

Reputation Mechanisms

Chrysanthos Dellarocas

R. H. Smith School of Business

University of Maryland

College Park, MD 20742

cdell@rhsmith.umd.edu

June 2005

Revised: August 2005

Abstract

Reputation mechanisms harness the bi-directional communication capabilities of the Internet in order to engineer large-scale word-of-mouth networks. Best known so far as a technology for building trust and fostering cooperation in online marketplaces, such as eBay, these mechanisms are poised to have a much wider impact on organizations. This paper surveys our progress in understanding the new possibilities and challenges that these mechanisms represent. It discusses some important dimensions in which Internet-based reputation mechanisms differ from traditional word-of-mouth networks and surveys the most important issues related to their design, evaluation, and use. It provides an overview of relevant work in game theory and economics on the topic of reputation. It discusses how this body of work is being extended and combined with insights from computer science, marketing, and psychology in order to take into consideration the special properties of online environments. Finally, it identifies opportunities that this new area presents for information systems research.

1 Introduction

The Internet offers buyers and sellers unprecedented opportunities to access resources that were difficult, or impossible, to locate in previous times. Whether these resources are obscure

books, highly specialized services, or faraway trading partners, the global connectivity of the Web has the potential to bring them within everyone's reach, significantly enriching our economic and cultural lives.

Several of the mechanisms through which cooperation is induced in offline settings, such as the legal system and stable partnerships, do not work as well on the global, decentralized Internet (Kollock, 1999). To fully reap the benefits of global connectivity and arms-length transactions with faraway partners our societies thus need to develop new trust mechanisms capable of ensuring cooperation and efficiency in a universe of strangers.

Reputation networks constitute an ancient solution to the problem of trust-building. The historical appeal of these networks has been their power to induce cooperation without the need for costly enforcement institutions. Before the establishment of formal law and centralized systems of contract enforcement backed by the sovereign power of a state, most ancient and medieval communities relied on reputation as the primary enabler of economic and social activity (Benson, 1989; Greif, 1993; Milgrom, North, and Weingast, 1990). Many aspects of social and economic life still do so today (Klein, 1997).

These most ancient of mechanisms are emerging as one of the most promising solutions to the problem of building trust on the Internet. *Online reputation mechanisms*, also known as *reputation systems* (Resnick et al., 2000; Dellarocas, 2003a), are using the Internet's bi-directional communication capabilities in order to artificially engineer large-scale word-of-mouth networks where individuals share opinions and experiences on a wide range of topics, including companies, products, services, and even world events.

For example, eBay's feedback mechanism is the primary means through which eBay elicits honest behavior and, thus, facilitates transactions among strangers over the Internet (Resnick and Zeckhauser, 2002). Several other communities also rely on reputation mechanisms to promote trust and cooperation. Examples include eLance (online community of freelance professionals), Slashdot (online discussion forum where reputation scores help prioritize and filter postings), and Epinions (online consumer report forum where user feedback helps evaluate the quality of product reviews). Table 1 lists several noteworthy examples of such mechanisms in use today.

Online reputation mechanisms have a lot in common with their offline counterparts. Their design and implementation, thus, has a lot to gain from a substantial body of prior work on reputation formation in economics and psychology. On the other hand, online mechanisms possess a number of unique properties, whose implications are not yet fully understood. Specifically:

Web Site	Category	Summary of feedback mechanism	Format of solicited feedback	Format of published feedback
Citysearch	Entertainment guide	Users rate restaurants, bars, clubs, hotels and shops.	Users rate multiple aspects of reviewed items from 1-10 and answer a number of yes/no questions; readers rate reviews as “useful”, “not useful”, etc.	Weighted averages of ratings per aspect reflecting both user and editorial ratings; user reviews can be sorted according to “usefulness”
eBay	Online auction house	Buyers and sellers rate one another following transactions	Positive, negative or neutral rating plus short comment; ratee may post a response	Sums of positive, negative and neutral ratings received during past 6 months (see Section 3)
eLance	Professional services marketplace	Contractors rate their satisfaction with subcontractors	Numerical rating from 1-5 plus comment; ratee may post a response	Average of ratings received during past 6 months
Epinions	Online opinions forum	Users write reviews about products/services; other members rate the usefulness of reviews	Users rate multiple aspects of reviewed items from 1-5; readers rate reviews as “useful”, “not useful”, etc.	Averages of item ratings; % of readers who found a review “useful”
Google	Search engine	Search results are ordered based on how many sites contain links that point to them (Brin and Page 1998)	A Web page is rated based on how many links point to it, how many links point to the pointing page, etc.	No explicit feedback scores are published; ordering acts as an implicit indicator of reputation
Slashdot	Online discussion board	Postings are prioritized or filtered according to the ratings they receive from readers	Readers rate posted comments	

Table 1: Examples of commercial reputation mechanisms (in use as of June 2005).

Global reach enables new applications. Scale is essential to the effectiveness of reputation networks. In an online marketplace, for example, sellers care about buyer feedback primarily to the extent that they believe that it might affect their future profits; this can only happen if feedback is provided by a sufficient number of current customers and communicated to a significant portion of future prospects. Theory predicts that a minimum degree of participation in reputation communities is required before reputation effects can induce any cooperation. Once this threshold is reached, however, the power of reputation immediately springs to life and high levels of cooperation emerge in a discontinuous fashion (Bakos and Dellarocas, 2002). Therefore, the vastly increased scale of Internet-based reputation mechanisms is likely to render them powerful institutions in environments where traditional word-of-mouth networks were heretofore considered ineffective devices¹. The social, economic and perhaps even political consequences of such a trend deserve careful study.

Information technology enables systematic design. In offline settings word-of-mouth emerges naturally and evolves in ways that are difficult to control or model. The Internet allows this powerful social force to be precisely measured and controlled through proper engineering of the information systems that mediate online reputation communities. Such automated *feedback mediators* specify who can participate, what type of information is solicited from participants, how it is aggregated and what type of information is made available to them about other community members. Through the proper design of these mediators, mechanism designers can exercise precise control over a number of parameters that are very difficult or impossible to influence in brick-and-mortar settings. For example, feedback mediators can replace detailed feedback histories with a wide variety of summary statistics; they can apply filtering algorithms to eliminate outlier or suspect ratings; they can weight ratings according to some measure of the rater's trustworthiness, etc. Such degree of control can impact the resulting social outcomes in non-trivial ways. Understanding the full space of design possibilities and the impacts of specific design choices on the resulting social outcomes is an important research challenge introduced by these new systems.

Online interaction introduces new challenges. The disembodied nature of online environments introduces several challenges related to the interpretation and use of online feedback.

¹Three recent incidents illustrate the growing power of online opinion forums to exert influence on corporations and other powerful institutions of our society. In December 2002 criticism of controversial remarks made by U.S. Senator Trent Lott by authors of Web Logs (blogs) eventually led to his resignation from his post as majority leader. In 2003 Intuit Corporation was forced to remove unpopular copy protection spyware from its TurboTax software following a wave of very negative reviews posted by customers in online product forums. In September 2004 scrutiny by blog authors revealed inaccuracies in a story aired by long-serving and respected CBS anchor Dan Rather. The ensuing events culminated into Dan Rather's retirement from the channel.

Some of these challenges have their roots in the subjective nature of consumer feedback. Offline settings usually provide a wealth of contextual cues that assist in the proper interpretation of opinions and gossip (such as familiarity with the person who acts as the source of that information, the ability to draw inferences from the source's facial expression or mode of dress, etc.). Most of these cues are absent from online settings. Readers of online feedback are thus faced with the task of evaluating the opinions of complete strangers. Other challenges to feedback interpretation have their root in the ease with which online identities can be changed. This opens the door to various forms of strategic manipulation. For example, community members can build a good reputation, milk it by cheating other members and then disappear and reappear under a new online identity and a clean record (Friedman and Resnick, 2001). They can use fake online identities to post dishonest feedback and thus try to inflate their reputation or tarnish that of their competitors (Dellarocas, 2004b). Finally, the mediated nature of online reputation mechanisms raises questions related to the trustworthiness of their operators. An important prerequisite for the widespread acceptance of online reputation mechanisms is, therefore, a better understanding of how such systems can be compromised, as well as the development of adequate defenses.

This chapter surveys our progress so far in understanding the new possibilities and challenges that these mechanisms represent. Section 2 introduces a framework for understanding the role of reputation mechanisms in various settings. Section 3 provides an overview of relevant past work in game theory and economics. Section 4 then discusses how this stylized body of work is being extended in order to take into consideration the special properties of online environments. Sections 5 surveys empirical and experimental work on reputation mechanisms. Finally, Section 6 summarizes the main points of the paper and discusses the opportunities that this new area presents for information systems research.

2 Signaling and sanctioning role of reputation mechanisms

The primary objective of reputation mechanisms is to enable efficient transactions in communities where cooperation is compromised by post-contractual opportunism (moral hazard) or information asymmetries (adverse selection). It is instructive to distinguish the role of reputation mechanisms with respect to moral hazard from their role with respect to adverse selection.

Moral hazard can be present any time two parties come into agreement with one another. Each party in a contract may have the opportunity to gain from acting contrary to the principles laid out by the agreement. For example, on eBay, the buyer typically sends money to the seller before receiving the goods. The seller then is tempted to keep the money and not ship the goods, or to ship goods that are inferior to those advertised.

Reputation mechanisms can deter moral hazard by acting as *sanctioning devices*. If the community follows a norm that punishes traders with histories of bad behavior (by refusing to buy from them, or by reducing the price they are willing to pay for their products) and if the present value of punishment exceeds the gains from cheating, then the threat of public revelation of a trader's cheating behavior in the current round provides rational traders with sufficient incentives to cooperate.

Adverse selection is present in situations where sellers have information (about some aspect of their innate ability, product quality, etc.) that buyers don't (or vice versa). Such situations often arise in markets for experience goods. Consider, for example, an online hotel booking site where hotels of different qualities advertise rooms. Consumers cannot be certain about the true quality offered by each hotel until they have actually stayed there. On the other hand, hotels do not have an incentive to advertise any of their weak points. Knowing this, consumers will assume that all hotels are of average quality and will not be willing to pay more than the average price. Akerlof (1970) shows that such a situation will eventually drive all, except the lowest quality sellers, out of the market.

Reputation mechanisms alleviate adverse selection issues by acting as *signaling devices*. For example, by soliciting and publishing experiences of consumers who have stayed in advertised hotels, they help the community learn the true quality of each hotel. This, in turn, allows a better matching of buyers and sellers and a more efficient market.

The most important distinction between (pure) moral hazard and (pure) adverse selection settings is that, in the former, all sellers are capable of the same type of behavior (e.g. cooperate, cheat), whereas in the latter case seller behavior is completely constrained by their innate "type." The role of reputation mechanisms in pure moral hazard settings is to constrain behavior whereas the role of such mechanisms in pure adverse selection settings is to induce learning.

In some real-life settings, moral hazard and adverse selection considerations are simultaneously present: Sellers differ in their intrinsic ability levels but, in addition, have a choice of behavior (which is partially, but not completely, conditioned by their type). For example,

certain attributes of the customer experience (location, size of rooms, etc.) can be considered as part of a hotel’s immutable “type”, whereas other attributes (cleanliness of facilities, professionalism and politeness of staff, etc.) are the result of the hotel’s level of “effort” and can be varied strategically on a daily basis. In such settings, reputation mechanisms play both a sanctioning and a signaling role, revealing the hotel’s true immutable attributes while providing incentives to the hotel to exert reasonable effort.

In other settings, one of the two roles is dominant. For example, *Amazon Reviews* primarily serves a signaling role: it spread information about the (initially privately known, but essentially “immutable”) qualities of the products (books, CDs, DVDs, etc.) being reviewed. eBay, on the other hand, is an example of a mechanism that primarily acts as a sanctioning device. Under the assumption that all eBay sellers are equally capable of acting in honest and dishonest ways, eBay’s problem is to deter moral hazard. Accordingly, eBay users do not rate sellers on the absolute quality of their products but rather on how well they were able to deliver what was promised on the item description. The role of eBay’s reputation mechanism is to promote honest trade rather than to distinguish sellers who sell high quality products from those that sell low quality products.

The distinction between sanctioning and signaling is central in reputation mechanisms. Throughout this article we shall see that several principles of reputation mechanism design depend on this distinction. Designers should, therefore, be conscious of their mechanism’s primary objective and carefully make design choices that maximize the resulting market efficiency, given that objective.

3 Reputation in game theory and economics

Reputation formation has been extensively studied by economists using the tools of game theory. This body of work is perhaps the most promising foundation for developing an analytical discipline of online reputation mechanisms. This section surveys past work on this topic, emphasizing the results that are most relevant to the design of online feedback mechanisms. Section 4 then discusses how this stylized body of work is being extended to address the unique properties of online environments.

3.1 Basic concepts

According to Wilson (1985), reputation is a concept that arises in repeated game settings when there is uncertainty about some property (the “type”) of one or more players in the mind of other players. If “uninformed” players have access to the history of past stage game outcomes, reputation effects then often allow informed players to improve their long-term payoffs by gradually convincing uninformed players that they belong to the type that best suits their interests. They do this by repeatedly choosing actions that make them appear to uninformed players as if they were of the intended type, thus “acquiring a reputation” for being of that type.

The existence of some initial doubt in the mind of uninformed players regarding the type of informed players is crucial in order for reputation effects to occur. To see this, consider a repeated game between a long-run player and a sequence of short-run (one-shot) opponents. In every stage game, the long-run player can choose one of several actions but cannot credibly commit to any of those actions in advance. If there is no uncertainty about the long-run player’s type², rational short-run players will then always play their stage-game Nash equilibrium response. Such behavior typically results in inefficient outcomes.

Consider, for example, the following stylized version of a repeated “online auction” game. A long-lived seller faces an infinite sequence of sets of identical one-time buyers in a marketplace where there are only two kinds of products:

- low-quality products that cost 0 to the seller and are worth 1 to the buyers, and
- high-quality products that cost 1 to the seller and are worth 3 to the buyers.

Each period the seller moves first, announcing the quality of the product he promises to buyers. Since high quality products are more profitable, the seller will always promise high quality. Buyers then compete with one another in a Vickrey auction and therefore bid amounts equal to their expected valuation of the transaction outcome. The winning bidder sends payment to the seller. The seller then has the choice of either “cooperating” (delivering a high quality good) or “cheating” (delivering a low quality good). It is easy to see that this game has a unique subgame perfect equilibrium. In equilibrium the seller always cheats

²In other words, if short-run players are convinced that the long-run player is a rational utility-maximizing player whose stage-game payoffs are known with certainty.

(delivers low quality), buyers each bid 1, each buyer's expected payoff is zero and the seller's expected payoff is 1.

The ability to build a reputation allows the long-run player to improve his payoffs in such settings. Intuitively, a long-run player who has a track record of playing a given action (e.g. cooperate) often enough in the past acquires a reputation for doing so and is "trusted" by subsequent short-run players to do so in the future. However, why would a profit-maximizing long-term player be willing to behave in such a way and why would rational short-term players use past history as an indication of future behavior?

To explain such phenomena, Kreps, Milgrom, Roberts and Wilson (1982), Kreps and Wilson (1982), and Milgrom and Roberts (1982) introduced the notion of "commitment" types. Commitment types are long-run players who are locked into playing the same action³. An important subclass of commitment types are Stackelberg types: long-run players who are locked into playing the so-called Stackelberg action. The Stackelberg action is the action to which the long-run player would credibly commit if he could. In the above "online auction" example the Stackelberg action would be to cooperate; cooperation is the action that maximizes the seller's lifetime payoffs if the seller could credibly commit to an action for the entire duration of the game⁴. Therefore, the Stackelberg type in this example corresponds to an "honest" seller who never cheats. In contrast, an "ordinary" or "strategic" type corresponds to an opportunistic seller who cheats whenever it is advantageous for him to do so.

Reputation models assume that short-run players know that commitment types exist, but are ignorant of the type of the player they face. An additional assumption is that short-run players have access to the entire history of past stage game outcomes⁵. A player's reputation at any given time then consists of the conditional posterior probabilities over that player's type, given a short-run player's prior over types and the repeated application of Bayes' rule on the history of past stage game outcomes.

In such a setting, when selecting his next move, the informed player must take into account not only his short-term payoff, but also the long-term consequences of his action based on

³Commitment types are sometimes also referred to as "irrational" types because they follow fixed, "hard-wired" strategies as opposed to "rational" profit-maximizing strategies. An alternative way to justify such players is to consider them as players with non-standard payoff structures such that the "commitment" action is their dominant strategy given their payoffs.

⁴If the seller could commit to cooperation (production of high quality), buyers would then each bid 2 and the seller's expected per period payoff would be 2.

⁵The traditional justification for this assumption is that past outcomes are either publicly observable or explicitly communicated among short-run players. The emergence of online reputation mechanisms provides, of course, yet another justification (but see discussion of complications arising from the *private* observability of outcomes in such systems in Section 4.1).

what that action reveals about his type to other players. As long as the promised future gains due to the increased (or sustained) reputation that comes from playing the Stackelberg action offset whatever short-term incentives he might have to play otherwise, the equilibrium strategy for an “ordinary” informed player will be to try to “acquire a reputation” by masquerading as a Stackelberg type (i.e. repeatedly play the Stackelberg action with high probability).

In the “online auction” example, if the promised future gains of reputation effects are high enough⁶, ordinary sellers are induced to overcome their short-term temptation to cheat and to try to acquire a reputation for honesty by repeatedly delivering high quality. Expecting this, buyers then place high bids, thus increasing the seller’s long-term payoffs.

In general, reputation effects benefit the most patient player in the game: the player who has the longest time horizon (discounts future payoffs less) is usually the one who is able to reap the benefits of reputation. Fudenberg and Levine (1992) show that this result holds even when players can observe only noisy signals of each other’s actions, so that the game has imperfect public monitoring. They prove that, if short-run players assign positive prior probability to the long-run player being a Stackelberg type, and if that player is sufficiently patient, then an ordinary long-run player achieves an average discounted payoff close to his commitment payoff (i.e., his payoff if he could credibly commit to the Stackelberg action). In order to obtain this payoff, the ordinary player spends long periods of time choosing the Stackelberg action with high probability⁷.

3.2 Reputation dynamics

In most settings where reputation phenomena arise, equilibrium strategies evolve over time as information about the types of the various players accumulates. In general, the derivation of closed-form solutions in repeated games with reputation effects is complicated. Nevertheless, a small number of specific cases have been studied. The general lesson is that reputation-based performance incentives are highly dynamic: agents tend to behave differently in different phases of the game.

⁶In this type of game this requires that the remaining horizon of the seller is long enough and that the profit margin of a single transaction is high enough relative to the discount factor.

⁷This result also requires that the stage game is either a simultaneous move game, or, in a sequential-move game, that the short-run players always observe whether or not the Stackelberg strategy has been played.

Initial phase In most cases, reputation effects begin to work immediately and in fact are strongest during the initial phase, when players must work hard to establish a reputation. Holmstrom (1999) discusses an interesting model of reputational considerations in the context of an agent's "career" concerns. Suppose that wages are a function of an employee's innate ability for a task. Employers cannot directly observe an employee's ability. However, they can keep track of the average value of past task outputs. Outputs depend both on ability and labor. The employee's objective is to maximize lifetime wages while minimizing the labor put in. At equilibrium, this provides incentives to the employee to work hard right from the beginning of a career in order to build a reputation for competence. In fact these incentives are strongest at the very beginning of a career when observations are most informative.

During the initial phase of a repeated game, it is common that some players realize lower or even negative profits, while the community "learns" their type. In those cases players will only attempt to build a reputation if the losses from masquerading as a Stackelberg type in the current round are offset by the present value of the gains from their improved reputation in the later part of the game. In trading environments, this condition usually translates to the need for sufficiently high profit margins for "good quality" products so that the promise of future gains from sustaining a reputation is persuasive enough to offset the short-term temptation to cheat. This was first pointed out by Klein and Leffler (1981) and explored more formally by Shapiro (1983).

Another case where reputation effects may fail to work is when short-run players are "too cautious" vis-à-vis the long-run player and therefore update their beliefs too slowly in order for the long-run player to find it profitable to try to build a reputation. Such cases may occur when, in addition to Stackelberg ("good") types, the set of commitment types also includes "bad" or "inept" types: players who always play the action that the short-run players like least. In the "online auction" example, a "bad" type corresponds to a player who always cheats (because, for example, he lacks the capabilities that would enable him to deliver high quality.) If short-run players have a substantial prior belief that the long-run player may be a "bad" type, then the structure of the game may not allow them to update their beliefs fast enough to make it worthwhile for the long-run player to try to acquire a reputation.

Diamond's (1989) analysis of reputation formation in debt markets presents an example of such a setting. In Diamond's model there are three types of borrowers: safe borrowers, who always select safe projects (i.e. projects with zero probability of default); risky borrowers, who always select risky projects (i.e. projects with higher returns if successful but with nonzero probability of default); and strategic borrowers who will select the type of project

that maximizes their long term expected payoff. The objective of lenders is to maximize their long term return by offering competitive interest rates, while at the same time being able to distinguish profitable from unprofitable borrowers. Lenders do not observe a borrower's choice of projects, but they do have access to her history of defaults. In Diamond's model, if lenders believe that the initial fraction of risky borrowers is significant, then, despite the reputation mechanism, at the beginning of the game, interest rates will be so high that strategic players have an incentive to select risky projects. Some of them will default and will exit the game. Others will prove lucky and will begin to be considered as safe players. It is only after lucky strategic players have already acquired some initial reputation (and therefore begin to receive lower interest rates) that it becomes optimal for them to begin "masquerading" as safe players by consciously choosing safe projects in order to sustain their good reputation.

Steady state (or lack thereof) In their simplest form, reputation games are characterized by an equilibrium in which the long-run player repeatedly plays the Stackelberg action with high probability and the player's reputation converges to the Stackelberg type.

The existence of such steady states crucially depends on the ability to perfectly monitor the outcomes of individual stage games. For example, consider the "online auction" game that serves as an example throughout this section with the added assumption that buyers perfectly and truthfully observe and report the seller's action. In such a setting, the presence of even a single negative rating on a seller's feedback history reveals the fact that the seller is not honest. From then on, buyers will always choose the low bid in perpetuity. Since such an outcome is not advantageous for the seller, reputation considerations will induce the seller to cooperate forever.

The situation changes radically if monitoring of outcomes is imperfect. In the online auction example, imperfect monitoring means that even when the seller produces high quality, there is a possibility that a buyer will post a negative rating, and, conversely, even when the seller produces low quality, the buyer may post a positive rating. A striking result is that in such "noisy" environments reputations cannot be sustained indefinitely: if a strategic player stays in the game long enough, short-run players will eventually learn his true type and the game will inevitably revert to one of the static Nash equilibria (Cripps, Mailath and Samuelson, 2004).

To see the intuition behind this result, note that reputations under perfect monitoring are typically supported by a trigger strategy. Deviations from the equilibrium strategy reveal

the type of the deviator and are punished by a switch to an undesirable equilibrium of the resulting complete-information continuation game. In contrast, when monitoring is imperfect, individual deviations neither completely reveal the deviator's type nor trigger punishments. A single deviation has only a small effect on the beliefs of the short-term players. As a result, a player of normal type trying to maintain a reputation as a Stackelberg type incurs only a small cost (in terms of altered beliefs) from indulging in occasional deviations from Stackelberg play. In fact, it is clear that always playing the Stackelberg action cannot be an equilibrium strategy, because if the short-term players expect long-term players of normal type to behave that way, then they can actually deviate at no cost, since any bad outcome will be interpreted by the short-run players as a result of imperfect monitoring. But the long-run effect of many such small deviations from the commitment strategy is to drive the equilibrium to full revelation.

These dynamics have important repercussions for reputation systems in settings with both moral hazard and adverse selection (for example, eBay, under the assumption that there exist rational and honest seller types). According to the Cripps, Mailath and Samuelson result, if eBay makes the entire feedback history of a seller available to buyers and if an eBay seller stays on the system long enough, once he establishes an initial reputation for honesty he will be tempted to cheat buyers every now and then. In the long term, this behavior will lead to an eventual collapse of his reputation and therefore of cooperative behavior. I revisit the implications of this result for reputation mechanism design in Section 4.2.

Endgame considerations Since reputation relies on a tradeoff between current “restraint” and the promise of future gains, in finitely repeated games, incentives to maintain a reputation diminish and eventually disappear as the end of the game comes close.

One solution to this problem is to assign some post-mortem value to reputation, so that players find it optimal to maintain it throughout the game. For example, reputations can be viewed as assets that can be bought and sold in a market for reputations. Tadelis (1998) shows that a market for reputations is indeed sustainable. Furthermore, the existence of such a market provides “old” agents and “young” agents with equal incentives to exert effort (Tadelis, 2002). However, the long-run effects of introducing such a market can be quite complicated since good reputations are then likely to be purchased by “inept” agents for the purpose of depleting them (Mailath and Samuelson, 2001; Tadelis, 2002). Further research is needed in order to fully understand the long-term consequences of introducing markets for reputation as well as for transferring these promising concepts to the online domain.

3.3 When is reputation bad?

In traditional reputation theory, publication of a long-term player's past history of outcomes is good for the long-term player. One, therefore, is tempted to assume that implementation of reputation mechanisms is always a "good thing". Ely, Fudenberg, and Levine (2005), henceforth referred to as EFL, challenge this assumption and show that there exist settings where the presence of public histories of past outcomes is unambiguously bad.

EFL generalize an example provided by Ely and Valimaki (2003), henceforth referred to as EV. EV describe a setting where a mechanic of unknown "character" (rational, dishonest) is facing a sequence of customers who bring their cars to him for repair. Each car might need a tune-up or an engine change; only the mechanic is able to determine the correct type of repair. Rational mechanics are assumed to have a payoff structure that induces them to perform the correct type of repair in a one-stage game. Dishonest mechanics, on the other hand, always perform engine changes.

Assume, now, that customers have access to the history of past repairs performed by the mechanic and will only contract with him if they are sufficiently confident that he is going to perform the correct type of repair. Given the above assumptions, if there is a positive prior probability that the mechanic is dishonest, after histories with many engine changes, the short-run players will become sufficiently convinced they are facing such a bad type and exit. In order to avoid these histories, a rational mechanic who has had the "bad luck" of having many customers that need an engine change may then begin to recommend tune-ups to customers who need engine changes; foreseeing this, the short-run players will choose not to enter.

Observe that, whereas, in the absence of reputation, rational types play friendly actions, the presence of a public history of past outcomes induces them to behave in ways that short-run players like least. Foreseeing this, short-run players choose not to enter; the presence of reputational information then causes the market to break down.

EFL show that settings where reputation is bad are characterized by the following properties: (i) a long-run player of privately-known type is facing an infinite sequence of short-run opponents, (ii) commitment types include "unfriendly" types who play "unfriendly" actions, that is, actions that short-run players dislike, (iii) short-run players will only enter the game if they are sufficiently confident that the long-run player is going to play one of the friendly actions, (iv) there exist "bad signals" that are most likely to occur when unfriendly actions are played but also occur with positive probability when friendly actions are played (in EV

the bad signal is “engine change”). Finally, (v) there are some actions that are not friendly, but reduce the probability of bad signals (such as “always perform tune-up” in EV).

The main result of EFL is that, in a bad reputation game with a sufficiently patient long-run player and likely enough unfriendly types, in any Nash equilibrium, the long-run player gets approximately the payoff that results from non-participation (exit) of short-run players.

EFL show that bad reputation games arise in various “expert advice” settings. This includes consulting a doctor or stockbroker, or in the macroeconomics context, the decision whether or not to turn to the IMF for assistance. In EV, the short-run players observe only the advice, but not the consequences of the advice. EFL consider what happens when the short-run players observe the consequences as well. They show that the bad reputation effect persists so long as this additional information is not perfectly accurate.

The EFL/EV result delivers a pessimistic message with respect to the application of reputation mechanisms in expert advice settings, an area where the author believes that such mechanisms have a lot to offer. On the other hand, the result crucially depends on the assumption that short-term players will exit as soon as the long-term player’s reputation falls below a threshold. If the long-term player can persuade them to participate (by charging lower prices or, perhaps, paying them a participation reward) the result breaks down: if a rational player who has had a stream of “bad luck” has a way to restore his reputation, he will continue to engage in friendly actions. Nevertheless, reputation mechanism designers must be aware of the EFL result and its implications. It is important to understand the range of practical settings in which these results might apply, and therefore to identify classes of settings for which the development of online reputation mechanisms may not be a good idea⁸.

3.4 Other extensions to the basic theory

The basic theory assumes that uninformed players are short-term. Facing longer-lived opponents may be worse for the informed player and generally results in less sharp predictions

⁸A related, but not identical, problem arises when the establishment of a reputation mechanism induces long-run players to change their behavior in ways that improve their payoff but reduce social welfare. Dranove et al. (2003) examine the consequences of public disclosure of patient health outcomes at the level of the individual physician and/or hospital. The intention behind this measure was to address informational asymmetries in markets for health care. However, it also gave doctors and hospitals incentives to decline to treat more difficult, severely ill patients. Using national data on Medicare patients at risk for cardiac surgery, Dranove et al. find that cardiac surgery report cards in New York and Pennsylvania led both to selection behavior by providers and to improved matching of patients with hospitals. On net, this led to higher levels of resource use and to worse health outcomes, particularly for sicker patients.

about reputation effects (Cripps and Thomas, 1995; Cripps, Schmidt and Thomas, 1996). Quite interestingly, however, in repeated games where a patient player faces one or more long-lived but less patient opponents, if the more patient player does not observe the less patient players' intended actions but only sees an imperfect signal of them, reputation effects once again become strong and result in lower bounds that are even higher than in the case where all opponents are myopic (Celentani, Fudenberg, Levine and Pesendorfer, 1996). This last case is equivalent to a situation where a long-run player faces a sequence of long-run but "infrequent" players. This is, perhaps, an even more realistic model of relationships in online communities and therefore an area that deserves further study.

Another assumption underlying most game theoretic models of reputation is that all players have identical prior beliefs and that behavior is consistent with the concept of Bayesian Nash equilibrium. These assumptions are probably too stringent and unrealistic in environments as diverse as large-scale online communities. Fortunately, reputation phenomena arise under significantly weaker assumptions on the knowledge and behavior of players. Watson (1993; 1996) and Battigalli and Watson (1997) demonstrated that reputation effects do not require equilibrium. They are implied by a weak notion of rationalizability along with two main conditions on the beliefs of players: First, there must be a strictly positive and uniform lower bound on the subjective probability that players assign to the Stackelberg type. Second, the conditional beliefs of short run players must not be "too dispersed".

4 New opportunities and challenges of online mechanisms

I began this chapter by discussing a number of differences between online reputation mechanisms and traditional word-of-mouth networks. This section surveys our progress in understanding the opportunities and challenges that these special properties imply.

4.1 Eliciting sufficient and honest feedback

Most game theoretic models of reputation formation assume that stage game outcomes (or imperfect signals thereof) are *publicly* observed. Most online reputation mechanisms, in contrast, rely on *private* monitoring of stage game outcomes and voluntary self-reporting⁹.

⁹For a comprehensive introduction to games with imperfect private monitoring see Kandori (2002) and other papers contained in the same issue.

This introduces two important new considerations: (a) ensuring that sufficient feedback is, indeed, provided, and (b) inducing truthful reporting.

Economic theory predicts that voluntary feedback will be underprovided. There are two main reasons for this. First, feedback constitutes a public good: once available, everyone can costlessly benefit from it. Voluntary provision of feedback leads to suboptimal supply, since no individual takes account of the benefits that her provision gives to others. Second, provision of feedback presupposes that the rater will assume the risks of transacting with the ratee. Such risks are highest for new products. Prospective consumers may, thus, be tempted to wait until more information is available. However, unless somebody decides to take the risk of becoming an early evaluator, no feedback will ever be provided.

Avery, Resnick and Zeckhauser (1999) analyze mechanisms whereby early evaluators are paid to provide information and later evaluators pay in order to balance the budget. They conclude that, of the three desirable properties for such a mechanism (voluntary participation, no price discrimination and budget balance), any two can be achieved, but not all three¹⁰.

Since in most reputation mechanisms, monitoring of transaction outcomes is private, an additional consideration is whether feedback is honest. No generally applicable solution to this important problem currently exists. Nevertheless, several researchers have proposed mechanisms that induce truth-telling in restricted settings.

Jurca and Faltings (2004) propose a mechanism that limits false reporting in settings with pure moral hazard. They consider a bilateral exchange setting where long-run buyers and long-run sellers transact repeatedly. Each period a buyer is asked to rate a transaction only if the corresponding seller has claimed to have successfully provided the service. If the two reports disagree, at least one of the traders must be lying; the center then fines both transacting parties different (fixed) amounts¹¹. Jurca and Faltings show that, if a buyer always reports feedback on a particular seller truthfully, the seller also finds it optimal to truthfully report transaction outcomes, with the exception of a finite number of transactions.

¹⁰Empirical evidence has, so far, not confirmed the (rather pessimistic) predictions of theory. Consumer participation in online feedback mechanisms is surprisingly high, even though, in most cases, such mechanisms offer no concrete participation incentives. Such behavior is consistent with a large body of empirical evidence (Dichter, 1966; Engel, Blackwell and Miniard, 1993; Sundaram, Mitra and Webster, 1998; Hennig-Thuray et al., 2004) that has identified a variety of extra-economic motivations to explain why consumers engage in (offline and online) word-of-mouth (desire to achieve social status, utility from engaging in social interaction, altruism, concern for others, easing anger, dissonance reduction, vengeance, etc.). I return to this point in Section 5.

¹¹Fines can be implemented by levying refundable listing fees from the two traders at the beginning of each transaction and confiscating these fees in the case of conflicting reports.

Papaioannou and Stamoulis (2005) propose a similar mechanism that is suitable for Peer-to-Peer environments. In such environments side-payments are usually not possible and peers are able to exchange roles. After each transaction both peers submit a rating, with each peer not knowing the rating submitted by the other. A credibility metric is maintained for each peer regarding his overall truthfulness record in rating transactions. In case of agreement between two peers, the credibility metric of each peer is improved. In case of disagreement, the credibility metric of each peer is deteriorated and both peers are punished. Punishment amounts to not allowing a peer to transact with others for a period that is exponential to the peer’s credibility metric. This is enforced by publicly announcing a peer’s punishment and by punishing other peers when they transact with him. Performance of this mechanism is analyzed by means of experiments in dynamically evolving peer-to-peer systems with renewed populations.

Miller, Resnick and Zeckhauser (2005) propose mechanisms for eliciting honest feedback in environments with pure adverse selection. Their mechanisms are based on the technique of proper scoring rules (Cooke, 1991). A scoring rule is a method that induces rational agents to truthfully reveal their beliefs about the distribution of a random variable by rewarding them on the basis of how a future realization of the random variable relates to the distribution they announced. A proper scoring rule has the property that the agent maximizes his expected score when he truthfully announces his beliefs.

Assuming that: (i) users rate a set of products, whose types (e.g. qualities) remain fixed over time, (ii) each type maps to a distribution of outcomes (e.g. satisfaction levels perceived by consumers) that can be statistically distinguished from that of every other type, and (iii) all raters and the center have a common set of prior beliefs over types, or, alternatively, each rater’s prior beliefs have been communicated to the center, MRZ propose a side-payment mechanism that asks each rater to announce the outcome a_1 she observed, and rewards her by an amount proportional to $\log p(a_2|a_1)$ where a_2 is the outcome reported by a future rater. MRZ show the existence of a Nash equilibrium where such a side-payment rule induces truth-telling. They also show that their mechanism can be extended to not only induce truth-telling, but also ratings of a given precision.

The above mechanisms represent promising first steps towards ensuring the credibility of online feedback. On the other hand, they have several limitations. First, such mechanisms only work as long as raters are assumed to act independently; all break down if raters collude. Second, in addition to the desirable truth-telling equilibria., all three mechanisms induce additional equilibria where agents do not report the truth. Equilibrium selection,

thus, becomes an important consideration in practical implementations. The development of robust mechanisms for eliciting truthful feedback thus remains one of the most important open areas of research.

4.2 Exploring the design space of feedback mediators

Information technology has added a large degree of flexibility and precision to the design of reputation mechanisms. Online mechanism designers can control a number of important parameters that are difficult, or impossible, to influence in offline settings. Examples of such parameters include the format of solicited feedback (eBay allows traders to rate a transaction as "positive", "negative" or "neutral", Amazon Auctions supports integer ratings from 1 to 5, other systems support even higher levels of detail), the amount and type of information included in a trader's reputation profile (most systems publish the sum or arithmetic mean of all posted ratings, some systems highlight recent ratings, other systems provide access to a trader's entire ratings history) as well as the frequency with which reputation profiles are updated with new information (most current systems make new ratings publicly available as soon as they are posted). These parameters impact the consequences of a trader's current behavior on the community's perception of him in the future, and thus, his incentives to cooperate. This section summarizes our current understanding of how such design choices affect trader behavior and market efficiency.

Granularity of feedback. eBay's mechanism solicits ternary feedback (positive, neutral, negative) while Amazon Auctions asks traders to rate transactions on a finer-grained scale of 1-5. What is the impact of the granularity of feedback on the seller's incentives to cooperate? In the special case of settings with two seller actions (cooperate, cheat) and pure moral hazard, Dellarocas (2005a) shows that the equilibrium that maximizes cooperation occurs when buyers divide feedback of arbitrary granularity into two disjoint subsets (the "good ratings" subset and the "bad ratings" subset) and behave as if feedback was binary, i.e. they reward the seller by the same amount if he receives any rating that belongs to the "good" set and they punish him by the same amount if he receives any rating that belongs to the "bad" set. Maximum efficiency is then inversely proportional to the minimum likelihood ratio of punishment if the seller cooperates vs. if he cheats over all possible ways of dividing ratings into "good" and "bad" subsets¹².

¹²In more general settings where both adverse selection and moral hazard are present the answer is generally more complex. Soliciting (and publishing) finer-granularity feedback might increase or decrease a seller's incentives to cooperate. See Dewatripont, Jewitt, and Tirole (1999) for a precise statement of the relevant conditions.

Length of published feedback history. How much history should a mechanism publish about a trader’s past behavior? The answer here crucially depends on the type of setting (moral hazard, adverse selection, or combined). In pure adverse selection settings the goal of the mechanism is to promote social learning. More information is, thus, always better. In contrast, in pure moral hazard settings Dellarocas (2005a) shows that the maximum efficiency that can be induced by a reputation mechanism is independent of the length of published history: a mechanism that publishes a trader’s entire history performs as well as a mechanism that only publishes the trader’s single most recent rating. The intuition behind this result is that, in pure moral hazard settings, the goal of the reputation mechanism is to threaten players with *future* punishment if the public outcome of the *current* transaction indicates that they cheated. The length of time during which a rating persists in the seller’s reputation profile only affects the duration of future punishment but not the total amount of punishment. This makes efficiency independent of the length of published history.

Finally, in settings where both moral hazard and adverse selection are present, if the objective of the mechanism is to induce long-term seller cooperation, Cripps, Mailath and Samuelson’s result (see Section 3.2) suggests that reputation mechanisms should not publish a seller’s entire history. (Otherwise, once sellers have established a good reputation, they will be tempted to occasionally cheat; in the long-run, this will reveal their opportunistic nature and will drive the system to the undesirable one-shot Nash equilibrium.) On the other hand, since their role is to promote some degree of learning about the seller’s type, reputation mechanisms should publish some amount of past history. Fan, Tan and Whinston (2005) explore the idea of discounting past ratings in such settings. Using simulation, they show that such policies can sustain a seller’s incentives to cooperate. General design guidelines for deciding what is the optimal length of history (or, equivalently, what is the optimal discount factor of past ratings) constitute an interesting open area of research.

Frequency of feedback profile updating. Most systems in use today update a trader’s reputation profile with new evaluations as soon as these are posted by users. Dellarocas (2005b) shows that, in pure moral hazard settings, this is not necessarily the optimal architecture. Specifically, he shows that, if ratings are noisy and the per-period profit margin of cooperating sellers sufficiently high, a mechanism that does not publish every single rating it receives but rather, only updates a trader’s public reputation profile every k transactions with a summary statistic of a trader’s most recent k ratings, can induce higher average levels of cooperation and market efficiency than a mechanism that publishes all ratings as soon as they are posted. The intuition behind the result is that basing each profile update on a trader’s most recent k -transaction rating history reduces the impact of spurious negative

ratings (because these are amortized over k transactions). On the other hand, it also increases the trader's temptation to cheat (because he can cheat for k period before news of his behavior become public). The optimal updating period k is derived from a tradeoff between these two opposite forces.

4.3 Coping with cheap online identities

In online communities it is usually easy for members to disappear and re-register under a completely different online identity with zero or very low cost. Friedman and Resnick (2001) refer to this property as “cheap pseudonyms.” This property hinders the effectiveness of reputation mechanisms. Community members can build a reputation, milk it by cheating other members and then vanish and re-enter the community with a new identity and a clean record.

Friedman and Resnick discuss two classes of approaches to this issue: either make it more difficult to change online identities, or structure the community in such a way that exit and re-entry with a new identity becomes unprofitable. The first approach makes use of cryptographic authentication technologies and is outside the scope of this paper. The second approach is based on imposing an upfront cost to each new entrant, such that the benefits of “milking” one's reputation are exceeded by the cost of subsequent re-entry. This cost can be an explicit entrance fee or an implicit cost of having to go through a reputation building (or “*dues paying*”) stage with low or negative profits. Friedman and Resnick show that, although dues paying approaches incur efficiency losses, such losses constitute an inevitable consequence of easy name changes.

Dellarocas (2005a) shows how such a “dues paying” approach can be implemented in pure moral hazard trading environments. He proves that, in the presence of easy name changes, the design that results in optimal social efficiency is one where newcomers pay an entrance fee and the mechanism only publishes a trader's single most recent rating. Dellarocas further demonstrates that, although this design incurs efficiency losses relative to the case where identity change is not an issue, in settings with two possible transaction outcomes and opportunistic sellers that can freely change identities, its efficiency is the highest attainable by any mechanism.

4.4 Understanding the consequences of strategic manipulation

As online reputation mechanisms begin to exercise greater influence on consumer behavior the incentive for strategically manipulating them becomes correspondingly stronger. The low cost of submitting online feedback coupled with the relative anonymity of the rater makes such manipulation a real problem that needs to be studied and addressed.

Dellarocas (2000; 2004a) has pointed out a number of manipulation scenarios and has proposed a number of immunization mechanisms, based on ideas from clustering and robust statistics, that reduce the effect of such attacks if the fraction of unfair raters is reasonably small (up to 20-30% of the total population). A number of commercial reputation mechanisms (for example, Amazon Reviews and Epinions) are attempting to address such problems through the concept of “rate the rater”: members can rate how useful other members’ feedback has been to them. Whereas this technique is somewhat effective for separating high-quality from low-quality postings, it is not effective for reducing strategic manipulation. Determined manipulators can manipulate the “rate the rater” ratings as much as they can manipulate the ratings themselves.

The ability to “cheaply” create multiple online identities further complicates attempts to combat manipulation attacks. For example, an early technique through which eBay traders attempted to manipulate their reputation profile was for colluding traders to form a ring, engaging in repeated “fake” transactions with each other for the purpose of artificially boosting each other’s feedback score. To combat such attacks, eBay is now only counting unique ratings received from frequent partners towards a trader’s reputation profile. Such defenses can be defeated, however, if traders can easily create multiple identities. Douceur (2002) has called this the *sybil attack*. Cheng and Friedman (2005) provide formal definitions of *sybil-proofness*, i.e. resistance to sybil attacks, and prove the impossibility of developing sybilproof reputation mechanisms in a broad class of settings.

Since it appears that complete elimination of reputation mechanism manipulation will remain a formidable task for some time, it is useful to explore the consequences of such activity in strategic settings. In settings with pure adverse selection, Mayzlin (2003) has analyzed the impact of strategic manipulation in the context of Usenet groups where consumers discuss products and services. Mayzlin’s setting involves two rival firms that wish to influence consumer beliefs regarding the quality of their respective products by posting costly “fake” promotional messages in Usenet groups. Mayzlin’s basic result is that, if the ratio of profits to manipulation cost is high enough, there exists an equilibrium in which both firms manipulate

but the low quality firm manipulates more. Promotional chat thus decreases, but does not completely destroy, the informativeness of online forums. Dellarocas (2004b) extended Mayzlin’s results in more general settings. His most interesting new result is that there exist equilibria where strategic manipulation can *increase* forum informativeness. Such equilibria arise in settings where firm revenues are sufficiently convex functions of consumer perceptions of their quality. In such settings, the presence of honest consumer opinions can induce firms to reveal their own, more precise, knowledge of product qualities by manipulating the forums at relative intensities that are proportional to their actual qualities. The impact of manipulation then is to further separate the published ratings of the high quality firm from those of the low quality firm, making it easier for consumers to infer each firm’s true quality.

4.5 Distributed reputation mechanisms

Traditional reputation theory assumes the existence of a public history of past outcomes. This implies a centralized architecture where outcomes are either automatically recorded or explicitly self-reported. It also implies the presence of a trusted mediator who controls feedback aggregation and distribution. Though the design possibilities of even that simple architecture are not yet fully understood, centralized reputation mechanisms do not nearly exhaust the new possibilities offered by information technology.

The growing importance of Peer-to-Peer (P2P) networks (Oram, 2001) is introducing new challenges to reputation mechanism design. In P2P networks every entity can act both as a provider and consumer of resources. Entities are assumed to be self-interested and cannot be trusted to engage in cooperative behavior unless concrete incentives are in place. For example, in file-sharing P2P networks, self-interested entities have short-term incentives to free-ride (consume as much content as they can without contributing any content themselves) or to contribute low-quality content. Furthermore, there is, usually, no central, universally trusted entity that can act as a repository of reputational information.

To cope with these challenges, several researchers have proposed decentralized reputation mechanisms. Two lines of investigation stand out as particularly promising:

Reputation formation based on analysis of “implicit feedback.” Traditional reputation mechanisms rely on explicit solicitation of feedback from transaction participants. If reliable explicit feedback is not available, information about an agent’s type can often be inferred by analyzing publicly available attributes of the network in which the agent is embedded.

Perhaps the most successful application of this approach to date is exemplified by the Google search engine. Google’s PageRank algorithm assigns a measure of reputation to each web page that matches the keywords of a search request. It then uses that measure to rank order search hits. Google’s page reputation measure is based on the number of links that point to a page, the number of links that point to the pointing page, and so on (Brin and Page, 1998). The underlying assumption is that if enough people consider a page to be important enough in order to place links to that page from their pages, and if the pointing pages are “reputable” themselves, then the information contained on the target page is likely to be valuable. Google’s success in returning relevant results is testimony to the promise of that approach.

Pujol , Sangüesa and Delgado (2002) apply network flow techniques in order to propose a generalization of the above algorithm that “extracts” the reputation of nodes in a general class of social networks. Sabater and Sierra (2002) describe how direct experience, explicit and implicit feedback can be combined into a single reputation mechanism.

Basing reputation formation on implicit information is a promising solution to problems of eliciting sufficient and truthful feedback. Careful modeling of the benefits and limitations of this approach is needed in order to determine in what settings it might be a viable substitute or complement of voluntary feedback provision.

Distributed reputation mechanisms. The majority of decentralized reputation mechanisms proposed so far are based on variations on the theme of *referral networks* (Aberer and Despotovic, 2001; Kamvar et al., 2003; Xiong and Liu, 2004; Yu and Singh, 2002). Since no centralized reputation repository exists, each agent in a P2P network solicits referrals about a target agent from a set of neighbors who might, in turn, ask their neighbors, and so on. Referrals are weighted by the relative amounts of trust that the soliciting agent places on his neighbors’ advice and combined with personal experiences that the soliciting agent might have had with the target agent in the past. Based on his subsequent experience with the target agent, the soliciting agent dynamically adapts both his beliefs regarding the target agent, as well as the amounts by which she trusts her neighbors’ recommendations¹³.

Most work in this area is currently based on heuristics and evaluated using simulation. The development of rigorous results on the efficiency bounds and design principles of distributed reputation mechanisms constitutes an important open area of research.

¹³See Despotovic and Aberer (2004) for a more comprehensive overview of trust-building mechanisms in P2P networks.

5 Empirical and experimental studies

A growing literature of empirical and experimental studies constitutes an essential complement to game-theoretic and simulation-based analyses of reputation mechanisms. This section surveys some of the main themes of such work.

5.1 Empirical studies and field experiments

Most empirical studies of reputation mechanisms have focused on eBay's feedback mechanism. The majority of early studies has looked at the buyers' response to seller reputation profiles. In particular, a large number of studies estimate cross-sectional regressions of sale prices and probabilities of sale on seller feedback characteristics. Table 2 summarizes the main results of these studies.

The following points summarize the principal conclusions derived from a collective reading of these works:

- Feedback profiles seem to affect both prices and the probability of sale. However, the precise effects are ambiguous; different studies focus on different components of eBay's complex feedback profile and often reach different conclusions.
- The impact of feedback profiles on prices and probability of sale is relatively higher for riskier transactions and more expensive products.
- Among all different pieces of feedback information that eBay publishes for a member, the components that seem to be most influential in affecting buyer behavior are the overall number of positive and negative ratings, followed by the number of recently (last 7 days, last month) posted negative comments.

Resnick, Zeckhauser, Swanson and Lockwood (2005) point out the potential for significant omitted variable bias in these cross-sectional regressions. They argue that the price difference commanded by sellers of higher reputation might be due to several other factors that tend to exhibit positive correlation with a seller's eBay score (for example, the clarity and aesthetics of the item listing, the professionalism of the seller's email communications, etc.). To better assess the premium attached to reputation they conduct a controlled field experiment in which a seasoned seller sells identical postcards using his real name and an assumed name.

Shorthand	Citation	Items sold	Remarks
BP	Ba and Pavlou 2002	Music, Software, Electronics	Positive feedback increased estimated price, but negative feedback did not have an effect
BH	Bajari and Hortacsu 2003	Coins	Both positive and negative feedback affect probability of modeled buyer entry into the auction, but only positive feedback had a significant effect on final price
DH	Dewan and Hsu 2004	Stamps	Higher net score increases price
E	Eaton 2002	Electric guitars	Negative feedback reduces probability of sale, but not price of sold items
HW	Houser and Wooders 2005	Pentium chips	Positive feedback increases price; negative feedback reduces it
KM	Kalyanam and McIntyre 2001	Palm Pilot PDAs	Positive feedback increases price; negative feedback reduces price
KW	Kauffman and Wood 2000	Coins	No significant effects, but negative feedback seems to increase price (!) in univariate analysis
LIL	Lee, Im and Lee 2000	Computer monitors and printers	Negative feedback reduces price, but only for used items
L	Livingston 2002	Golf clubs	Positive feedback increases both likelihood of sale and price; effect tapers off once a record is established
LBDP	Lucking-Reiley et al. 2000	Coins	No effect from positive feedback; negative feedback reduces price
MA	Melnik and Alm 2002	Gold coins	Positive feedback increases price; negative feedback decreases price
MS	McDonald and Slawson 2002	Dolls	Higher net score (positives -negatives) increases price
RZ	Resnick and Zeckhauser 2002	MP3 players, Beanie babies	Both forms of feedback affect probability of sale but not price contingent on sale
RZSL	Resnick Zeckhauser, Swanson and Lockwood 2005	Vintage postcards	Controlled field experiment; established seller commands higher prices than newcomers; among newcomers, small amounts of negative feedback have little effect

Table 2: Summary of early empirical studies on eBay (adapted from Resnick et al., 2005).

They find an 8% premium to having 2000 positive feedbacks and 1 negative over a feedback profile with 10 positive comments and no negatives.

In another field experiment, Jin and Kato (2004) assess whether the reputation mechanism is able to combat fraud by purchasing un-graded baseball cards with seller-reported grades, and having them evaluated by the official grading agency. They report that, while having a better seller reputation is a positive indicator of honesty, reputation premia or discounts in the market do not fully compensate for expected losses due to seller dishonesty.

Using panel data, Cabral and Hortacsu (2005) analyze the impact of reputation on eBay sellers. They find that, when a seller first receives negative feedback, his weekly sales rate drops from a positive 7% to a negative 7%; subsequent negative feedback ratings arrive 25% more rapidly than the first one and don't have as much impact as the first one. They also find that a seller is more likely to exit the lower his reputation is; and that, just before exiting, sellers receive more negative feedback than their lifetime average.

Cabral and Hortacsu then consider a series of theoretical models (pure moral hazard, pure adverse selection, combined moral hazard and adverse selection) and measure them against these empirical results. They are not able to draw definite conclusions as to which theoretical model best explains the data; they do, however, conclude that eBay's reputation system gives way to noticeable strategic responses from both buyers and sellers.

Another interesting area of empirical work relates to the motivations for participation in online feedback mechanisms. Economic theory predicts that, being public goods, evaluations will be underprovided unless evaluators are provided with concrete incentives (Avery, Resnick and Zeckhauser, 1999). Nevertheless, in systems such as eBay, more than 50% of transactions receive feedback even though no direct incentives to rate are in place. Dellarocas, Fan and Wood (2003b) use survival analysis techniques to study the motivations and dynamics of voluntary feedback submission on eBay. They find that reciprocity is an important driver of participation on eBay's mechanism: a trader's propensity to rate a transaction increases after her partner posts a rating. They also find that silence is often an indicator of an unsatisfactory transaction and propose an empirical method for inferring the true number of unsatisfactory transactions on eBay from the observed patterns of feedback submission. Using their method, they estimate that, whereas the percentage of negative feedback on eBay is less than 1%, the fraction of transactions where, at least one of the two parties remained unsatisfied is closer to 10%.

Product review forums (such as *Amazon Reviews*) have, so far, received less attention. Nevertheless, their increasing popularity is attracting more empirical research. Hennig-Thurau

et al. (2004) conduct an online survey to identify what motivates consumers to rate products in such forums. Their study identifies four significant motives. In order of decreasing significance the four motives are: social benefits (genuine fun that results from the social experience of participating in an online dialog), economic incentives (rewards offered by some sites for contributing content), concern for other consumers (urge to help others by recommending good products and warning against bad products) and extraversion/self-enhancement (positive feelings that result from sharing one’s successes with others; enhancement of one’s self-image by projecting oneself as intelligent shopper).

Another question of theoretical and practical consequence is the extent to which such forums influence consumer behavior. Obtaining reliable answers to this question is challenging. eBay has the advantage of being a “closed universe”: all transactions and feedback are mediated through the site and are publicly visible. On the other hand, in the general case consumers can obtain product ratings from one forum and purchase the product from another site, or from an offline store. It is, thus, difficult to establish causality between ratings and sales.

The approach of Chevalier and Mayzlin (2003) constitutes a step in the right direction. Chevalier and Mayzlin examine the effect of consumer reviews on relative sales of the same set of books on Amazon.com and BarnesandNoble.com, two large online booksellers¹⁴. They find that an improvement in a book’s reviews on one site leads to an increase in relative sales at that site. They also find that reviews are overwhelmingly positive and, thus, that the impact of (less common) negative reviews is greater than the impact of positive reviews.

5.2 Controlled experiments

Controlled experiments are a useful complement to empirical studies, especially in relation to exploring individual-level perceptions of and reactions to reputation mechanisms.

Keser (2003) reports a number of experiments based on the “trust game.” In this game, one player (the “buyer”) can choose to send money to a second (the “seller”). This amount is then substantially increased and the seller can choose to share some of the gain with the buyer. By removing many of the complexities involved in market transactions, this game provides a simple context to study the effect of different information policies about revealing past behaviors. Keser finds that the presence of a reputation mechanism significantly increases both the buyers’ level of investment (trust) as well as the fraction that sellers share with

¹⁴Book sales are inferred from the publicly available “sales rank” metric using an approach introduced by Chevalier and Goolsbee (2003).

their buyers (trustworthiness). Furthermore, she finds that both trust and trustworthiness are higher when the reputation mechanism publishes the entire history of each player's past behavior than when it publishes each player's behavior in the most recent transaction only.

Questions involving reputation's effect on market efficiency require more complex experimental scenarios. Bolton, Katok and Ockenfels (2002) study trading in a fixed-price market where buyers can choose whether to send the purchase price to the seller and sellers have the option of not fulfilling their contracts. They compare transaction completion rates in a setting with random matching of players and public histories of trader fulfillment (reputation), to a setting with random matching without reputation, as well as to a market where the same people interact with each other repeatedly (partners market). They find that, while the presence of the reputation mechanism induces a substantial improvement in trading efficiency, it falls short of the efficiency achieved in the partners market.

Chen, Hogg, and Wozny (2004) conduct experiments similar to those performed by BKO but provide a broader set of endogenous choices for the players. First, players can explicitly decide who they wish to do business with rather than being paired with a single other player by the experimenter. Second, both buyers and sellers make fulfillment choices and so face a moral hazard for which reputations are relevant. Third, in addition to settings with automatically generated reputation, CHW examine games where players self-report transaction outcomes, allowing them to misreport their experiences as possible punishment for a poor report on their own reputation. Fourth, prices and trading volumes are determined endogenously. The findings of CHW are consistent with the previous experiments: the presence of reputational information led to a substantial increase of transaction fulfillment. Interestingly, CHW found fulfillment rates to be almost as high when traders self-reported transaction outcomes as when reliable reputational information was automatically generated, indicating that, in practice, private observability of outcomes might not be as big a problem as theory suggests.

6 Conclusions: Opportunities for IS research

Online reputation mechanisms harness the remarkable ability of the Web to not only disseminate, but also collect and aggregate information from large communities at very low cost, in order to artificially construct large-scale word-of-mouth networks. Best known so far as a technology for building trust and fostering cooperation in online marketplaces, these mechanisms are poised to have a much wider impact on organizations.

For example, in a marketing context, certain types of reputation mechanisms (e.g., product discussion forums) can serve as a low-cost channel for acquiring and retaining customers, complementary to advertising (Mayzlin, 2003). In addition, they can assist organizations better understand consumer reactions to their current product line. At the same time, these systems affect an industry's competitive dynamics because they reveal the same information to all competitors. In a supply chain context, reputation mechanisms can assist organizations better assess the trustworthiness of prospective first-time suppliers, and thus enable more flexible supply chain management. At the same time they can become a powerful disciplining mechanism that can ensure fulfillment of contractual obligations without recourse to litigation; they, thus, have the potential of lowering the legal costs of doing business in a variety of industries (Bakos and Dellarocas, 2002).

The design of online reputation mechanisms can greatly benefit from the insights produced by more than twenty years of economics and game theory research on the topic of reputation. These results need to be extended to take into account the unique new properties of online environments, such as their unprecedented scalability, the ability to precisely design the type of feedback information that is solicited and distributed, and the relative ease with which traders can change online identities. The following list contains what the author considers to be the most important open areas of research in reputation mechanism design:

- Scope and explore the design space and limitations of online reputation mechanisms. Understand what set of design parameters work best in what settings. Develop models and prototype implementations of such systems.
- Develop effective solutions to the problems of sufficient participation, easy identity changes and strategic manipulation of online feedback.
- Conduct theory-driven experimental and empirical research that sheds more light on buyer and seller behavior vis-à-vis such mechanisms.
- Compare the relative efficiency of reputation mechanisms to the efficiency of more established mechanisms for dealing with moral hazard and adverse selection (such as the legal system and advertising);¹⁵ develop theory-driven guidelines for deciding which set of mechanisms to use when.

¹⁵See Cooter and Rubinfeld (1989) for an introduction to economic models of legal disputes. See Milgrom and Roberts (1986) for the role of advertising in revealing a firm's unknown quality.

- Understand how decision-makers must adapt their strategies to react to the presence of such mechanisms in areas such as marketing, product development, and customer service.

The power of online reputation mechanisms has its roots in the strategic side effects brought about by the increased interdependencies they create among firms, their customers, their partners, and their competitors. As the mathematical study of interaction of self-interested agents, game theory is the natural foundation for the study of online reputation mechanisms.

Other established paradigms of information systems research can play an important role in translating the conceptual insights of game theoretic models into concrete guidelines for building (and reacting to) large-scale reputation mechanisms that can influence the dynamics of entire industries or societies. For example, computational methods can help analyze games that may be too complex to solve analytically. Laboratory experiments can inform about how people will behave when confronted with these mechanisms, both when they are inexperienced and as they gain experience. Finally, game-theoretic models can often be approximated by generally more tractable, decision-theoretic approaches¹⁶.

There is much work to be done. But it is important that research be conducted now, in the formative phases of this technology and the social practices surrounding it. There are likely to be path-dependent effects in the deployment and use of online reputation mechanisms, so it is important that researchers develop insights into the functioning and impacts of these systems while they can still have a large impact on practice.

References

Aberer, K. and Despotovic, Z., 2001, Managing trust in a Peer-2-Peer information system, *Proc. 2001 ACM International Conf. on Information and Knowledge Management (CIKM)*, 310-317.

Akerlof, G., 1970, The market for lemons: Quality uncertainty and the market mechanism, *Quarterly Journal of Economics* 84, 488-500.

¹⁶See Zacharia et al. (2001) for an example of using simulation modeling to study the effects of reputation mechanisms on markets with dynamic pricing. See Dellarocas (2001) and Shapiro (1982) for examples of how the adoption of “reasonable” (common-sense or empirically-driven) assumptions about the behavior of other players permits the use of decision-theoretic approaches to model reputation phenomena.

- Avery, C., Resnick, P., R. Zeckhauser, 1999, The Market for Evaluations. *American Economics Review* 89 (3) 564-584.
- Ba, S., P. Pavlou, 2002, Evidence of the Effect of Trust Building Technology in Electronic Markets: Price Premiums and Buyer Behavior. *MIS Quarterly* 26 (3) 243-268.
- Bajari, P., A. Hortacsu, 2003, Winner's Curse, Reserve prices and Endogenous Entry: Empirical Insights from eBay Auctions. *Rand Journal of Economics* 34 (2).
- Bakos, Y., C. Dellarocas, 2002, Cooperation without Enforcement? A Comparative Analysis of Litigation and Online Reputation as Quality Assurance Mechanisms, L. Applegate, R. Galliers, J. I. DeGross, eds. *Proc. 23rd Int. Conf. on Information Systems (ICIS 2002)*. Association for Information Systems, Barcelona, Spain, 127-142.
- Battigalli, P., and Watson, J., 1997, On Reputation Refinements with Heterogeneous Beliefs, *Econometrica* 65, 369-374.
- Benson, B., 1989, The Spontaneous Evolution of Commercial Law, *Southern Economic Journal* 55 (January) 644-661. Reprinted in D. B. Klein, ed. 1997. *Reputation: Studies in the Voluntary Elicitation of Good Conduct*. University of Michigan Press, Ann Arbor, 165-189.
- Bolton, G. E., Katok, E., A. Ockenfels, 2004, How Effective are Online Reputation Mechanisms? An Experimental Study, *Management Science* 50 (11), 1587-1602.
- Brin S., L. Page, 1998, The Anatomy of a Large-Scale Hypertextual Web Search Engine. *Computer Networks and ISDN Systems* 30 (1-7) 107-117.
- Cabral, L., and Hortacsu, A. 2004, The Dynamics of Seller Reputation: Theory and Evidence from eBay (April 2004), CEPR Discussion Paper No. 4345. <http://ssrn.com/abstract=541161>
- Celentani, M., Fudenberg, D., Levine, D., and Pesendorfer, W., 1996, Maintaining a Reputation against a Long Lived Opponent, *Econometrica* 64, 691-704.
- Chen, K-Y, Hogg, T., and Wozny, N., 2004, Experimental study of market reputation mechanisms, *Proc. ACM Conference on Electronic Commerce 2004*, 234-235.
- Cheng, A. and Friedman, E., 2005, Sybilproof Reputation Mechanisms, *Proc. of the 3rd Workshop on the Economics of Peer-2-Peer Systems (P2PECON)*, Philadelphia, PA.
- Chevalier, J. A. and Goolsbee, A., 2003, Price Competition Online: Amazon Versus Barnes And Noble, *Quantitative Marketing and Economics* 1(2), 203-222.

Chevalier, J. A. and Mayzlin, D., 2003, The Effect of Word of Mouth on Sales: Online Book Reviews, Yale SOM Working Paper No's. ES-28 & MK-15. <http://ssrn.com/abstract=432481>

Cooke, R. M., 1991, *Experts in Uncertainty: Opinion and Subjective Probability in Science*. Oxford University Press, New York, NY.

Cooter, R.D. and Rubinfeld, D. L., 1989, Economic Analysis of Legal Disputes and Their Resolution, *Journal of Economic Studies* 27 (3), 1067-1097.

Cripps, M., Mailath, G., L. Samuelson, 2004, Imperfect Monitoring and Impermanent Reputations, *Econometrica* 72, 407-432.

Cripps, M., Schmidt, K., and Thomas, J., 1996, Reputation in Perturbed Repeated Games, *Journal of Economic Theory* 69, 387-410.

Cripps, M., and Thomas, J., 1995, Reputation and Commitment in Two-Person Repeated Games without Discounting, *Econometrica* 63, 1401-1419.

Dellarocas, C., 2000, Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior, *Proceedings of the 2nd ACM Conference on Electronic Commerce*. Association for Computing Machinery, Minneapolis, MN, 150-157.

Dellarocas, C., 2001, Analyzing the economic efficiency of eBay-like online reputation reporting mechanisms, *Proceedings of the 3rd ACM Conference on Electronic Commerce*. Association for Computing Machinery, Tampa, FL, 171-179.

Dellarocas, C., 2003a, The Digitization of Word-of-Mouth: Promise and Challenges of Online Feedback Mechanisms, *Management Science* 49 (10), 1407-1424.

Dellarocas, C., Fan, M., and Wood, C., 2003b, Self-Interest, Reciprocity, and Participation in Online Reputation Systems, *2003 Workshop in Information Systems and Economics (WISE)*, Seattle, WA.

Dellarocas, C., 2004a, Building Trust On-Line: The Design of Robust Reputation Mechanisms for Online Trading Communities. G. Doukidis, N. Mylonopoulos, N. Pouloudi, eds. *Social and Economic Transformation in the Digital Era*. Idea Group Publishing, Hershey, PA.

Dellarocas C., 2004b, Strategic Manipulation of Internet Opinion Forums: Implications for Consumers and Firms, MIT Sloan Working Papers No. 4501-04. <http://ssrn.com/abstract=585404>

- Dellarocas C., 2005a, Reputation Mechanism Design in Online Trading Environments with Pure Moral Hazard, *Information Systems Research* 16 (2), 209-230.
- Dellarocas C., 2005b, How Often Should Reputation Mechanisms Update a Trader's Reputation Profile? Working Paper.
- Despotovic, Z. and Aberer, K., 2004, Possibilities for Managing Trust in P2P Networks. Swiss Federal Institute of Technology (EPFL) Technical Report IC/2004/84, Lausanne, Switzerland.
- Dewatripont, M. Jewitt, I. and Tirole, J., 1999, The Economics of Career Concerns, Part I: Comparing Information Structures, *Review of Economic Studies* 66 (1), 183-198.
- Dichter, E., 1966, How Word-of-Mouth Advertising Works, *Harvard Business Review* 44 (Nov-Dec), 147-166.
- Dewan, S. and Hsu, V., 2004, Adverse Selection in Reputations-Based Electronic Markets: Evidence from Online Stamp Auctions, *Journal of Industrial Economics* 52 (4), 497-516.
- Diamond, D., 1989, Reputation Acquisition in Debt Markets, *Journal of Political Economy* 97 (4) 828-862.
- Douceur, J. 2002, The sybil attack, *Proc. 1st International Workshop on Peer-to-Peer Systems (IPTPS '02)*, Cambridge, MA.
- Dranove, D., Kessler, D., McClellan, M., and Satterthwaite, M., 2003, Is More Information Better? The Effects of "Report Cards" on Health Care Providers, *Journal of Political Economy* 111 (3), 555-588.
- Eaton, D. H., 2002, Valuing Information: Evidence from Guitar Auctions on eBay, Department of Economics Working Paper #0201, Murray State University, Murray, KY.
- Ely, J. C., Fudenberg, D. and Levine, D. K., 2005, When is Reputation Bad? Working Paper.
- Ely, J. C. and Valimaki, J., 2003, Bad Reputation, *Quarterly Journal of Economics* 118 (3), 785-814.
- Fan, M., Tan, Y., and Whinston, A. B., 2005, Evaluation and Design of Online Cooperative Feedback Mechanisms for Reputation Management, *IEEE Trans. Knowledge and Data Engineering* 17 (3), 244-254.

- Friedman, E., P. Resnick, 2001, The Social Cost of Cheap Pseudonyms, *Journal of Economics and Management Strategy* 10 (1), 173-199.
- Fudenberg, D., D. Levine, 1992, Maintaining a Reputation when Strategies are Imperfectly Observed, *Review of Economic Studies* 59 (3) 561-579.
- Greif, A., 1993, Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition, *American Economic Review* 83 (June) 525-548.
- Hennig-Thurau T., Gwinner K.P., Walsh G., Gremler D.D., 2004, Electronic Word-of-Mouth via Consumer-Opinion Platforms: What Motivates Consumers to Articulate Themselves on the Internet? *Journal of Interactive Marketing* 18 (1), 38-52.
- Holmstrom, B., 1999, Managerial Incentive Problems: A Dynamic Perspective, *Review of Economic Studies* 66 (1) 169-182.
- Houser, D., J. Wooders, 2005, Reputation in Auctions: Theory and Evidence from eBay. Forthcoming *Journal of Economics and Management Strategy*.
- Jin, G. Z., and Kato, A., 2004, Blind Trust Online: Experimental Evidence from Baseball Cards, University of Maryland Working Paper.
- Jurca, R. and Faltings, B., 2004, CONFESS: Eliciting Honest Feedback without Independent Verification Authorities. *Sixth International Workshop on Agent Mediated Electronic Commerce (AMEC VI 2004)*, New York, USA, July 19 2004.
- Kalyanam, K., S. McIntyre, 2001, Returns to Reputation in Online Auction markets. Retail Workbench Working Paper W-RW01-02, Santa Clara University, Santa Clara, CA.
- Kamvar, S. D., Schlosser, M. T. and Garcia-Molina, H., 2003, The EigenTrust Algorithm for Reputation Management in P2P Networks, *Proc. of the 12th International World Wide Web Conference*, 640-651.
- Kandori, M., 2002, Introduction to Repeated Games with Private Monitoring, *Journal of Economic Theory* 102 (1), 1-15.
- Kauffman, R. J., C. Wood, 2000, Running up the Bid: Modeling Seller Opportunism in Internet Auctions. *Proc. Americas Conference on Information Systems (AMCIS 2000)*, Association for Information Systems, Long Beach, CA, 929-935.

- Keser, C., 2003, Experimental Games for the Design of Reputation Management Systems, *IBM Systems Journal* 42 (3), 498-506.
- Klein, B., K. Leffler, 1981, The Role of Market Forces in Assuring Contractual Performance, *Journal of Political Economy* 89 (4) 615-641.
- Klein, D., ed. 1997, *Reputation: Studies in the Voluntary Elicitation of Good Conduct*, University of Michigan Press, Ann Arbor, MI.
- Kollock, P., 1999, The Production of Trust in Online Markets. E.J. Lawler, M. Macy, S. Thyne, and H.A. Walker, eds. *Advances in Group Processes* (Vol. 16). JAI Press., Greenwich, CT.
- Kreps, D., Milgrom, P., Roberts, J., R. Wilson, 1982, Rational Cooperation in the Finitely Repeated Prisoners' Dilemma, *Journal of Economic Theory* 27 (2) 245-252.
- Kreps, D., R. Wilson, 1982, Reputation and Imperfect Information, *Journal of Economic Theory* 27 (2) 253-279.
- Lee, Z., Im, I., S. J. Lee, 2000, The Effect of Negative Buyer Feedback on Prices in Internet Auction Markets. W. J. Orlikowski, S. Ang, P. Weill, H. C. Krcmar, J. I. DeGross, eds. *Proc. 21st Int. Conf. on Information Systems* (ICIS 2000), Association for Information Systems., Brisbane, Australia, 286-287.
- Livingston, J., 2002, How Valuable is a Good Reputation? A Sample Selection Model of Internet Auctions. Working Paper, University of Maryland, College Park, MD.
- Lucking-Reiley, D., D. Bryan, et al., 2000, Pennies from eBay: the Determinants of Price in Online Auctions, Working Paper, University of Arizona, Tucson, AZ.
- Mailath, G. J., L. Samuelson, 2001, Who Wants a Good Reputation? *Review of Economic Studies* 68 (2) 415-441.
- Mayzlin, D., 2003, Promotional Chat on the Internet, Working Paper #MK-14, Yale School of Management, New Haven, CT.
- McDonald, C. G., V. C. Slawson, Jr., 2002, Reputation in an Internet Auction Market, *Economic Inquiry* 40 (3) 633-650.
- Melnik, M. I., J. Alm, 2002, Does a Seller's Reputation Matter? Evidence from eBay Auctions, *Journal of Industrial Economics* 50 (3) 337-349.

- Milgrom, P. R., North, D., B. R. Weingast, 1990, The Role of Institutions in the Revival of Trade: The Law Merchant, Private Judges, and the Champagne Fairs, *Economics and Politics* 2 (1) 1-23. Reprinted in D. B. Klein, ed. 1997. *Reputation: Studies in the Voluntary Elicitation of Good Conduct*. University of Michigan Press, Ann Arbor, 243-266.
- Milgrom, P., J. Roberts, 1982, Predation, Reputation and Entry Deterrence, *Journal of Economic Theory* 27 (2) 280-312.
- Milgrom, P. and Roberts, J., 1986, Price and Advertising Signals of Product Quality, *Journal of Political Economy* 94 (4), 796-821.
- Miller, N., Resnick, P., R. Zeckhauser, 2005, Eliciting Honest Feedback: The Peer-Prediction Method, forthcoming *Management Science*.
- Oram, A. (ed.), 2001, *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*, O'Reilly Media, Sebastopol, CA.
- Papaioannou, T. and Stamoulis, G. D., 2005, An Incentives' Mechanism Promoting Truthful Feedback in Peer-to-Peer Systems, *Proc. of IEEE/ACM CCGRID 2005 (Workshop on Global P2P Computing)*, May 2005.
- Pujol, J. M., Sanguesa, R., J. Delgado, 2002, Extracting reputation in multi agent systems by means of social network topology, *Proc. of the First Int. Joint Conference on Autonomous Agents and Multiagent Systems*, Association for Computing Machinery, Bologna, Italy, 467-474.
- Resnick, P., Zeckhauser, R., Friedman, E., Kuwabara, K., 2000, Reputation Systems, *Communications of the ACM* 43 (12) 45-48.
- Resnick, P., R. Zeckhauser, 2002, Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System. Michael R. Baye, ed. *The Economics of the Internet and E-Commerce (Advances in Applied Microeconomics, Vol. 11)*. JAI Press.
- Resnick, P., Zeckhauser, R., Swanson, J., K. Lockwood, 2005, The Value of Reputation on eBay: A Controlled Experiment, forthcoming *Experimental Economics*.
- Sabater, J., C. Sierra, 2002, Reputation and social network analysis in multi-agent systems, *Proc. of the First Int. Joint Conference on Autonomous Agents and Multiagent Systems*, Association for Computing Machinery, Bologna, Italy, 475-482.

- Shapiro, C., 1982, Consumer Information, Product Quality, and Seller Reputation, *Bell Journal of Economics* 13 (1) 20-35.
- Shapiro, C., 1983, Premiums for High Quality Products as Returns to Reputations, *Quarterly Journal of Economics* 98 (4) 659-680.
- Sundaram, D.S., Mitra, K., and Webster, C., 1998, Word-of-Mouth Communications: A Motivational Analysis, *Advances in Consumer Research* 25, 527-531.
- Tadelis, S., 1999, What's in a Name? Reputation as a Tradeable Asset, *The American Economic Review* 89 (3) 548-563.
- Tadelis, S., 2002, The Market for Reputations as an Incentive Mechanism, *Journal of Political Economy* 92 (2) 854-882
- Watson, J., 1993, A Reputation Refinement without Equilibrium, *Econometrica* 61, 199-205.
- Watson, J., 1996, Reputation in Repeated Games with No Discounting, *Games and Economic Behavior* 15, 82-109.
- Wilson, R., 1985, Reputations in Games and Markets. A. Roth, ed. *Game-Theoretic Models of Bargaining*, Cambridge University Press, Cambridge, UK, 27-62.
- Xiong, L., and L. Liu, 2004, PeerTrust: Supporting Reputation-Based Trust in Peer-to-Peer Communities, *IEEE Transactions on Knowledge and Data Engineering* 16 (7), 843-857.
- Yu, B. and M. Singh, 2002, An evidential model of distributed reputation management, *Proc. of the First Int. Joint Conference on Autonomous Agents and Multiagent Systems*, Association for Computing Machinery, Bologna, Italy, 294-301.
- Zacharia, G., Moukas, A., P. Maes, 2000, Collaborative Reputation Mechanisms in Electronic Marketplaces, *Decision Support Systems* 29 (4) 371-388.
- Zacharia, G., Evgeniou, T., Moukas, A., Boufounos, P., P. Maes, 2001, Economics of Dynamic Pricing in a Reputation Brokered Agent Mediated Marketplace, *Electronic Commerce Research* 1 (1) 85-100.