1 Introduction

Today we work for a reputation.
Tomorrow our reputation will work for us.

Russian saying

EBay and Other Stories

EBay's electronic market is one of the most successful Internet applications, generating 15 million new listings on the average day and allowing more than 90 million users worldwide to buy and sell all sorts of new and secondhand items (eBay 2009, 2010). On eBay, traders almost invariably do not know each other and usually will not have an encounter more than once. It would be easy to cheat on eBay—for example, you could fail to send the merchandise after having received the requested payment, or send a product of lower value than the one advertised, or change your mind on a purchase and not follow it through. If dishonesty were widespread, people would not trust each other and would not use the market. Beyond a certain level of mutual mistrust, eBay would not function.

However, an essential characteristic of eBay prevents this from happening. At the end of a transaction, both buyer and seller may (and often do) write a "feedback" (as it is called) on each other, which can be positive, neutral, or negative. The sum of each type of feedback received by each participant during the previous six months forms what amounts to an "index of reputation" that is visible to all. The availability of such an index provides a strong incentive to be honest: a seller who receives negative feedbacks from previous buyers would find it difficult to continue

trading, because prospective buyers would be wary of doing business with him. A buyer would also have a hard time doing his shopping, because sellers would not believe in his resolve. On eBay, fear of acquiring a bad reputation represents a strong incentive to be honest and efficient, and the observed outcome is that most people, in fact, honestly describe the merchandise that they plan to sell, ship it quickly, and pay their bills. The result is the success of eBay, one of the poster cases of the Internet age.

Beyond what may be concluded from impressionistic observations, rigorous studies have assessed the role of reputational effects, using data on eBay transactions together with appropriate econometric techniques. Resnick et al. (2006, table 1) summarize fifteen such studies. The broad picture that emerges is one where such effects exist and are relevant. More recent research also points in the same direction. For example, Resnick et al. (2006) set up an experiment where an established dealer with a good reputation is compared with a new dealer, with no reputation at all. They both provide exactly the same good and quality of service. The authors find that the seller with an established reputation enjoys a significant price premium. Cabral and Hortaçsu (2010) find that when a seller receives a negative feedback for the first time, his sales drop significantly, and they conclude that overall "the eBay reputation system gives way to noticeable strategic responses from both buyers and sellers."[1]

The case of eBay is a good starting point for our inquiry into the role of reputation in public governance. True, the outcome that we observe—a viable and thriving market—does not apply to *public* governance but to the private domain. It also turns out that there are more examples available of *private* governance where reputation plays an important role. One of the main theses of this book is that this divergence occurs precisely because the role of reputational considerations in *public* governance today is not as important as it could, and should, be. However, the eBay example does hint at two issues that we will encounter over and over again in our reasoning on public governance. First, reputational considerations may induce people to act in useful ways even without the presence of a formal institution threatening to punish them should they misbehave. What induces most users of eBay to be honest is not the fear of the police knocking at their door, should they cheat. For many, and possibly for most, an interiorized sense of honesty may certainly play a role in

this respect. However, the presence of eBay's widely visible feedbacks, and their consequent reputational effects, has a much more compelling role in guaranteeing the viability of eBay as a marketplace.

Secondly, the success of eBay hinges upon the presence of a communication technology: Internet and the Web. In general, all solutions to governance problems need appropriate technologies in order to function. For example, the Roman Empire would not have existed as we know it without its network of roads and an efficient postal system, allowing the transmission of information and orders from the capital to the legions stationed in the provinces. Today's aviation technology, which permits, for example, government representatives to meet frequently, is an integral part of the system of international relations. In the case of eBay in particular, the Internet is used to transmit information on the reputation of traders. Also, the eBay case illustrates that this reputational information typically has to be appropriately *organized*: though individual comments that people post on concluded transactions are accessible to all, it is the aggregation of information provided by a reputation index that is most useful.

We should exercise care when considering the role of technology within a governance model. If the Internet, or something similar to it, were not available, eBay could not exist. However, while being able to record information and to communicate is indispensable for reputation to play a role, obviously reputation information can also spread using different and much less sophisticated technologies. The following description of a well-studied historical episode serves as a convincing example of the relevance of informal reputation considerations in very different technological contexts.

In the early Middle Ages, a group of Jewish traders, originally based in Egypt, formed a coalition to better pursue their business (see Greif 1989 and 1993, on which this account is based). We know quite a lot about the dealings of the Maghribi traders, as they are known, because they are documented in an archive found in Fustat (Old Cairo, in Egypt), which contained letters and accounts of various types. The merchants lived scattered in several trade centers in the Muslim area of the western Mediterranean, and they helped one another in their work. Each trader could act either as a merchant or as a merchant's agent, in this case

supplying the merchant with the services needed to operate long-distance trade, such as loading and unloading the ship, paying customs and transportation fees, and so forth. Depending on the occasion, one member of the coalition acted as a merchant or as an agent. For a merchant, operating through agents had some important advantages over taking his own goods to market, because it kept costs down and reduced the overall risk, as many parallel undertakings could go on in different markets.

Where long-distance trade was administered through agents, however, there was ample opportunity for cheating, and the anticipation of widespread dishonesty would have been enough to discourage any trade. Just as eBay would be unviable as a market if potential participants did not trust each other enough, so it would have been for the activities of the Maghribi traders. As Avner Greif (1993) wrote, "To gain from cooperation, there was a need for an institution capable of surmounting this commitment problem, an institution through which an agent could commit himself *ex-ante*, before receiving the merchant's capital, to be honest *ex post*."

The legal system of the time was not capable of offering effective protection to merchants, and there was no *formal* institution that was capable of forcing agents to respect pacts. The solution the Maghribi traders came up with was to inflict a collective punishment in the form of a boycott on agents who cheated. Agents were also allowed to cheat those merchants who had been dishonest, without themselves being subject to collective punishment, thus reinforcing the threat to cheaters. The existence of this "reputational equilibrium," supported by historical documents, allowed the Maghribi traders to prosper.

The example of the Maghribi traders, to which we will return in the next chapter, illustrates that successful reputational mechanisms may use very different communication technologies. In fact, it has been suggested that the Internet "digitalizes word-of-mouth" (Dellarocas 2003), thus enhancing the possibilities of passing information (which traditionally was communicated informally) to others on the behavior of people or organizations. In the past, before choosing a restaurant or a hotel, we asked friends in the know for a tip, but now we access specialized Web sites providing ratings by former customers, such as www.zagat.com for restaurants or any of the many online hotel booking services available. In all

cases, the presence of a communication channel, transmitting information about past performances, stimulates the provision of a good service.

However, if we simply relegate the Internet to the category of instruments that allow us to digitize activities that humans have always carried out—for example, exchanging information on the performance of people and institutions—we run the risk of missing the novelty that it represents in this and many other respects. Today, it is not only the availability of Internet technologies that permits the working of information systems, but also the fact that they treat digitized reputational information in a coherent manner, compute useful measures of reputation, and disseminate the relevant information to all concerned parties, all at negligible cost. Also, one distinctive trait of these systems is that they allow the gathering and processing of reputation information to be specifically designed, unlike the spontaneous informal arrangements of the past. Today, reputational systems may be *engineered* far more than before, because their working is mapped into the architectural characteristic of a technological artifact, an Internet-based information system. The emerging possibility of crafting reputational systems to serve particular governance goals is an important theme of this book to which we will dedicate much thought.

The example of the Maghribi traders is evidence of the fact that reputational considerations have always played an important role in governance and, more generally, in human affairs. We can distinguish between two main beneficial effects of reputation: one that we may call "static" and the other "dynamic." Let us start with the static effect. A good reputation, for the owner, is the result of past behavior and is analogous to a capital good, which can be kept and increased by respecting one's obligations, but also easily wasted when shortsightedness prevails. Rational individuals are constantly worried about their reputations, and such worries are reflected in the working of the governance institutions to which they contribute. We think twice before deciding to squander, in a moment of madness, a string of past investments in socially acceptable behavior. The ancient Romans understood this well when they crafted the adage *semel in anno licet insanire*, or "it is acceptable to be crazy once a year," which does grant the right to insanity, but only within well-specified

limits on the number of transgressions allowed. In this respect, self-control, which in part is certainly motivated by reputational considerations, translates into predictability of human behavior, which in itself is an important prerequisite for most social intercourse.

Implicit in the private calculations people make on what course of action to take is a weighing up of the short-term gain that they could obtain from dishonest and predatory behavior, with the long-term damage that would follow from losing the trust of their acquaintances. The more important reputational considerations are, the more likely it is that the balance will tilt in favor of virtuous behavior. When this happens, the individual actors of governance are better off, and often society as a whole also benefits. We can define this as a *static* effect of reputation, in the sense that it applies to a given situation, that is, to the relevant actors as they are at a given time.

Reputational incentives, however, also have a *dynamic* effect, in that over time they tend to cause an improvement in the characteristics of the relevant actors. They do so in two distinct ways.

By rewarding good performances, they encourage actors to invest resources in improving their skills. If professionals know they will move up the career ladder if they do well in the job, they may decide to spend evenings taking classes to learn new skills. On the other hand, if they know that all they need are good connections, then they are quite likely to spend their evenings in networking activities. One effect of the presence of reputational incentives is to curb unproductive, rent-seeking activities, as economists call them.

Another, more radical way in which reputational incentives have a dynamic effect on improving the quality of the actors of governance is by supporting a process of selection. Being able to observe past performances helps us to discriminate between high- and low-quality suppliers. For example, given a choice of movies with similar plots, we would probably choose one directed by Martin Scorsese rather than one directed by some unknown director. Similarly, we would rather buy a car manufactured by Audi than one produced by a less reputable firm, even though the appearance and price tags were similar. Such reputational effects set in motion selection forces that weed out the least fit from the competition. Film directors who are unable to secure an audience for their creations have to

find new jobs, and uncompetitive firms lose importance and, eventually, go out of business.

One consequence of these static and dynamic effects is that they also lead to more predictable human interactions—if I can read the strategic considerations that inform the choices of others, I can predict them better. An alternative route to achieving better predictability of behavior is through contractualization, with courts intervening when necessary to force people to respect their obligations. Indeed, large parts of our lives are influenced by contracts of various types, as the high—and possibly increasing—litigiousness of our societies demonstrates. However, courtrooms and hefty lawyers' fees provide only a partial solution to the problem. Most relationships are not easily contractualized because they are fraught with unforeseeable contingencies and are simply too complex to be fully described in writing. Also, the transaction costs involved in a contract may be too high, and, last but not least, the courts may be unavailable, as we saw in the case of the Maghribi traders, or ineffective. The desire to maintain and improve one's reputation, and to build or safeguard trusting relationships with others, in many instances constitutes an alternative to the contractualization of behavior, and in many others supplements the capacity of contracts and formal institutions to guarantee the necessary smoothness of human interactions.[2]

An important theme of this book is that there is a complex relation between reputational incentives, their nature, and some key characteristics of governance. To clarify this point, we look at the example of open-source software production. In this production method, highly relevant reputational considerations go hand in hand with a mode of governance that is rather horizontal and nonhierarchic.

Open-source software[3] is very different from proprietary software where the code cannot be modified, or even accessed. Often, it also reflects a specific mode of production, characterized by the presence of very little (apparent) structure, and of horizontal and nonhierarchical relationships among participants, within which collaborative effort and experimentation play an important role. The tight relationship between production and experimentation is shown by one of the slogans of the open-source community: "Deliver early, deliver often." There is no clear distinction

between the planning and the production of a product and, in general, open-source software developers do not pay much attention to the codified tenets of the software-engineering discipline, which establish how software projects should be conducted and determine, among other things, that the requirements of the software to be developed should be analyzed formally and at length (Crowston et al. 2005). Open-source software development, on the other hand, often sets out to solve a problem faced by the developers themselves, and does not include a proper analysis of requirements.

This horizontal mode of production is in sharp contrast with the more traditional and vertical structure of proprietary software projects. A fortunate description of this contrast can be found in the title of the account of open-source software by Eric Raymond (2000), *The Cathedral and the Bazaar*. It is in fact rather amazing that this "bazaar" successfully deals with the degree of complexity of some of the open-source projects. A good example is the way the operating system Linux grew from the project of a college student, Linus Torwalds, to become a serious competitor to products by firms like Microsoft and Sun Microsystems.[4]

While some contributors to open-source software projects are employed by important corporations (IBM and Sun Microsystems being noteworthy cases), most of them participate in their spare time without receiving any monetary compensation. An important question, then, is what makes such highly skilled people work for free—or, to express it in the words of an economist, what is the incentive structure of the open-source mode of production? The presence of an altruistic motive seems to be contradicted by the fact that a truly altruistic person would probably prefer to spend time on finding a solution to one of the many great problems plaguing humanity, and the development of a software product (which may later be adopted by possibly greedy corporations [Lerner and Tirole 2002]) would not appear on this list. The personal satisfaction of being recognized by one's peer group certainly plays a role. However, accounts written by insightful advocates (as an example, see Raymond 2000) and academic enquirers (as in Dalle et al. 2005) agree on the pivotal role of reputation in the governance of open-source projects. Having a good reputation, besides advancing one's position within the open-source community, also signals to potential employers that one has the qualities of a good programmer (Lerner and Tirole 2002, 2005).

Previously, we noted that there is a relation between technology and the presence of reputational considerations in governance. A given technology may *support* and *allow* the functioning of a (reputation-based) governance model, just as the Internet is necessary for eBay to work, and a network of roads was essential for the functioning of the Roman empire. However, the Internet does not *determine* eBay, nor did the technology that was available to the Romans two thousand years ago *cause* the characteristics of their empire. The analysis of the relation between the available technology and governance will have to be careful and nuanced.[5]

We would now like to discuss some examples of *public* governance where reputational incentives are relevant. Obviously, reputational incentives also matter in the public sphere in general—for example, citizens in democracies, when electing representatives, care about the past record of the available candidates. Today, we are witnesses to attempts to translate forms of public governance where reputational incentives play an important role, that are Internet-based, and, incidentally, that are rather horizontal in character. One example, besides being interesting in its own right, also received the important endorsement of being cited in the program of U.S. president Barack Obama (Waters 2008). The project Peer-to-Patent (Allen et al. 2009), running from June 2007 until June 2009, was developed by the New York Law School Institute for Information Law and Policy in cooperation with the U.S. Patent and Trademark Office (USPTO), and aimed to try out a collaborative solution to fix some of the many problems that affect the patent system, in the United States and elsewhere.

Patents are legal documents that provide the assignee with limited monopoly over an invention. They do so in order to reward inventors for their successful endeavors and to encourage their efforts. The monopoly, however, is limited, because after the invention is done, it is in the best interests of society to allow people to use it freely. A patent is then a compromise between two opposing needs: to reward inventors, and to allow inventions to be adopted as widely as possible.[6] According to many observers, the patent system is in crisis, and most acutely so in the United States. There, the number of patent applications has risen conspicuously of late, giving rise to a "patent inflation" that has put strain on the USPTO and has resulted in lengthening waiting times for patent applications to

be examined. Moreover, there is a widespread perception that over the last few years, more than in the past, patents have been granted for inventions of very dubious utility, and that in general the technological and economic relevance of granted patents has decreased (Bessen and Meurer 2008).

In order to assess the fix that the Peer-to-Patent project advocates, we should remember that a patent in the United States is required by law to have three main characteristics. First, it should describe an invention that is nonobvious. Secondly, the invention should be potentially useful. Last, it should be novel. Patent examiners, who are employees of the USPTO, have to examine each application to determine if all these criteria are met. In particular, to determine the novelty of the invention, patent examiners have to consider what is called the "prior art"—that is, all relevant existing knowledge that in principle was available to the inventor. If, in the prior art, there already is a description of the invention, the novelty condition is not satisfied, and the patent should not be granted.

However, patent examiners typically have little time at their disposal to be so diligent, and the analysis of the previous art often turns out to be one of the trickiest aspects of their work. The Peer-to-Patent project addressed precisely this problem by attempting to get outside experts to participate in the search. Anyone, in fact, could participate in this experimental project by accessing its Web site (www.peertopatent.org). The Peer-to-Patent Web site allowed people to discuss the claims made by a given patent application, within an available selection, and suggest instances of relevant prior art. At the end of the discussion, all the material, together with a selection of up to ten suggestions of prior art, was transmitted to the patent examiners, who were then free to use any, or none, of it.

After the conclusion of the project, a formal assessment of its results indicated an encouraging degree of participation, and both the external participants and the examiners expressed favorable opinions (Allen et al. 2009). Outside experts contributed on a voluntary basis and did not receive any pecuniary compensation for their work. Certainly, many of them were employees of big companies like IBM, GE, Intel, and Sun Microsystems, which were also stakeholders in the project, and we may expect these people to have contributed in their working hours. However, others were independent reviewers who made their contributions in

their free time, like the independent programmers who work on open-source software projects. Reputational considerations were important, at least in the view of the designers of the system: whenever the patent office examiners chose a prior art that was suggested by one of the outside experts, this expert was awarded a symbolic "prior artist award," and his name appeared on the Peer-to-Patent Web site.

The workings of a patent system, if not the need for it, may seem rather arcane to most. National security, however, is a sector of public action that certainly feels closer to our daily lives, particularly in the aftermath of the terrorist attacks of September 11, 2001. Many observers, and also the official commission that was established in the United States to investigate the circumstances surrounding the attacks, coincided in chastising the relative lack of collaboration among the many agencies that had responsibilities in the field of national security (National Commission on Terrorist Attacks upon the United States 2004, Section 13). Sweeping reforms to correct the problem followed, the most visible of which was the creation in 2002 of the U.S. Department of Homeland Security, which aims to coordinate the work of over forty federal agencies (Moynihan and Roberts 2002).[7]

One initiative that was then taken deserves our attention. In 2005, the National Intelligence Council launched a pilot project called Intellipedia, as a way of increasing the sharing of classified information among participants in the intelligence community and putting a check on the "silo mentality" that almost naturally arises along the borders of separate organizations. This online collaborative environment runs on a classified network but is based on the same software used by Wikipedia. It is divided into three parts corresponding to increasing levels of classification, and it allows members of the intelligence community to share information across agency boundaries and to develop intelligence on particular topics collaboratively (Thompson 2006; Mazzetti 2007). Intellipedia has been a success, at least judging by the number of users and the amount of information they circulate and generate over the network. Charles Fingar, at the time deputy director of National Intelligence for Analysis, made it very clear that reputational incentives had to have an important role in Intellipedia, that it was to be "the Wikipedia on a classified network,

with one very important difference: it's not anonymous. We want people to establish a reputation. If you're really good, we want people to know you're good. If you're making contributions, we want that known. If you're an idiot, we want that known too" (Fingar 2008).[8]

In the public sphere, the Peer-to-Patent project and Intellipedia reflect an increasing interest in forms of governance that are open, that are relatively horizontal, and where reputational considerations play a role. However, given the ubiquity of reputational incentives in governance in general, it is striking to note that in *public* governance, it is possible to find cases where reputational considerations are almost completely absent. We consider one such case—the buying of a product or service—that is very often treated in diametrically opposed ways in the private and public spheres. In particular, let us consider the purchase of a service that cannot be fully contractualized. The difficulty in doing so may be due to the complexity of the service, making it close to impossible to describe it in full detail, and because the specifics and timing of its delivery depend on outside factors that are also complex and not fully predictable, so that the simple enumeration beforehand of all possible contingencies would be a daunting prospect. In fact, it does not take anything very fancy to fall into the category of "incomplete contracts," as economists call them, and it may be argued that all contracts are, to some extent, incomplete, precisely because they cannot contemplate all possible contingencies.

In real-life situations, the more incomplete a contract is, the more we observe the presence of postcontractual agreements of various types. As a simple and mundane example, consider refurbishing a family kitchen. Assume that the maker of the tiles that were specified in the contract goes bankrupt. In such cases, the owner of the house would typically talk with the workers and agree on a substitute—a case of postcontractual agreement. There are often more compelling reasons it may be impossible to fully contractualize the quality of a product or service. Its quality frequently becomes known in detail only after its delivery has taken place or sometimes after prolonged use. All products are to some extent "experience goods": We fully learn about our car after having used it for several years, and we need to watch a movie to the end in order to judge it.