

Reputation: The New Palstave Dictionary  
Martin Cripps, March 2006 Revised July 2006

Abstract: We explain what reputation effects are, how they arise and the factors that limit or strengthen them.

JEL classification: C72,C73,D43,L13.

In a dynamic setting signals sent now may affect the current and future behavior of other players, thus signals can have effects unrelated to their current costs and benefits. It is the interplay between signals and their long run consequences that is studied in the literature on reputation.

The literature on reputation has two main themes. The first is that introducing a small amount of incomplete information in a dynamic game can dramatically change the set of equilibrium payoffs: introducing something to signal can have big implications in a dynamic model. These kind of results can also be interpreted as providing a robustness check. Dynamic and repeated games typically have many equilibria and reputation results allow us to determine which equilibria continue to be played when a game is “close” to complete information. The second theme of the literature on reputations is that introducing incomplete information in a dynamic game may introduce new and important signalling dynamics in the players’ strategies. Thus reputation effects tells us something about behavior. This theme is particularly important in applications to macroeconomics and to industrial organization, for example. For either of these themes to be relevant it is necessary to have a dynamic game with incomplete information, so work on reputation has been influenced by, and influences, the larger literature on repeated and dynamic games of incomplete information. An excellent detailed treatment of reputation can be found in Mailath and Samuelson (2006).

## 1. An example

Most of the results below will be described in the context of a simple infinitely repeated trading game. The row player is a seller who can produce high or low quality. The column player is a buyer. Producing high quality is always expensive for the seller, so she would rather produce low quality, the buyer, however, only wants to buy a high quality product. The only non-standard element is that the buyer regrets not buying a high quality product. The trading game (Figure 1) has a unique equilibrium  $(L, N)$ .

		Buy(B)	Not Buy(N)
High Quality	(H)	(1, 1)	(-1, -1)
Low Quality	(L)	(2, -1)	(0, 0)

Figure 1: A Trading Game

Let us record some facts about this game. The set

$$V \equiv \{ (x, y) : x > 0, y > -1/3, y \leq x, y \leq 3 - 2x \} \subset \mathbb{R}^2,$$

illustrated in Figure 2 below, is the set of feasible and strictly individually rational payoffs for the trading game.




Figure 2: Sets of Equilibrium Payoffs and Reputation Bounds

(In this figure the axes are drawn through the minmax payoffs to make  $V$  clear.) If the seller could commit to a pure strategy, she would prefer to choose  $H$  as the buyer's best response to this is  $B$ . However, she could do even better by committing to a mixed strategy; playing  $(3/4, 1/4)$  for example would also ensure the buyer played  $B$  and give the seller a bigger payoff. Reputation arguments can provide ways for these commitment payoffs to be achieved by sellers who are not actually committed to anything.

The trading game is played in each of the periods  $t = 1, 2, \dots$  with perfect monitoring; at the end of the period the players get to observe all payoffs and the pure action taken by their opponent. If both players' discount factors,  $\delta < 1$ , were sufficiently large, any point in  $V$  could be sustained as an equilibrium payoff — see entry on repeated games. If the seller is long

lived but faces an infinite sequence of buyers who each live one period, then, any point on the line segment joining  $(0, 0)$  to  $(1, 1)$  is an equilibrium payoff. (No seller payoff above 1 is achievable if mixed actions are not observable, see Fudenberg, Kreps and Maskin (1990).)

The stage is now set. To understand how reputation works we will need to introduce something for the seller to signal: Its commitment to high quality? Its low cost of high quality? Its commitment to always ripping off customers...? At this stage it is unnecessary to be specific, and we will concentrate on the general issues of learning. There are two types of sellers, “strong” and “normal”, that the buyer may face in a game. The seller is told their type by nature at time  $t = 0$ . The buyer, however, is unaware of nature’s selection and spends the rest of the game looking at the seller’s behavior and trying to figure out what type she is. The normal seller plays action  $a \in \{H, L\}$  with probability  $\hat{\sigma}^t(a)$  at time  $t$  and the strong seller plays  $a$  with probability  $\tilde{\sigma}^t(a)$  at time  $t$ . Everything we say in the section below applies to the case where normal and strong sellers follow history-dependent strategies. (These behavior strategies do depend on the (public) history of play before time  $t$ , but let us keep this out of our notation.) An initially uninformed buyer attaches probability  $p^t$  to the strong type and  $1 - p^t$  to the normal type at time  $t$ , again this depends on the observed history. Our buyer expects the seller to play  $a$  with probability  $\bar{\sigma}^t(a) = p^t \tilde{\sigma}^t(a) + (1 - p^t) \hat{\sigma}^t(a)$  and as time passes the buyers observe the outcomes of this strategy and revise their prior accordingly.

## 2. Tricks with Bayes’ rule and martingales

Now we generate three properties of learning that are extensively used in the reputations literature. We will call them the “merging” property, the “right ballpark” property and the “finite surprises” property. These properties are based on some simple facts about how Bayesian agents revise their beliefs. That is, how uncertainty about the seller’s type is processed by the buyers or any other observer of its behavior. A more advanced treatment of these results can be found in Sorin (1999). We will defer any derivation of reputation results to the next section so a reader could skip this section.

How does the buyer revise their beliefs in the light of an observed action  $a^t$ ? A plain application of Bayes’ rule tells us

$$p^{t+1} = \frac{\Pr(a^t \cap \text{Strong})}{\Pr(a^t)} = \frac{p^t \tilde{\sigma}^t(a^t)}{\bar{\sigma}^t(a^t)}.$$

Or, in terms of the change in the beliefs

$$p^{t+1} - p^t = \frac{p^t[\hat{\sigma}^t(a^t) - \bar{\sigma}^t(a^t)]}{\bar{\sigma}^t(a^t)} = \frac{p^t(1 - p^t)[\hat{\sigma}^t(a^t) - \bar{\sigma}^t(a^t)]}{\bar{\sigma}^t(a^t)}.$$

These equalities are powerful tools when combined with the properties of the priors.

**Merging property** This tells us exactly how the long-run behavior of the sellers is related to the buyer's long run beliefs. Either,  $p^t(1 - p^t) \rightarrow 0$  and the buyer eventually learns the type of the seller and can perfectly predict their actions. Or, all types of the seller end up behaving in the same way  $\hat{\sigma}^t(a^t) - \bar{\sigma}^t(a^t) \rightarrow 0$  and again the buyer can perfectly predict their actions. Nothing else can happen!

The stochastic process  $\{p^t\}$  is a martingale on  $[0, 1]$  with respect to public histories. To see this there is a simple calculation we can do.

$$E(p^{t+1}|h_t) = \sum_{a^t} \Pr(a^t)p^{t+1} = \sum_{a^t} \bar{\sigma}^t(a^t) \frac{p^t \hat{\sigma}^t(a^t)}{\bar{\sigma}^t(a^t)} = p^t$$

(The expectation  $E(\cdot)$  is taken with respect to the buyer's beliefs about future play.) Bounded martingales converge almost surely (see Williams (1991) for example), which implies  $|p^{t+1} - p^t| \rightarrow 0$  almost surely. Applying this to the second equality above (noting that  $|\bar{\sigma}^t(a^t)| \leq 1$ ) we get

$$p^t(1 - p^t)|\hat{\sigma}^t(a^t) - \bar{\sigma}^t(a^t)| \rightarrow 0, \quad (\text{Merging})$$

almost surely. This kind of result is extensively used in Hart (1985) and the literature that stems from his work.

**Right ballpark property** The strong seller knows that the future will evolve according to the strategy  $\hat{\sigma}$  (we use  $\hat{\Pr}(\cdot)$  and  $\hat{E}(\cdot)$  to denote her probability measure and its expectation). This seller might ask, as she plays out an equilibrium, how little probability can the buyers attach to the strong seller? Or, how low could  $p^t$  get when she plays  $\hat{\sigma}$ ? Of course, when the seller is in fact the strong type it is very unlikely that  $p^t$  becomes low — beliefs must stay in the right ballpark. (For example, if  $\hat{\sigma}$  was actually a pure strategy the strong seller cannot ever believe  $p^t$  will decrease. As she plays  $\hat{\sigma}$  there will be periods in which the normal type of seller could have done something different so observing the actions of  $\hat{\sigma}$  will cause buyers to revise  $p^t$  upwards.)

From the perspective on the strong seller, the likelihood ratio is a martingale:

$$\hat{E} \left( \frac{1 - p^{t+1}}{p^{t+1}} \mid h_t \right) = \frac{1 - p^t}{p^t}.$$

(The calculation is just like the earlier one for  $p^t$ , where we use  $\hat{\text{Pr}}(a^t) = \hat{\sigma}(a^t)$ .) Let  $\tau$  be the first time,  $s$  say, that  $p_s \leq \nu$  and let  $C^t$  be the event that  $\tau \leq t$ . That is, sometime in the first  $t$  periods  $p_s \leq \nu$ . Then the martingale property combined with the Optional Stopping Theorem (for example, Williams 1991) implies

$$\frac{1 - p^0}{p^0} = \hat{E} \left( \frac{1 - p^{t+1}}{p^{t+1}} \right) \geq \hat{\text{Pr}}(C^t) E \left( \frac{1 - p^\tau}{p^\tau} \mid C^t \right) \geq \hat{\text{Pr}}(C^t) \frac{1 - \nu}{\nu}.$$

The above gives an upper bound on  $\hat{\text{Pr}}(C^t)$  that is independent of  $t$ . Thus it also bounds the probability that  $p_t$  is ever below  $\nu$ :

$$\hat{\text{Pr}}(\exists t \text{ s.t. } p_t < \nu) \leq \frac{\nu}{p^0}. \quad (\text{Right Ballpark})$$

Hence, the strong seller knows that it is very unlikely that the buyer's posterior will ever be close to certain she is actually the normal seller.

**Finite surprises property** The strong seller might also ask how many times (as she plays  $\hat{\sigma}$ ) will the uninformed buyers make a big mistake in predicting her strategy? That is, how many periods is  $\|\hat{\sigma}^t - \bar{\sigma}^t\| > \nu$  when the seller actually plays  $\hat{\sigma}$ ? Here we are helped by the fact that our seller has only two actions, so the variation distance between the mixed actions is just twice the difference in probability of the realized action  $\|\hat{\sigma}^t - \bar{\sigma}^t\| = 2|\hat{\sigma}^t(a^t) - \bar{\sigma}^t(a^t)|$ . Let  $M_N$  be the event that there are more than  $N$  mistakes,  $\|\hat{\sigma}^t - \bar{\sigma}^t\| > \nu$ , before time  $T$ . The finite surprises property is that independently of the equilibrium  $\hat{\text{Pr}}(M_N) \rightarrow 0$  as  $T, N \rightarrow \infty$ . Thus it is very unlikely that there are many periods where the buyers do not think the seller will play as the strong type if the seller is indeed this type.

Jensen's inequality applied to the likelihood ratio above implies that the prior is a submartingale, that is,  $\hat{E}(p^{t+1}|h_t) \geq p^t$ . There is a second property of martingales we can now use — they cannot move around very much:  $\sum_{t=1}^T \hat{E}((p^{t+1} - p^t)^2) \leq 1$ . (A proof of this fact follows from  $\hat{E}(p^{t+1} - p^t)^2 \leq \hat{E}((p^{t+1})^2 - (p^t)^2)$ .) A substitution from the first Bayes' rule equality above then tells us

$$1 \geq \sum_{t=1}^T \hat{E} \left( (p^t [\hat{\sigma}^t(a^t) - \bar{\sigma}^t(a^t)])^2 \right).$$

It is immediate that only a few of the (non-negative) terms in the sum above can be much above zero, otherwise the upper bound will be violated. The right ballpark property tells us it is very unlikely that  $p^t < \nu$ . On the event  $\{p^t \geq \nu \forall t\} \cap M_N$  the  $p^t$  in the above expectation is greater than  $\nu$  and there are at least  $N$  differences that are bigger than  $\nu/2$ , so the sum is at least  $N\nu(\nu/2)^2$ , hence

$$1 \geq \sum_{t=1}^T \hat{E}(p^t[\hat{\sigma}^t(a^t) - \bar{\sigma}^t(a^t)])^2 \geq \hat{\Pr}(\{p^t \geq \nu \forall t\} \cap M_N) \frac{N\nu^3}{4}.$$

Using the fact that  $\Pr(A \cap B) \geq \Pr(B) - \Pr(A^c)$  we now have an upper bound on  $\hat{\Pr}(M_N)$ .

$$\frac{4}{N\nu^3} + \hat{\Pr}(\exists t \text{ s.t. } p_t < \nu) \geq \hat{\Pr}(M_N)$$

The right ballpark property gives us

$$\hat{\Pr}(M_N) \leq \frac{\nu}{p^0} + \frac{4}{N\nu^3}. \quad (\text{Finite Surprises})$$

As the size of the surprises becomes small  $\nu \rightarrow 0$  and the number of surprises becomes large  $N\nu^3 \rightarrow \infty$  the strong seller must attach smaller and smaller probability to  $M_N$ . Fudenberg and Levine (1989; 1992), for example, invoke this property.

### 3. Basic reputation results: behavior

The three tools above are sufficient to establish most well known reputation results. The arguments below are entirely general, and are widely applied, but we will only use them in the trading game. To make things simple suppose that for some reason the strong seller is committed to playing  $(b, 1 - b)$ , that is, in every period  $t$  the strong buyer provides high quality with probability  $b$ . We will reserve the discussion of more complicated types of reputations for a later section.

From the perspective of the buyer any equilibrium will consist of two phases: An initial phase where there is learning and signalling about the seller's type (this is sometimes called reputation building although often reputation destruction is what occurs). And a terminal phase where the learning has virtually settled down. It is the merging property that tells us there must be this latter phase. The play in the game moves into this second phase either because the buyer is almost sure he knows the type of

the seller (reputation considerations have vanished) or because the sellers are playing in the same way. Thus the equilibria of dynamic signalling games are inherently non-stationary, which is in contrast to much of the work on repeated games. Of course, Markovian equilibria can be calculated, but these too will exhibit the two phases of play. The initial learning, when reputation builds or is destroyed, depends on the particular equilibrium and the game being studied. This phase may last only one period (if a once and for all revealing action is taken by a seller) but frequently it is long and has a random duration (if both types of seller randomize, for example).

Let us first examine reputation destruction in the case where  $b \approx 1$ , so the strong seller is committed to high quality and only very occasionally slips up. There is an equilibrium of this game where the normal type of seller will offer low quality more often than the strong type and thereby gradually reveal her type (destroy her reputation for being good). Nevertheless, as this occurs she will enjoy heightened payoffs. The trade-offs our normal seller experiences in this game are what drive the reputation destruction. A seller offering low quality today enjoys the benefit of a higher payoff now, but the observation of low quality typically leads the buyers to revise downwards their probability of the strong seller and buy less in the future. Whereas, a seller offering high quality will lead the buyer's posterior on the strong seller to be revised upwards and an increased likelihood of buying in the future. Exactly how the normal seller chooses to trade off long run benefits and short run costs is unclear. It is possible that pooling dominates and that future buying is so strong that the normal seller prefers to offer high quality today even if it costs something in the short run. However, in this equilibrium the normal seller perceives the long run benefits to be relatively small and prefers to offer low quality today. The normal seller can be thought of exploiting, or cashing-in, the value of her accumulated reputation. We also know, from the finite surprises property, that there will be finite opportunities for the normal seller to do this. Relatively soon there will be a time where the buyers know the seller is normal and purchase accordingly.

The nature of the trade-offs above depend on the form of the incomplete information and do not arise in games where there is complete information. If there is no strong seller present, there is no particular reason for a seller to expect that good quality today would necessarily be rewarded by more buying tomorrow. The uncertainty about seller types and the way buyers revise their beliefs, however, makes this occur. Thus the presence of uncertainty has eliminated those equilibria of the complete information game where good quality today is followed by a reduced probability of good quality in the future. This is an important restriction on behavior in equilibrium.

Reputation building (as opposed to destruction) is more likely in a world where there is the possibility that you are thought to be bad: for example if the strong type is committed to ripping customers off and only occasionally produces a good product ( $b \approx 0$ ). In such a world the normal seller wants to tell buyers she is not this type, because by playing as the strong type she is doomed to never trade. She is building a reputation for *not* being the strong type. To do this the normal type will have to incur the cost of repeatedly offering high quality, even if the buyer is not buying. This is expensive and will drag the normal seller's equilibrium payoff down. But, as above, it will increase the likelihood of future buying by decreasing the likelihood of a strong seller. In contrast to the reputation destruction case there are short run costs borne by the normal type to achieve long run gains. Again, the nature of these costs and benefits rely on the uncertainty the buyers have about the seller's type.

#### 4. Basic reputation results: payoffs

Reputation issues can have an extreme effect on payoffs and this is what first came to the attention of economists. The general question, how does the presence of something to signal in the repeated game affect the equilibrium payoffs, could be answered in a number of ways. One way would be to calculate equilibria explicitly. This is usually difficult and would not establish results that hold for *all* equilibria.

Instead a different approach is taken that is described in the following recipe:

1. If the seller is strong, then in finite time the buyers will believe they face a seller who plays arbitrarily close to  $(b, 1 - b)$  forever.
2. Figure out what will the buyers do when the seller is strong.
3. Use Step Two to evaluate the normal seller's payoff if she pretends to be strong forever.
4. At a Nash equilibrium the answer to Step Three is a lower bound on the normal seller's equilibrium payoff.

Step 1 is independent of the model and is a result of our earlier calculations. The right ballpark property tells us that  $p^t$  does not tend to zero when the seller is strong. The merging property then implies either  $p^t \rightarrow 1$ , or eventually all remaining normal types of buyer are also playing arbitrarily



close to  $(b, 1 - b)$ . In either of these cases, at a large but finite time the buyers believe that they face a seller who will always play  $(b, 1 - b)$ .

Before proceeding to apply this recipe we will illustrate its power with the remarkable results we expect to get. Let us first consider a world where buyers are short run. We will show that introducing arbitrarily small probability that there is a strong seller places a lower bound on the normal seller's equilibrium payoffs of  $2 - b$  (when  $b > 1/3$ ). Thus for  $b$  close to  $1/3$  the equilibrium payoffs in the complete information game (the segment joining  $(0, 0)$  and  $(1, 1)$ ) and the incomplete information game are disjoint! Moreover, the normal seller can get almost his maximum feasible payoff at every equilibrium. In the second case, where buyers are also long run we will get less strong conclusions, nevertheless, we will show that the normal type of seller must get at least  $2/3 - b$  when  $b > 1/3$ . These payoffs are illustrated in Figure 2.

The really difficult part of our recipe is Step Two, because we have to understand how the buyers will behave in equilibrium. We, therefore, need to consider as separate cases what happens if buyers are short run or long run. Also, the amount of discounting that the sellers do affects the answer to Step Three, so we need to consider different arguments for different amounts of discounting. The following catalog moves from simple to more elaborate arguments and from stronger to weaker reputation effects.

**Reputation without discounting: short term buyers** When a buyer lives only one period he plays a best response to the seller's current action. By Step One in the very long run this will be  $B$  if  $b > 1/3$  and  $N$  if  $b < 1/3$ . Step Three is simple, by playing  $\hat{\sigma}$  forever the normal seller knows that in a large but finite time she can ensure the buyer will behave as above and so she will receive a stage game payoff approximately  $R^*(b)$ , where

$$R^*(b) := \begin{cases} 2 - b & b > 1/3, \\ -b & b < 1/3. \end{cases}$$

If there is no discounting, and limits are correctly taken,  $R^*(b)$  will equal the normal type's payoff from playing  $\hat{\sigma}$  forever. Thus we have, Step Four, at any Nash equilibrium the normal type must get at least  $R^*(b)$ .

In a general game  $R^*$  is equal to the seller's payoff from playing the strong type's stage game strategy when the buyer plays their unique best response. (If the best response is not unique this is not correct.)

**Reputation with discounting: short term buyers** Step Two is as above — we still have short term buyers. When the normal seller discounts payoffs, however, playing  $\hat{\sigma}$  and eventually getting  $R^*(b)$  every period does not tell us what her payoff discounted to time zero will be. There is an order of limits issue; as the discounting of the seller becomes weaker ( $\delta \rightarrow 1$ ) it could be that the equilibria change and there are more and more periods where the seller is not getting  $R^*(b)$ . It is now that the finite surprises property plays an important role. First notice that when  $\nu$  is chosen appropriately and  $\|\bar{\sigma} - \hat{\sigma}\| < \nu$ , then playing a best response to  $\bar{\sigma}$  is the same as playing a best response to  $\hat{\sigma}$ . Hence, it is only when a surprise occurs that the normal seller is not getting  $R^*(b)$  from playing  $\hat{\sigma}$ . But the probability of more than  $N$  surprises can be made very small *independently of the discounting*. So, as the discounting becomes weak and  $N$  periods have a small effect on total discounted payoff, there is a small probability of the normal seller of getting anything less than  $R^*(b)$  when she plays  $\hat{\sigma}$ . Any Nash equilibrium, therefore, gives the normal seller at least  $R^*(b)$ . This is the kind of argument first made in specific cases by Kreps and Wilson (1982); Milgrom and Roberts (1982) and generalized in Fudenberg and Levine (1989; 1992).

**Reputation without discounting: one long run buyer** If the buyer lives for many periods, he will not necessarily play a short run best response to  $(b, 1 - b)$  even if he expects it to be played forever. We can, however, use some weaker information. At an equilibrium the buyer must on average get at least  $-1/3$  (his minmax payoff) against  $(b, 1 - b)$ . This implies that the buyer has to buy with at least probability  $1/3$  when  $b > 1/3$  and buy with at most probability  $1/3$  when  $b < 1/3$ . There are, consequently, some bounds on the normal seller's payoff when she has played  $\hat{\sigma}$  for a sufficiently long time. While playing  $(b, 1 - b)$  she gets  $2 - b$  when the buyer buys and  $-b$  if not, thus if the buyer buys with probability greater than  $1/3$  she expects to receive a payoff of at least  $2/3 - b$ . If the buyer buys with at most probability  $1/3$  she expects to get at least  $-b$ . The seller is not discounting, so what she gets in the long run from playing  $\hat{\sigma}$  is also what she expects to get at time zero. Our answer to Step Three, therefore, is

$$R^\dagger(b) := \begin{cases} (2/3) - b & b > 1/3, \\ -b & b < 1/3; \end{cases}$$

and we have a weaker lower bound on the normal type's payoff.

In an arbitrary game  $R^\dagger$  is equal to the seller's worst payoff from playing as the strong type when the buyer plays a response that gives him more than

his minmax payoff. In certain cases this can be a very strong restriction. For example, if the seller has a pure strategy that minmaxed the buyer and there is a unique response for the buyer that ensured he received his minmax payoff. Certain games, known as games of conflicting interests, have the property that the best action for the seller to commit to is pure and minmaxes the buyer.  $R^\dagger$  is a very tight bound for such games.

**Reputation with discounting: one long run buyer** This final case combines most of the above issues. If the seller discounts the future much less than the buyer, then in the long run the seller must get  $R^\dagger(b)$  from playing  $\hat{\sigma}$ . If a normal seller pretends to be strong the buyers think there are at most  $N$  periods when the strong strategy is not played. Imagine now we have a buyer who cares only about what happens in the next  $t'$  periods. Such a buyer can think there are at most  $t'N$  periods in which  $\hat{\sigma}$  is not played for the next  $t'$  periods. (This kind of argument is due to Schmidt 1993). As the seller becomes very patient  $Nt'$  periods becomes of vanishing importance and the normal seller's payoff is bounded below by  $R^\dagger(b)$ . If the seller and buyer discount equally, however, reputation effects cannot be found except in some very special cases.

## 5. Imperfect monitoring: temporary and bad reputation

The analysis of reputation given above presupposes perfect monitoring by the buyers and sellers of each other's actions. In many dynamic and repeated games this not likely — see entry on repeated games. To what extent do the above results continue to hold when the players are not able to see exactly what their opponent did in any one period? Perhaps reputations are harder to establish if the observed behavior is noisy? On the other hand maybe deviations from the strong type's action are harder to detect and so reputations last longer and are more valuable.

The merging, right ballpark and finite surprises properties all hold true under imperfect monitoring, with a suitable redefinition, provided there is enough statistical information for the buyer to eventually identify the seller's behavior. (This is a full-rank condition on the players' signals.) As a result, the bounds on payoffs given in the previous section continue to hold.

Under imperfect monitoring with adequate statistical information there is one new behavioral feature of these games — reputation is almost always temporary, that is, the buyer will eventually get to know the seller's type. To see why this is so, let us amend the game in Figure 1 by restricting the buyer to imperfectly observe the seller's action. With probability  $1 - \varepsilon$

the buyer observes the seller's true action in the current period, but with probability  $\varepsilon$  he observes the reverse action. (We must also assume the buyer does not see his own payoffs, otherwise he can deduce the seller's action from his payoff.) Consider a game where the seller always provides high quality ( $b = 0$ ) and suppose that reputation is permanent in such a game. Then  $p$  would, at least some of the time, converge to a number that is not zero or one. (Remember beliefs have to converge.) The merging property tells us that, in this case when the limit of beliefs is between zero and one, the buyer will be certain the normal seller is always providing high quality. Such buyers will ignore the occasional low quality product as just unlucky outcomes and there will be no loss of seller reputation if the buyer ever receives low quality. The normal type of seller can, therefore, deviate from always providing high quality, gain one unit of profit, and not face any costs in terms of loss of reputation. This cannot be an equilibrium. The initial claim that reputation is permanent has to be false as a result of this contradiction. The details of this argument can be found in Cripps, Mailath and Samuelson (2004).

When the monitoring is not statistically informative "bad reputation", due to Ely and Valimaki (2003), is a possibility. Uninformative monitoring is a particular problem in repeated extensive form games, because players do not get to see the actions their opponent would have taken on other branches of the game tree. Bad reputation may arise in our example if the buyer could take an action (such as not buy) that stopped the seller being able to signal her type. Then, the normal seller might find herself permanently stuck in a situation where she cannot sell. This is not particularly surprising if the buyers were very convinced they faced a strong seller that almost always provided low quality. However, in certain circumstances this problem is much more severe: even if the buyers were almost certain the seller were normal every equilibrium has trade ending in a bounded and finite time. Thus it is possible that introducing something for the seller to signal has huge negative costs for her equilibrium payoffs. To illustrate this, suppose the seller were a restaurant with imperfect control over quality although it does have a strategy (for example doubling the butter and salt content!) that makes it more likely the buyer will think the meal he received is good — but is actually damaging to the buyer. When play has reached the position where just one more bad meal will lead the buyer to permanently avoid the restaurant, then the restaurant will choose to use this unhealthy strategy. Knowing this, the buyer will choose to go elsewhere for his last but one meal too and there is an unraveling of the putative equilibrium. Buyers only eat at the restaurant if they get very few bad meals, because they know they are

in for clogged arteries and high blood pressure after that. Bad reputation arises because the seller cannot resist the temptation of taking actions that are actually unfavorable to the buyer in an effort to regain his good opinion. They actually have the reverse effect of ultimately driving the buyers away.

## 6. Reputation for what?

In our discussion we consider a strong type of seller who is committed to playing a particular fixed (random) action in each period. Is this form of uncertainty the only relevant one, or are there other potential types of strong seller that may do even better for our normal seller? There are two alternatives to consider: the strong seller is committed to playing a history-dependent strategy, or the strong player is equipped with a payoff function and her strategy is determined by an equilibrium.

If the seller faces a sequence of short term buyers then committing to a fixed stage game action is the best she could ever do, because each buyer's optimization focusses on what the seller does in the current period — the future is irrelevant. Even when the buyers are long lived there are circumstances where committing to play a fixed action imparts a strategic advantage in repeated play, for example in most coordination or common interest games. However, there are other repeated games, such as the prisoners' dilemma, and dynamic games where committing to a fixed stage action is worthless. What the seller would like to do is to commit to a strategy, such as tit-for-tat, which would persuade a sufficiently patient buyer to cooperate with the strong type. Provided some rather strong conditions are satisfied this is possible.

Our recipe for reputation results will break down when we consider strong sellers with payoffs rather than actions, nevertheless, reputation results are possible. For example, if the strong seller had payoffs of two for high quality and zero for low quality he would be strategically identical to a seller who always provided high quality.

## 7. Many players: social reputation and other considerations

Thus far we have resolutely stuck to a model of two players, but it is clear that reputation is a pervasive social and competitive phenomenon. Below we will sketch some of the issues in many-player reputation. The literature on this area is in its infancy, very little can be said with much certainty now.

The easiest case to deal with is what happens as the number of uninformed players (the buyers in our example) increase. Here the benefit to

the seller of building a reputation for high quality increases, as providing a good product today means the seller is more likely to trade with many buyers tomorrow. In a way, increasing the number of buyers is like making the seller more patient and so we would expect the seller to be more inclined to build a reputation in this case.

A second case would be where there are very large numbers of informed buyers trying to acquire reputations for individual or group characteristics. Models of career concerns are similar to reputation models and have many workers trying to acquire reputations for individual characteristics. Also, there are models of group reputation, such as Tirole (1996), where a particular class of individuals behaves in a particular way to perpetuate the “group’s” reputation. In both these types of model the large numbers assumption allows one individual’s reputation decision to be treated as virtually independent of others. Thus they can be analyzed using quite simple tools.

A final case is where a few informed agents are in competition or collusion with each other. Collusion in team reputation obviously introduces a public goods issue. If one player contributes to the good name of the group he or she does not get to enjoy the full benefits of the contribution. Typically, therefore, reputations for such teams are harder to establish. One might conjecture that competition appears to drive a player towards excessive investment in reputation, but there are many effects at work that we do not completely understand. For example, competitors’ may also act to undermine their rival’s reputation and to interfere with its development. This is a fertile region for applied and theoretical investigations.

## References

- Cripps, M.W., G.J. Mailath and L. Samuelson 2004. Imperfect monitoring and impermanent reputations. *Econometrica*, 72, 407–432.
- Ely, J. and J. Valimaki 2003. Bad reputation. *Quarterly Journal of Economics*, 118, 785–814.
- Fudenberg, D., D. Kreps and E. Maskin 1990. Repeated games with long-run and short-run players. *Review of Economic Studies*, 57, 555–574.
- 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.
- Hart, S. 1985. Nonzero-sum two-person repeated games with incomplete information. *Mathematics of Operations Research*, 10, 117–153.

- Kreps, D. and R. Wilson 1982. Reputation and imperfect information. *Journal of Economic Theory*, 27, 253–279.
- Mailath, G.J. and L. Samuelson 2006. *Repeated games and reputations: long-run relationships*. Oxford: Oxford University Press.
- Milgrom, P. and J. Roberts 1982. Predation, reputation and entry deterrence. *Journal of Economic Theory*, 27, 280–312.
- Schmidt, K. 1993. Reputation and equilibrium characterization in repeated games of conflicting interests. *Econometrica*, 61, 325–351.
- Sorin, S. 1999. Merging, reputation, and repeated games with incomplete information. *Games and Economic Behavior*, 29, 274–308.
- 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.
- Williams, D. 1991. *Probability with Martingales*. Cambridge: Cambridge University Press.