods'. *RAND Journal of Economics* **25**, 590–607.
- Garella, P. and M. Peitz (2007), 'Alliances between competitors and consumer information'. *Journal of the European Economic Association* **5**, 823–845.
- Gibbons, R. and K. J. Murphy (1992), 'Optimal incentive contracts in the presence of career concerns: Theory and evidence'. *Journal of Political Economy* **100**, 468–505.
- Greif, A. (1989), 'Reputation and coalitions in medieval trade: Evidence on the maghribi traders'. *Journal of Economic History* **49**, 857–882.
- Greif, A. (2006), *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*. Cambridge University Press.
- Greif, A., P. Milgrom, and B. Weingast (1994), 'Commitment, coordination, and enforcement: The case of the merchant guilds'. *Journal of Political Economy* **102**, 745–776.
- Hakenes, H. and M. Peitz (2007), 'Observable reputation trading'. *International Economic Review* **48**, 693–730.
- Hakenes, H. and M. Peitz (2008a), 'Umbrella branding and external certification'. *European Economic Review*, forthcoming.
- Hakenes, H. and M. Peitz (2008b), 'Umbrella branding and the provision of quality'. *International Journal of Industrial Organization* **26**, 546–556.
- Harrington, J. E. (1993), 'Economic policy, economic performance, and elections'. *American Economic Review* **83**, 27–42.
- Harrington, J. E. (1995), 'Cooperation in a one-shot prisoners' dilemma'. *Games and Economic Behavior* **8**, 364–377.
- Harstad, B. (2007), 'Organizational form and the market for talent'. *Journal of Labor Economics* **25**, 581–611.
- Holmström, B. (1999), 'Managerial incentive problems — A dynamic perspective'. *Review of Economic Studies* **66**, 169–182. (Originally published in 1982 in *Essays in Honor of Professor Lars Wahlback*).

- Holmström, B. and Ricart-i-Costa (1986), 'Managerial incentives and capital management'. *Quarterly Journal of Economics* **101**, 835–860.
- Hong, H. and J. D. Kubik (2003), 'Analyzing the analysts: Career concerns and biased earnings forecasts'. *Journal of Finance* **58**, 313–351.
- Hong, H., J. D. Kubik, and A. Solomon (2000), 'Security analysts' career concerns and herding of earnings forecasts'. *RAND Journal of Economics* **31**, 121–144.
- Hörner, J. (2002), 'Reputation and competition'. *American Economic Review* **92**, 644–663.
- Hubbard, T. N. (2002), 'How do buyers motivate experts? Reputational incentives in an auto repair market'. *Journal of Law and Economics* **45**, 437–468.
- Jackson, M. O. (2006), 'The economics of social networks'. In: R. Blundell, W. Newey, and T. Persson (eds.): *Advances in Economics and Econometrics, Theory and Applications: Ninth World Congress of the Econometric Society*, Vol. 1, Chapter 1. Cambridge University Press.
- Jeon, S. (1996), 'Moral hazard and reputational concerns in teams: Implications for organizational choice'. *International Journal of Industrial Organization* **14**, 297–315.
- Jin, G. Z. and P. Leslie (2003), 'The effect of information on product quality: Evidence from restaurants hygiene grade cards'. *Quarterly Journal of Economics* **118**, 409–451.
- Jin, G. Z. and P. Leslie (2008), 'Reputational incentives for restaurant hygiene'. *American Economic Journal (Microeconomics)*, forthcoming.
- Kandori, M. (1992), 'Social norms and community enforcement'. *Review of Economic Studies* **59**, 63–80.
- Kandori, M. (2006), 'Repeated games'. *The New Palgrave Dictionary of Economics*, second edition.
- Klein, B. and K. B. Leffler (1981), 'The role of market forces in assuring contractual performance'. *Journal of Political Economy* **89**, 615–641.

- Klein, T., C. Lambertz, G. Spagnolo, and K. Stahl (2006), ‘Last-minute feedback’. CEPR Discussion Paper No. 5693.
- Koch, A. and E. Peyrache (2005), ‘Tournaments, individualized contracts and career concerns’. IZA DP number 1841.
- Kovrijnykh, R. (2007), ‘Career uncertainty and dynamic incentives’. Working Paper, University of Chicago.
- Közsegi, B. and W. Li (2008), ‘Drive and talent’. *Journal of the European Economic Association* **6**(1), 210–236.
- Kranton, R. E. (2003), ‘Competition and the incentive to produce high quality’. *Economica* **70**, 385–404.
- Kreps, D. (1990), ‘Corporate culture and economic theory’. In: J. Alt and K. Shepsle (eds.): *Perspectives on Positive Political Economy*. Cambridge: Cambridge University Press, pp. 90–143.
- 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.
- Kreps, D. and R. Wilson (1982), ‘Reputation and imperfect information’. *Journal of Economic Theory* **27**, 253–279.
- Kreps, D. M. and J. Sobel (1994), ‘Signalling’. In: R. Aumann and S. Hart (eds.): *Handbook of Game Theory, Volume II*. Elsevier, pp. 849–868.
- Levin, J. and S. Tadelis (2005), ‘Profit sharing and the role of professional partnerships’. *Quarterly Journal of Economics* **120**, 131–171.
- Levy, G. (2004), ‘Anti-herding and strategic consultation’. *European Economic Review* **48**, 503–525.
- Levy, G. (2007), ‘Decision making in committees: Transparency, reputation and voting rules’. *American Economic Review* **97**, 150–168.
- Li, W. (2007), ‘Peddling influence through intermediaries: Propaganda’. Riverside working paper, University of California.
- Lippert, S. and G. Spagnolo (2007), ‘Networks of relations, word-of-mouth communication, and social capital’. SSE/EFI Working Paper in Economics and Finance No. 570.
- Liu, Q. (2007), ‘Information acquisition and reputation dynamics’. Mimeo, University of Pennsylvania.
- Lizzeri, A. (1999), ‘Information revelation and certification intermediaries’. *RAND Journal of Economics* **30**, 214–231.

- Macleod, B. (2007), 'Reputations, relationships and the enforcement of incomplete contracts'. *Journal of Economic Literature* **45**, 595–628.
- Mailath, G. J. and L. Samuelson (2001), 'Who wants a good reputation?'. *Review of Economic Studies* **68**, 415–441.
- Mailath, G. J. and L. Samuelson (2006), *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press.
- Martinez, L. (2008), 'Reputation, career concerns, and job assignments'. *B.E. Journal of Theoretical Economics*, forthcoming.
- Marvel, H. and L. Ye (2008), 'Trademark sales, entry, and the value of reputation'. *International Economic Review* **49**, 547–576.
- Maskin, E. and J. Tirole (2001), 'The politician and the judge: Accountability in government'. Working Paper, Toulouse University.
- McKenna, C. (2006), *The World's Newest Profession: Management Consulting in the Twentieth Century*. Cambridge University Press.
- Meyer, M. A. (1994), 'The dynamics of learning with team production: Implications for task assignment'. *Quarterly Journal of Economics* **109**, 1157–1184.
- Miklós-Thal, J. (2006), 'On the signaling and feedback effects of umbrella branding'. Mimeo, University of Toulouse.
- Milgrom, P., D. North, and B. Weingast (1990), 'The role of institutions in the revival of trade: The medieval law merchant, private judges, and the champagne fairs'. *Economics and Politics* **2**, 1–23.
- Milgrom, P. R. and D. J. Roberts (1982), 'Predation, reputation and entry deterrence'. *Journal of Economic Theory* **27**, 280–312.
- Milgrom, P. R. and J. Roberts (1986), 'Price and advertising signals of product quality'. *Journal of Political Economy* **94**, 796–821.
- Miller, N., P. Resnick, and R. Zeckhauser (2005), 'Eliciting informative feedback: The peer prediction method'. *Management Science* **51**, 1359–1373.
- Monte, D. (2007), 'Reputation and bounded memory'. Mimeo, Yale University.
- Montgomery, C. and B. Wernerfelt (1992), 'Risk reduction and umbrella branding'. *Journal of Business* **65**, 31–50.
- Morris, S. (2001), 'Political correctness'. *Journal of Political Economy* **109**, 231–265.

- Morrison, A. D. and W. J. Wilhelm, Jr. (2004), 'Partnership firms, reputation and human capital'. *American Economic Review* **94**, 1682–1692.
- North, D. C. (1990), *Institutions, Institutional Change and Economic Performance*. Cambridge University Press.
- Okuno-Fujiwara, M. and A. Postlewaite (1995), 'Social norms and random matching games'. *Games and Economic Behavior* **9**, 79–109.
- Ortega, J. (2003), 'Power in the firm and managerial career concerns'. *Journal of Economics and Management Strategy* **12**, 1–29.
- Ottaviani, M. and P. Sørensen (2006a), 'Professional advice'. *Journal of Economic Theory* **126**, 120–142.
- Ottaviani, M. and P. Sørensen (2006b), 'Reputational cheap talk'. *RAND Journal of Economics* **37**, 155–175.
- Ottaviani, M. and P. Sørensen (2006c), 'The strategy of professional forecasting'. *Journal of Financial Economics* **81**, 441–466.
- Peppall, L. and D. J. Richards (2002), 'The simple economics of brand-stretching'. *Journal of Business* **75**, 535–552.
- Perri, T. J. and E. B. Rasmusen (2001), 'Can high prices ensure product quality when buyers do not know the seller's cost?'. *Economic Inquiry* **39**, 561–567.
- Phelan, C. (2006), 'Public trust and government betrayal'. *Journal of Economic Theory* **130**, 27–43.
- Phelan, C. and A. Skrzypacz (2006), 'Private monitoring with infinite histories'. Federal Reserve Bank of Minneapolis, Research Department Staff Report 383.
- Prat, A. (2005), 'The wrong kind of transparency'. *American Economic Review* **95**, 862–877.
- Prendergast, C. and L. Stole (1996), 'Impetuous youngsters and jaded oldtimers'. *Journal of Political Economy* **104**, 1105–1134.
- Richman, B. (2006), 'How communities create economic advantage: Jewish diamond merchants in New York'. *Law and Social Inquiry* **31**, 383–420.
- Rob, R. and A. Fishman (2005), 'Is bigger better? Customer base expansion through word-of-mouth reputation'. *Journal of Political Economy* **113**, 1146–1162.

- Rosen, S. (1981), 'The economics of superstars'. *American Economic Review* **71**, 845–858.
- Rothschild, M. (1974), 'A two-armed bandit theory of market pricing'. *Journal of Economic Theory* **9**, 185–202.
- Sappington, D. E. M. and B. Wernerfelt (1985), 'To brand or not to brand? A theoretical and empirical question'. *Journal of Business* **58**, 279–293.
- Scharfstein, D. and J. Stein (1990), 'Herd behavior and investment'. *American Economic Review* **80**, 465–479.
- Shapiro, C. (1983), 'Premiums for high quality products as returns to reputations'. *Quarterly Journal of Economics* **98**, 659–680.
- Sobel, J. (1985), 'A theory of credibility'. *Review of Economic Studies* **52**, 557–573.
- Spence, M. (1973), 'Job market signaling'. *Quarterly Journal of Economics* **87**, 355–374.
- Spulber, D. (1999), *Market Microstructure: Intermediaries and the Theory of the Firm*. Cambridge University Press.
- Stein, J. (1989), 'Efficient capital markets, inefficient firms: A model of myopic corporate behavior'. *Quarterly Journal of Economics* **104**, 655–669.
- Sullivan, M. (1990), 'Measuring image spillovers in umbrella-branded products'. *Journal of Business* **63**, 309–329.
- Tadelis, S. (1999), 'What's in a name? Reputation as a tradeable asset'. *American Economic Review* **89**, 548–563.
- Tadelis, S. (2002), 'The market for reputations as an incentive mechanism'. *Journal of Political Economy* **92**, 854–882.
- Tadelis, S. (2003), 'Firm reputations with hidden information'. *Economic Theory* **110**, 854–882.
- 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.
- van Damme, E. (1992), 'Refinements of Nash equilibrium'. In: J.-J. Laffont (ed.): *Advances in Economic Theory: Sixth World Congress Vol I*. 28 June 2007: Cambridge University Press, Cambridge Collections Online, Cambridge University Press. http://cco.cambridge.org/extract?id=ccol0521416663_CCOL0521416663A003.

- Wang, T. (2007), 'The reputation of an organization and its dynamics'. Mimeo, London School of Economics.
- Wasserman, S. and K. Faust (1995), *Social Network Analysis: Methods and Applications*. Cambridge University Press.
- Wernerfelt, B. (1988), 'Umbrella branding as a signal of product quality: An example of signaling by posting a bond'. *RAND Journal of Economics* **19**, 458–466.
- Zwiebel, J. (1995), 'Corporate conservatism and relative compensation'. *Journal of Political Economy* **103**, 1–25.