



The MIT Press

Reputation Effects and the Limits of Contracting: A Study of the Indian Software Industry

Author(s): Abhijit V. Banerjee and Esther Duflo

Source: *The Quarterly Journal of Economics*, Vol. 115, No. 3 (Aug., 2000), pp. 989-1017

Published by: The MIT Press

Stable URL: <http://www.jstor.org/stable/2586901>

Accessed: 02/04/2010 19:54

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=mitpress>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The MIT Press is collaborating with JSTOR to digitize, preserve and extend access to *The Quarterly Journal of Economics*.

<http://www.jstor.org>

REPUTATION EFFECTS AND THE LIMITS OF CONTRACTING: A STUDY OF THE INDIAN SOFTWARE INDUSTRY*

ABHIJIT V. BANERJEE AND ESTHER DUFLO

This paper examines evidence of the role that reputation plays in determining contractual outcomes. We conduct an empirical analysis of the Indian customized software industry, using a data set we collected containing detailed information on 230 projects carried out by 125 software firms. We propose a model of the industry where reputation determines contractual outcomes. The evidence supports the view that reputation matters. Ex ante contracts, as well as the outcome after ex post renegotiation, vary with firms' characteristics plausibly associated with reputation. This holds after controlling for project, client, and firm characteristics.

I. INTRODUCTION

While the idea that there are severe limits to contracting and therefore reputation matters is now commonplace in economics,¹ little is known about the size of these effects. If these effects are shown to be large, they would imply a distinct view of the determinants of trade and growth, and, by implication, of policy: a view in which legal infrastructure and the availability of reputed firms could be as important as comparative advantage, human capital, and physical infrastructure.

This paper attempts to quantitatively assess the importance of reputation and, by implication, the seriousness of the limits on contracting, using data gathered from the Indian customized software industry. India is one of the largest exporters of software to the United States, but its legal infrastructure is widely seen as being quite primitive, limiting the scope for contracts. Contracts, however, play a vital role in the customized software industry since the structure of the industry invites holdups. The production process takes time—eight months on average—during which the

* We thank Daron Acemoglu, Philippe Aghion, Susan Athey, Roland Bénabou, Mathias Dewatripont, David Genesove, Jonathan Guryan, Oliver Hart, Douglas Miller, Dilip Mookherjee, Ashok Rai, Emmanuel Saez, Andrei Shleifer, Jean Tirole, Richard Zeckhauser, and two referees for helpful comments, and Paul Joskow for encouragement and support. We acknowledge financial support from the National Science Foundation, the Alfred P. Sloan Foundation, the Shultz Fund, and the John D. and Catherine MacArthur Foundation. We are especially grateful to the software professionals in India, who patiently answered our many questions, and in particular to Mr. N. R. Narayana Murthy, from Infosys Technologies Limited, whose help and insights were crucial throughout this project.

1. For an early expression of this view, see Fama [1980].

firm commits a substantial number of man-hours. The final product is typically quite specific to the client, so the client forgoes going to another firm. The client always keeps the property rights, and there is little general learning in the process of production.²

The extent to which contracts can protect, however, is limited by the fact that the desired end-product tends to be complex and difficult to describe ahead of time in a way that a court, especially an Indian court, would understand. At the outset, in most cases, neither the client nor the firm entirely understands what they are setting out to build. Disputes about the interpretation of the contract are not uncommon whenever there are cost overruns, i.e., costs in excess of what was envisaged at the onset of the project. Both sides in these disputes accuse the other of failing to adhere to the letter of the contract, with the client claiming the software firm (hereafter: the firm) has not delivered what it had contracted to deliver, while the firm argues that the client has changed his demands. Given that the description of the product in the contract is highly incomplete, courts cannot be relied on to adjudicate correctly between these conflicting claims. In other words, contracts that target the source of the overrun are probably unenforceable. This is consistent with the fact that in the Indian industry we only observe fixed-price contracts (where the firm is supposed to absorb the entire overrun) and time and materials contracts (where the client is supposed to pay for all of it).

The basic hypothesis of this paper is that a role for reputation arises because contracts are limited in this way. We model reputation for being reliable: reliable firms and clients always pay for any overrun that they have generated, as long as the other side does the same. While this is what we often heard (many people told us “it was our fault and so we paid for it”), and it is consistent with the data (see Section III), we recognize that other, related, forms of reputation are probably important in practice and that our empirical work cannot distinguish between them.

If all firms and clients were reliable, there would be no need for contracts about the sharing of overruns. However, it is usually impossible to be sure if the other side is reliable, hence the observation that all projects begin with either a fixed-price contract or a time and materials contract. While these contracts only protect one side from unreliable behavior by the other, that

2. In 90 percent of the cases we looked at, the platform and programming tools were known to the firm before starting the project.

still dominates having no contract at all.³ The choice of which side to protect depends on which side is more likely to be unreliable, i.e., on the reputations of the two sides.⁴ This generates correlations between reputation, the chosen contract, and contractual outcomes (such as the actual sharing of overrun) that we test in the empirical section.

The model suggested here is formally developed in Section II. In Sections III and IV we show that the model and its predictions are consistent with the observed facts. In particular, we find strong effects of reputation (measured in several different ways) on the choice of contracts and the contractual outcomes. Crocker and Reynolds' [1993] study of the procurement of airplane engines by the military is, to our knowledge, the only other paper that documents the effect of reputation (measured in their case by past disputes) on the choice of contracts (but not on the sharing of costs). The results of McMillan and Woodruff [1999] are also related. Using data from Vietnam and using several measures of trust including the length of the relationship and information from third parties, they show that interfirm trade credit is more likely when the delivering firm trusts its client.

II. A MODEL OF THE SOFTWARE INDUSTRY

The model we propose in this section is an attempt to capture, in as simple a way as possible, the story of the industry told in the introductory section. The model adds endogenous contractual choice to a basic model of reputation formation along the lines of Diamond [1989].

II.1. Disagreements, Overrun, and Contracts

The client (C) wants the firm (F) to build a piece of software that will be worth V to the client (we will assume risk neutrality

3. The ineffectiveness of the Indian court system also limits the extent to which these contracts protect. The costlier, more corrupt, and slower the courts, the larger will be the concessions that the plaintiff will be willing to make in an out-of-court settlement. However, this is not to say that contracts are irrelevant. The renegotiated outcomes in our sample are indeed correlated with the initial contract. Anecdotal evidence suggests that some disputes do end up in court (at least one of the firms in our sample went to court). An Indian legal expert suggested that unless the plaintiff provides prima facie evidence that the court can interpret, the case has a chance of getting thrown out, which in this case amounts to upholding the written contract.

4. For an alternative view of what determines the choice between fixed-cost and time and materials contracts based on ease of renegotiation, see Bajari and Tadelis [1999].

on both sides throughout, so that this is best thought of as a money payoff and the costs as money costs). In a world where the client can describe the product it wants perfectly and the firm also understands this description perfectly, the project should cost \bar{y} . We adopt the normalization that $\bar{y} = 0$. We assume, however, that in every real project the actual cost will be positive; i.e., there will be overrun. One reason for this is that it is very costly for the client to describe *ex ante* exactly what it wants, and even if the description is perfect, it is too costly for the firm to understand it perfectly. Thus, total overrun will be the sum of overrun caused by the firm (y_F), and overrun caused by the client (y_C). We assume that all overrun is initially paid for by the firm.

The amounts of the overrun, y_F and y_C , are a result of choices made, respectively, by the firm and the client. Specifically, we assume that firms face a choice between a high level \bar{y}_F and a low level \underline{y}_F . Likewise, the client faces a choice between \bar{y}_C and \underline{y}_C . Firms can reduce overrun by putting more effort into understanding what the client wants, by better management, and by assigning more able people to the project. Clients can reduce overrun by working more closely with the firm, and in particular, by being precise about what they want. Reducing overrun is therefore costly. The extra costs to the client and the firm from choosing a low level of overrun are, respectively, B_C and B_F . Assume that $\bar{y}_C - \underline{y}_C > B_C$ and $\bar{y}_F - \underline{y}_F > B_F$ so that it is always efficient to minimize overrun.

However, we will assume that both y_C and y_F are known only to the firm and its client. Third parties, such as the courts, only observe total overrun ($y_C + y_F$).⁵ Moreover, we restrict all contracts to being linear and additionally require that they do not involve throwing away any money. In other words, we only consider contracts where the client pays the firm an amount $P + (1 - s)(y_F + y_C)$, where P is a prespecified fixed payment and s is the share of the overrun borne by the firm ($s \in [0, 1]$)—and neither party makes any payments to third parties. Of particular interest to us will be two extreme contracts: the contract with $s = 1$ (corresponding to a fixed-price contract) and the contract with $s =$

5. In other words, the fact that, say, two people were assigned to the project for fourteen weeks longer than originally planned is verifiable, but not whether it took that long because the firm had not understood what it had to do or because the client had changed his mind about what he wanted.

0 (which corresponds roughly to a time and materials contract).⁶ It is key to our analysis that neither of these contracts obviously dominates the other. The fundamental trade-off comes from the fact that s is one number that is being used to give incentives to both parties. A high s gives good incentives to the firm but not to the client, while a low s contract does the reverse. As a result, it is typically not possible to implement the first best.⁷

II.2. Norms, Courts, and Renegotiation

When the first best cannot be achieved by contractual means, it is possible to improve on the outcome if the behavior of the firms and the clients is at least partly norm-governed. Specifically, assume that there are two types of firms and two types of clients. Of these, one type of firm and one type of client observe a norm of being reliable. Reliable firms and clients pay for the overrun that they generate, as long as the other side is reliable, but act to maximize current profits otherwise. Assume by contrast that unreliable firms and clients are myopic and always act to maximize their current profits, although it is possible that this involves simulating reliability.

When a firm and a client are matched, they do not directly observe each other's types. Rather, on the basis of what it knows of the client (firm), the firm (client) puts a probability θ_C (θ_F) on the client's (firm's) reliability. θ_C and θ_F therefore measure reputation.

Our formulation leaves out one of the key ingredients of the standard reputation model, the type that Tirole [1996] calls opportunists who choose to be reliable, not out of inner compulsions but because it pays in the long run. However, note that in a reputational equilibrium this type is behaviorally indistinguishable from the true reliable type. Therefore, we can accommodate the possibility that a fraction of the reliable types are actually reliable by choice as long as we assume that there is some probability that unreliable behavior will become public and that in equilibrium no one contracts with those who are known to have

6. The correspondence is not exact because a time and materials contract typically pays a markup on the realized costs, rather than a fixed payoff. Similar results hold for that case but the exposition is somewhat more cumbersome.

7. This basic tension is very general. Our restriction to linear contracts, while vital in the discrete case we have chosen here, can be relaxed if we are prepared to go to the model where overrun varies continuously: the impossibility of implementing the first best in that case is a consequence of the results in Holmstrom [1982]. More intuitively, the fact that both sides can generate overrun but the courts cannot observe who was responsible generates a free-rider problem. This is what drives our model.

been unreliable in the past (as in Kandori [1992] or Tirole [1996]). Introducing into our model opportunists who only sometimes behave reliably complicates the analysis but leaves the main results unchanged.

In our chosen setting, since the reliable firms and clients are going to be self-regulated, the function of the contract is to protect reliable clients against opportunism by unreliable firms and vice versa. When at least one side turns out to be unreliable, there is a dispute, and the contract becomes important. However, because going to court is costly, the contract is likely to be renegotiated. We assume that in the renegotiated outcome, the firm pays a share $s^*(s)$ of the overrun when there is a dispute. It is natural to assume that s^* is increasing in s . We also assume that s^* is bounded below by $\underline{s}^* > 0$ and above by $\bar{s}^* < 1$.⁸ A higher s^* and a lower \bar{s}^* correspond in a natural way to a less effective legal system.

The timing of the model is as follows: a firm and a client are matched and decide on a contract; the project is then carried out; people choose actions that lead to an overrun. At the end of the project, the overrun is observed, and the two parties decide whether or not to behave reliably. If either side is unreliable, there is a dispute, and the overrun is shared in the proportion $s^*(s)/(1 - s^*(s))$.

II.3. Reputation Formation

There are several mechanisms by which reputation can evolve. First, in those cases where the firm and the client have contracted at least once before, the presumption is that both had behaved reliably so that they both now have a better reputation with the other.

Second, the age of the firm should be a source of reputation. We have already noted that in order for there to be a reputational equilibrium, there must be some probability that unreliable behavior becomes publicly known in the future. Once that happens, no one will want to contract with the guilty firm, and it will probably end up going out of business. This selection process ensures that older firms (and clients) will typically be more reliable. This is reinforced by the fact that information is revealed

8. Given that we have assumed risk neutrality, the natural interpretation for s^* is that it is the *expected* share of the overrun borne by the firm in the event of a dispute. The actual share will presumably vary according to the exact circumstances of the negotiation process, and, on occasion, may turn out to be 1 or 0.

over time, and therefore a firm that has been in the industry a long time and does not have any black marks against it, is more likely to be reliable.

Finally, firms may be able to establish a good reputation by demonstrating that they follow processes that, in principle, should reduce overrun. Process certification by outside agencies, such as ISO 9001 certification, is therefore another potential source of reputation.

II.4. Analysis of the Model

Firm Behavior. Reliable firms have only one choice, to choose between \bar{y}_F and \underline{y}_F . Unreliable firms have the same choice, but in addition they can choose whether or not to mimic reliable firms in the current period. Clients face an exactly symmetrical choice. The following proposition gives sufficient conditions for there to be an equilibrium where reliable firms choose low overrun, and unreliable firms behave unreliably and choose high overrun.

CLAIM 1. The unique equilibrium behavior of reliable firms and clients is to choose $y_F = \underline{y}_F$ and $y_C = \underline{y}_C$, as long as $(\theta_C + s^*(1 - \theta_C))(\bar{y}_F - \underline{y}_F) > B_F$ and $(\theta_F + (1 - s^*)(1 - \theta_F))(\bar{y}_C - \underline{y}_C) > B_C$. The unique equilibrium behavior of unreliable firms and clients is to choose $y_F = \bar{y}_F$ and $y_C = \bar{y}_C$, subsequently going into dispute, as long as $\bar{s}^*(\bar{y}_F - \underline{y}_F) < B_F$, $(1 - \bar{s}^*) \cdot (\bar{y}_C - \underline{y}_C) < B_C$, and \underline{y}_F and \underline{y}_C are both sufficiently close to 0.

For these conditions to be jointly satisfied, θ_F and θ_C have to be high enough, and s^* must be sufficiently bounded away from both 0 and 1. If the first condition fails, i.e., if either θ_F or θ_C is not large enough, it does not pay to be reliable as the other side is very likely to be unreliable. When \bar{s}^* is very large, it implies that the contract is very effective in punishing unreliable behavior by firms and therefore even unreliable firms will behave reliably. The same holds when \underline{s}^* is very small. In the rest of the formal analysis, we will assume that these conditions hold.

The Optimal Contract and the Sharing of Overrun. The firm proposing the contract in our model knows its own type, and as a result the proposed contract can be used as a signaling device. A firm that plans to be unreliable will prefer a contract where it pays very little of the overrun. Therefore, by choosing to absorb most of the overrun, a firm may be able to signal that it is reliable. Given that we are in a signaling environment, we expect that there will be many equilibria. However, all such equilibria will involve

pooling because in any separating equilibrium no one will contract with unreliable firms.

Among the set of pooling equilibria, we focus on contractual outcomes where the total joint surplus of a firm and a client who are both reliable is maximized. This is always a Bayesian-Nash equilibrium (sustained by the belief that only opportunists deviate). The fact that it is also Pareto optimal from the point of view of the reliable types makes it an obvious focal outcome.

This expression for joint surplus is

$$W(s, \theta_C, \theta_F) = V - \theta_F \underline{y}_C - (1 - \theta_F)(1 - s^*(s))(\underline{y}_C + \overline{y}_F) \\ - \theta_C \underline{y}_F - (1 - \theta_C)s^*(s)(\underline{y}_F + \overline{y}_C).$$

The third term in this expression gives the total surplus that is lost because in a pooling equilibrium a reliable client must allow for the possibility that the firm is unreliable, while the fifth term is the surplus that is lost because the firm must allow for the possibility that the client is unreliable.

From the linearity of this expression in s^* , it immediately follows that the optimal contract is always either a fixed-price or a time and materials contract (s is either 0 or 1). For any fixed level of client reputation, the more reputed the firm, the more likely it is that it is a time and materials contract.⁹

The expected share of the overrun paid by the average firm with reputation θ_F that works for a client of reputation θ_C , is $\theta_F \theta_C [\underline{y}_F / (\underline{y}_F + \underline{y}_C)] + (1 - \theta_F \theta_C) s^*(s(\theta_F, \theta_C))$. Increasing θ_F has two effects on the expected share. First there is an effect through the choice of the contract. Second, keeping the choice of the contract fixed, there is a mechanical effect deriving from the fact that more reputed firms are more likely to behave reliably and keep their overrun under control. As a result, they end up paying a smaller share of the overrun as long as $s^* > [\underline{y}_F / (\underline{y}_F + \underline{y}_C)]$, which is likely for firms with low reputations since they face fixed-price contracts. However, note that for very reputed firms, as it is likely that $s^* < [\underline{y}_F / (\underline{y}_F + \underline{y}_C)]$, the net effect of an increase in reputation can be to increase their share of the overrun.

9. The reader may feel that the result that the contracts are always at one or the other extreme is driven by the assumption of risk neutrality. This is partly true in the sense that if the two parties are sufficiently risk-averse a less extreme contract will be chosen. However, note also that s^* is always strictly between zero and one; the effective contract is never very extreme. Therefore, it may well be the case that the initial choice of an extreme contract is consistent with optimal risk sharing.

The mean overrun generated by a firm of reputation θ_F and a client of reputation θ_C is $\theta_F \underline{y}_F + (1 - \theta_F) \overline{y}_F + \theta_C \underline{y}_C + (1 - \theta_C) \overline{y}_C$. This is clearly decreasing in θ_F . Finally, the share of total overrun caused by the firm is $\theta_F \theta_C [\underline{y}_F / (\underline{y}_F + \underline{y}_C)] + \theta_F (1 - \theta_C) [\underline{y}_F / (\underline{y}_F + \overline{y}_C)] + \theta_C (1 - \theta_F) [\overline{y}_F / (\overline{y}_F + \underline{y}_C)] + (1 - \theta_F)(1 - \theta_C) [\overline{y}_F / (\overline{y}_F + \overline{y}_C)]$. Simple calculations show that this expression is decreasing in θ_F .

So far, our results describe the effect of increasing θ_F , keeping θ_C fixed. However, our data do not contain measures of the client's reputation. The same results, however, hold without conditioning for the client's reputation under two scenarios. The first is one where most clients have similar reputations. This is not implausible because most clients are well-established companies. Alternatively, we could assume that there is random matching. This may be plausible in at least some cases because clients have quite specific demands. Finding an Indian firm with the right specialization that can work on the required schedule may not always be easy, forcing the client to choose the one firm that, by chance, happens to be available at that time.¹⁰ Under either of these scenarios we have

CLAIM 2. Firms with a better reputation are more likely to have time and materials contracts. In most cases they also pay for a smaller share of the overrun. There is less firm-generated overrun and less total overrun in contracts involving more reputable firms.

II.5. Discussion

Our model yields a number of clear-cut results, but only under rather stringent conditions. One key assumption was that all firms do similar projects. If firms in our model could choose their projects, low reputation firms would tend to gravitate

10. The random matching assumption is clearly indefensible in the case of repeat contracts. The fact that the client and the firm have agreed to work together again signals mutual respect. Compared with the initial match between them, both θ_C and θ_F should therefore be higher. The net effect is therefore potentially ambiguous. However, given that there are many more start-ups among firms than among clients, it is plausible the firm's reputation will improve by more than the client's reputation. In this case, repeat contracts will be more likely to be time and materials contracts, and the firm's share of the overrun in a repeated contract should be lower. For the case of first time matches, in a previous version of this paper we formalized an argument showing that, *ceteris paribus*, firms with low reputations will prefer to work with highly reputed clients. This is because they are particularly vulnerable to opportunistic behavior by the client. This kind of nonrandom matching reinforces our result that low reputation firms get fixed-price contracts and end up with a higher share of the overrun. However, it also weakens the relation between total overrun and the firm's reputation.

toward projects with low \bar{y}_C even at the cost of lower returns. This is because they have the most to lose if the client misbehaves.¹¹ This suggests that low reputation firms will tend to choose projects which are simple and well understood where the client does not have to do very much work to make clear what he wants. Firms that switch to a low \bar{y}_C project should be *more willing* to accept a fixed-cost contract. Moreover, the switch to low \bar{y}_C will tend to *raise* the fraction of the overrun generated by the firm, and therefore the firm will typically pay a *larger* part of the overrun. All of these effects go in the direction of our previous results. Total overrun may, however, no longer fall with reputation since low reputation firms choose low overrun projects.

Our model is also limited by the fact that, by assumption, contracts have no direct incentive effects. If the set of options faced by the firms and clients were less discrete, one would imagine that contracts would also have incentive effects. A high value of s would then have the additional effect of discouraging firms from generating very high levels of overrun. An incentive effect of this kind reinforces the correlation between low reputation firms and fixed-price contracts. However, it also means that low reputation firms do not necessarily generate more overrun than their more reputable counterparts. This, in turn, weakens the correlation between reputation and total overrun. It should also weaken the relation between reputation and the share of overrun that is *paid* by the firm. However, one would still expect this last correlation to be negative, as low reputation firms have fixed-price contracts and those who have fixed-price contracts pay for more of the overrun.

A generalized version of our basic model that allows for both project choice and the incentive effect of contracts would have the following predictions: 1) low reputation firms will choose projects that are simple and easy to define; 2) low reputation firms are more likely to face fixed-price contracts, and on average pay for a higher share of the overrun; and, 3) neither total overrun nor the overrun generated by the firm is necessarily correlated with its reputation.

The combination of the last two predictions is what distinguishes a model, like ours, that emphasizes incentive problems and reputation effects: low reputation firms in our model will face harsher terms *even if their actual performance (measured by*

11. This is formally shown in a previous version of the paper.

overrun) is indistinguishable from that of high reputation firms.¹² This is because what matters in this type of model is not what firms actually do, but rather, what they *could* do.

By contrast, theories where firms are paid based on performance, but not because of contracting problems—standard marginal product theories, for example—would predict a correlation between reputation and the share of overrun paid by the firm, *only if there was also a correlation between reputation and overrun caused by the firm* (which, in turn, suggests a correlation between reputation and total overrun).¹³ Theories that give a central role to bargaining between the firm and the client, for example, can generate a correlation between reputation and the share paid by the firm if bargaining power is correlated with reputation, but would tend to predict the wrong correlation between contract type and reputation. Intuitively, a firm that has a lot of bargaining power will value having a fixed-cost contract less than a firm that is in weak bargaining position. Finally, risk-sharing-based theories would tend to predict the wrong correlation between reputation and the share of the overrun paid by the firm. Less reputed firms are typically smaller and have fewer assets, and should bear less of the risk (overrun).

III. DESCRIPTIVE EVIDENCE

In this section and the following, we document that both the central assumptions and main implications of our model are consistent with what we observe in the industry. The data set used here was collected by interviews with the CEOs of 125 software firms in the three major centers of software production in India (Bangalore, Pune, and Hyderabad) in the winter of 1997–1998. We collected general information about each firm, and specific data about the last two projects completed by the firm. More information about the data collection process is presented in the Appendix.

12. There are of course other models within the general category of incentive/reputation-based models. For example, a model where the reputation is for quality of work rather than for reliability may have rather similar predictions.

13. One example of this class of theories would be one that says low reputation firms generate more overrun and therefore pay for them. An alternative would be the argument that low reputation firms set low estimates either deliberately or out of inexperience and end up paying more as a result.

III.1. Projects, Estimates, and Overruns

We focus in this paper on contracts for customized software projects. A *project* is a well-defined notion in the industry. It corresponds to a set of functions that the software must execute. These functions are (often loosely) described in the Request for Proposal (RFP) sent by the client when they first look for a firm. Most projects in the sample (75 percent) were for overseas clients. Common projects include the development of a customized application specific to the client, customization of a package, porting or reengineering of an existing solution, and testing of existing software. Seven percent of the projects in the sample were Y2K projects, and another six percent were other simple projects: data manipulation, web pages, and CAD projects.

In response to the client's RFP, several firms submit proposals. Then the client chooses one of these firms. Each proposal contains contractual terms and an *estimate*. Our understanding, based on descriptions in standard software management textbooks (e.g., Pressman [1997]) and conversations with people in the industry, is that the estimate is supposed to be the firm's best guess as to how many hours will be needed to complete the project, *assuming that the firm's current understanding of the project is correct and that the firm adheres to its own productivity norms*.¹⁴ The firm and the client understand that both these assumptions may be false. The project may take more effort than estimated either because the firm took the wrong direction on the project, or the client changed his demands, or because productivity in the firm was lower than it would normally be. This uncertainty is taken into account in the terms of the contract, but the estimate is still a useful benchmark.

We asked the firm what the initial estimate was for each project. Using a timeline representing the project (which we drew on a piece of paper as we conducted the interview), we then asked them to identify all the stages in the project when the estimate was revised and by how much (more details about the process are given in the Data Appendix). *Overrun* is defined to be the difference between the actual effort involved in completing the project and the estimate. Several different questions were asked about the level of the overrun, and we made sure that the

14. The software industry measures everything in the unit of man-hours. This reflects the fact that labor is the key input and all other inputs (computers, space, and phone time) are more or less linearly related to the number of employees.

respondent gave consistent answers to all the questions, probing him if there was an inconsistency. We are confident that the figure we obtained for the overrun is a reliable measure of the difference between actual cost and the estimate. It is, of course, possible that the estimate itself was manipulated. A high estimate would favor an unreliable firm and a low estimate an unreliable client, and moreover, the estimate could be used for signaling at the precontracting stage. However, the fact that the client gets several competing estimates and the fact that there are standard methods and computer packages (used in 69 percent of cases) for computing the estimate make it much less likely that a firm would get away with substantially manipulating this number (the evidence in Section IV on the level of overrun also gives us some confidence in the way we use these numbers).

Descriptive statistics are presented in Table I. Overruns are frequent (75 percent of the cases), large (23 percent of the estimate on average), and variable (the standard deviation is 30 percent). We asked the firms to report the reason for the overrun by asking them what exactly happened in as neutral a way as possible. Firms attribute, on average, 45 percent of the overrun to changes explicitly required by the client, a further 12.8 percent to delay in the client's response at critical points, 19.5 percent to their own misunderstanding of the specifications of the project, and 9.1 percent to internal problems in the firm (for example, the loss of the project's manager). According to this, the client is responsible for 58 percent of the overrun on average (changes and delays in providing critical inputs), and the firm for 29 percent (misunderstandings and internal problems).

III.2. Clients and Contracts

Clients fall into two categories, external and internal clients. An internal client can be the firm's mother company (an American firm or an Indian group), or a company with whom the firm has a very long-term and open-ended relationship. For example, some software firms have established "Off-shore Software Development Centers" (OSDC), for their foreign clients whereby the firm dedicates a part of its office and manpower to that client. Even in the case of internal clients, the work is defined in terms of projects and is governed by a contract. However, the scope for unreliable behavior in this type of relationship is presumably much more limited, since both parties share the control rights. We will therefore view this kind of relationship as a substitute for

TABLE I
DESCRIPTIVE STATISTICS

	External projects			Internal projects
			Standard	Mean
	Mean	Median	deviation	
	(1)	(2)	(3)	(4)
Number of observations	167			57
PANEL A: PROJECT CHARACTERISTICS				
Project size (man-months)	75	27	141	207
Application area is familiar to the firm	0.68	1	0.47	0.69
Platform is familiar to the firm	0.90	1	0.30	0.93
Programming tools are familiar to the firm	0.87	1	0.34	0.84
Y2K,CAD, Web pages	0.14	0	0.35	0.07
PANEL B: REPUTATION				
Founded 1993 or later	0.43	0	0.50	0.57
ISO certification	0.19	0	0.39	0.14
Worked with the client in the past	0.41	0	0.49	0.83
PANEL C: SOFTWARE FIRMS AND CLIENTS				
Number of employees in the software firm	147	60	279	153
Client is a Fortune 500 company or an Indian group	0.55	1	0.50	0.28
Client is an Indian firm	0.22	0	0.42	0.01
PANEL D: OVERRUNS				
Project cost > estimate (overrun)	0.75	1	0.43	0.70
Project cost < estimate (underrun)	0.02	0	0.15	0.10
Total overrun (percent of project cost)	23	15	30	27
Percent of overrun due to the client:	58	50	43	58
—Changes	45	50	43	47
—Client's delay	13	0	26	11
Percent of overrun due to the firm:	29	0	43	11
—Ambiguity in specifications	20	0	34	20
—Internal difficulties	9	0	25	−9
PANEL E: CONTRACTS (proportion)				
Fixed-cost contracts	0.58	1	0.49	0.23
Mixed contracts	0.26	0	0.44	0.07
Time and material contracts	0.15	0	0.36	0.67
PANEL F: SHARE OF OVERRUN PAID BY THE FIRM (percent)				
All contracts	57	77	45	20
Fixed-cost contracts	66	100	44	45
Time and material contracts	28	0	39	8
Fixed-cost contract, firm responsible	86	100	33	no obs.
Fixed-cost contracts, client responsible	61	100	46	45

reputation. Except for a few very large internal projects, internal and external projects are similar in nature.¹⁵ Firms with internal projects are on average smaller and more recently established than other firms. Most firms who work for internal clients also do some work for external clients. We will use this feature to contrast the features of internal and external contracts.

Consistent with the predictions of our model, we observe only three types of contracts. Time and material contracts are the least frequent (15 percent). Fixed-cost contracts are the most frequent (58 percent). Mixed contracts divide the project into a specification phase (which is the first and most open-ended part of the project, where the client's needs are analyzed to design the broad architecture of the piece of software), and an execution phase (design, coding, and testing). A separate contract (time and materials or fixed-price) is signed for each part. The most frequent configuration is a time and materials contract for the specifications, and a fixed-price contract for the execution.

As assumed in our model, contracts are frequently renegotiated, but do indeed matter for the final outcome. This is shown in Panel F of Table I. Firms bear on average 57 percent of the overrun. When they have a fixed-price contract, they bear on average 66 percent of the overrun (the median is 100 percent). When they have a time and materials contract, they bear on average 26 percent of it (the median is 0 percent). The data are consistent with our premise that some firms tend to pay for their mistakes no matter what the contract was. Even with fixed-price contracts, firms pay more (86 percent) of the overrun when they are fully responsible for it than when the client is fully responsible for it (61 percent). The ordering is the same for time and materials contracts.

III.3. Reputation

As discussed in the theory section, we consider three measures of reputation: the age of the firm, the fact that the firm has previously worked with the client (41 percent of the contracts), and ISO certification (19 percent of the contracts). In addition, we contrast internal and external contracts among firms that do some internal contracts.

15. The difference in mean effort (207 man-months for internal projects, 75 man-months for external projects) is driven by a few very large internal contracts. The median efforts are similar in both cases (25 and 27 man-months, respectively).




FIGURE I
Proportion of Fixed Cost Contracts

Preliminary evidence suggests that the principal empirical implications of our model are not rejected by the data. In Figure I we plot the proportion of fixed-cost contracts as a function of the age of the firm. In Figure II we plot the average share of overrun that is paid by the firm. Both the proportion of fixed cost contracts and the firm's share of the overrun strongly decrease with the age of the firm, in particular in the first five years (which is consistent with the prediction that the marginal effect of reputation declines as reputation increases). Table II shows the proportion of fixed-cost contracts, the firm's share of the overrun, and other project characteristics for each type of firm, and the difference between low and high reputation firms. Standard errors of the means and the difference in means are presented in parentheses. In columns (1) to (3) we show the contrast between young firms (created in 1994 or after) and old firms (created in 1993 or before). Young firms are significantly more likely to have fixed-price contracts (the probability is 26 percent higher). They also, on average, bear substantially more of the overrun (19 percent more). These differences are significant at the 95 percent level of confidence. ISO-certified firms (columns (4) to (6)), by contrast, are no less likely to get fixed-price contracts nor to pay less overrun than other firms. Firms engaged in a repeated relationship with their




FIGURE II
Share of Overrun Paid for by the Firm

client are 9 percent less likely than other firms to have fixed-price contracts (columns (7) to (9)), although the difference is not significant, and they pay significantly less of the overrun (20 percent less). Finally, firms that have internal contracts pay for more of the overrun when they deal with external clients than when they deal with internal clients. Almost half of their external contracts are fixed-price contracts, whereas only 23 percent of the internal contracts are fixed-price contracts. Firms pay a much smaller share of the overrun (20 percent instead of 47 percent) in internal contracts than in external contracts.

This set of contrasts among the contracts and contractual outcomes of the different types of firms is not explained by obvious differences in firms' behavior or abilities, at least to the extent that these would be reflected in different overruns or different probabilities of delay.¹⁶ In Table II we present the total overrun of the project, the total overrun due to the firm, and the proportion of projects that were delivered with a delay. None of the differences between high and low reputation firms are significant. If any-

16. Delay is clearly an important index of the firm's performance. For simple projects where quality can be controlled relatively easily, clients are primarily concerned about delays (in 54 percent of the cases the deadline was extremely important or critical to the client). Delay is useful for us because it provides a check on the firm's performance that does not rely on the *estimate* which, as we have noted, might be subject to some manipulation.

TABLE II
DESCRIPTIVE STATISTICS, BY REPUTATION

	External contracts												Firms that have internal projects		
	Firm foundation date			ISO 9001 certification			Relationship with the client			Type of project					
	Difference			No	Yes	Difference	New client	Repeated	Difference	External	Internal	Difference			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)			
Proportion of fixed-cost contracts	0.73 (0.053)	0.47 (0.052)	0.26* (0.089)	0.59 (0.042)	0.53 (0.090)	0.061 (0.12)	0.62 (0.049)	0.53 (0.061)	0.087 (0.092)	0.45 (0.061)	0.23 (0.051)	0.22* (0.080)			
Share of overrun paid for by the firm	68.6 (5.16)	49.3 (6.25)	19.4* (8.10)	56.9 (4.59)	57.8 (8.84)	-0.98 (10.4)	64.2 (4.94)	44.2 (6.73)	20* (8.27)	46.7 (6.36)	20.2 (5.82)	26.6* (8.67)			
Mean of overrun (percent of initial estimate)	19.9 (2.84)	24.8 (2.65)	4.9 (5.16)	20.1 (1.89)	33.75 (6.10)	13.65 (8.64)	23.3 (2.59)	22.02 (2.86)	1.28 (4.84)	15.9 (2.55)	4.4 -4.21	-11.5 (5.92)			
Overrun caused by the firm	7.52 (2.34)	7.03 (1.28)	0.49 (3.18)	6.97 (1.32)	10.9 (3.24)	-3.90 (4.24)	8.19 (1.81)	7.04 (1.49)	1.15 (3.39)	4.61 (1.64)	7.46 (2.86)	-2.85 (4.61)			
Project late	0.15 (0.043)	0.23 (0.043)	-0.074 (0.062)	0.18 (0.03)	0.28 (0.081)	-0.10 (0.078)	0.22 (0.042)	0.17 (0.045)	0.06 (0.063)	0.18 (0.047)	0.20 (0.049)	-0.02 (0.068)			

Standard errors are in parentheses. The standard errors of the differences in means are corrected for clustering at the firm level. An asterisk indicates that the difference is statistically different from 0 at the 95 percent level of confidence.

thing, total overruns and the probability of delay are somewhat smaller for young firms than for old firms, for non-ISO firms than for ISO firms, and for external contracts than for internal contracts.

We interpret these results as showing that reputation (except for ISO certification) does influence the way the overruns are shared between the client and the firm. To confirm this interpretation, we now turn to regressions estimates, which allow us to control for project, firm, and client characteristics which may have independent effects on both the contract and the outcome and may be correlated with our measures of reputation.

IV. EVIDENCE FROM REGRESSIONS

IV.1. Choice of Contract

The theory suggests that there is a threshold, such that when the reputation of a firm crosses it, the optimal contract switches from time and materials to fixed-price contracts. In practice, this threshold should vary by firm, client, and project characteristics, and perhaps also in response to other environmental features. Therefore, the natural model to estimate is a discrete choice model. In our setting there are three kinds of contracts that can be ordered. Fixed-cost contracts give the most incentives to the firms, time and materials contracts give the most incentives to the clients, and mixed contracts are in between. We therefore assign a “1” to a fixed-price contract, a “2” to a mixed contract, and a “3” to a time and materials contract, and estimate an ordered probit model for the choice of contract.

Empirically, potential determinants of contract choice include not only reputation characteristics (a dummy for whether the firm was created after 1993, a dummy for whether the firm has worked with the client in the past, and a dummy for whether the firm is ISO certified) but also other variables potentially correlated with reputation, which may affect the choice of contract.

First, if firms and clients were risk-averse (in contrast to what we assume in the theory section), contracts would reflect risk-sharing motives. The contract may then depend on the client's and firm's size (as a proxy for their risk aversion). The inherent uncertainty of the project (which should be correlated with estimated project size, the complexity of the project, the degree to which the firm and the client are familiar with the

project, etc.) should also influence the choice of contract. Since the client is generally much bigger than the firm, we might expect that characteristics which make the project riskier will also make it more likely that the chosen contract is a time and materials or a mixed contract.

Second, Crocker and Reynolds [1993] argue that fixed-price contracts are more likely when it is easier to draft an exhaustive list of requirements (a complete contract). Their empirical analysis confirms these predictions. Contracts tend to be fixed-price if the engine is well-known or the production cycle is short. In our context, this would also suggest that the contract is more likely to be a time and materials contract if the project is large and if the area of application or the platform is not familiar to the firm. We would also expect that extremely well-defined projects (such as CAD or Y2K projects) are more likely to have fixed-price contracts.

To summarize, denote by C_{ic} the type of contract obtained by firm i in contract c , by R_{ic} the vector of reputation variables, by X_{ic} the project characteristics, by Z_{ic} the firm's characteristics, and by M_{ic} the client's characteristics. We estimate the following model by ordered probit:¹⁷

$$(1) \quad C_{ic} = \alpha R_{ic} + \beta X_{ic} + \gamma Z_{ic} + \delta M_{ic} + v_i + w_{ic}.$$

Results are presented in Table III, columns (1) and (2). A negative coefficient on a variable indicates that the contract is moving away from time and materials toward fixed-price. Controlling for client and firm size, the client's country, and project characteristics, young firms are significantly more likely to have fixed-price contracts. Firms that work with a repeated client are more likely to have a time and materials contract, but the difference is not significant. ISO-certified firms are *more* likely to have a fixed-price contract, although this coefficient is not significant. These results confirm the descriptive statistics presented in Table II.¹⁸ In column (2) we focus on firms that have internal clients. When these firms work with an internal client, they are significantly more likely to have a time and materials contract than when they work with an external client.

Interestingly, project characteristics do not seem to influence

17. The standard errors are corrected for firm-level grouping.

18. We have also interacted the fact of working for a repeat client with the age of the firm. Working with a repeat client makes it more likely that they have a time and materials contract, and this is not true for old firms. This suggests that old firms may have a strong enough reputation in the market to make specific bilateral relationships relatively unimportant.

TABLE III
REGRESSION RESULTS: CHOICE OF CONTRACTS AND SHARE OF OVERRUN
PAID BY THE FIRM

	Share of overrun paid by the firm							
	Choice of contract Ordered probit		Unconditional				Conditional	
			Random effect		Fixed effect		Random effect	Fixed effect
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Reputation								
Young firm	-0.69*		15*				9.0	
	(0.25)		(8.5)				(8.6)	
Repeated contract	0.22		-17*		-20		-15*	-19
	(0.24)		(8.8)		(16)		(8.7)	(17)
ISO-certified firm	-0.27		17				16	
	(0.32)		(13)				(13)	
Internal project		0.87*		-25*		-64*		
		(0.31)		(11)		(26)		
Contract								
Fixed-cost contract							13	12
							(9.4)	(25)
Time and material contract							-12	15
							(13)	(34)
Firm and client characteristics								
Number of employees (/100)	0.44	-0.44	-4.1	-2.4			-4.8	
	(0.98)	(1.1)	(7.7)	(5.0)			(4.5)	
Client is big	0.15	-0.18	-17	-13	-16	-40	16	-16
	(0.22)	(0.30)	(8.5)	(10)	(16)	(32)	(8.3)	(16)
Client is Indian	-0.43*	-0.76	13	23	-46*	14	9.3	-45*
	(0.27)	(0.63)	(9.5)	(18)	(20)	(41)	(9.5)	(21)
Project characteristics								
Estimated project size (man-months/10)	0.01	0.00	-0.12	0.35	-2.4	2.1	-0.13	-2.1
	(0.017)	(0.011)	(7.7)	(.38)	(1.9)	(2.6)	(0.76)	(2.0)
Area is familiar	0.08	0.09	-9.0	-4.0	18.9	-1.5	-6.3	16
	(0.25)	(.027)	(9.4)	(12)	(21)	(28)	(9.2)	(22)
Platform is familiar	0.60	0.20	26	1.9	48	61	29	49
	(0.48)	(0.59)	(20)	(12)	(35)	(74)	(19)	(36)
Y2K, data manipulation, etc.	-0.13	0.15	19	-4.8	-47		15	-41
	(0.32)	(0.35)	(17)	(25)	(36)		(16)	(39)

Standard errors (corrected for clustering at the firm level in columns (1) and (2) are in parentheses.

An asterisk indicates that the coefficient is significant at the 90 percent level of confidence.

All regressions include the following additional variables: number of employees squared, project size squared, and an indicator for whether tools are familiar to the firm. Sample in columns (1), (3), (5), (7), and (8) include all external projects (167 projects in 92 firms). Sample in columns (2), (4), and (6) include all projects of firms that do some internal work (88 projects in 58 firms).

the choice of contract. Firm and client size have no impact either. There is no strong evidence that risk-sharing or the difficulty in drafting the contract (as in Crocker and Reynolds [1993]) plays an important role in determining the contract.¹⁹

IV.2. Share of the Overrun Paid by the Firm

A quantitative measure of the importance of reputation is obtained by estimating the effect of reputation on the share of the overrun that is actually paid by the firm. Our model predicts that reputation influences the share of the overrun paid by the firm both without conditioning on the contract, and potentially also after conditioning on the contract. We therefore run both regressions. Clearly, when we do not condition for the contract, we should control for all the characteristics that affect contractual choice. Once we condition on the contract, realized overrun and share paid by the firm are jointly determined. We focus on the estimation of a reduced form, as we do not have instruments for unpredicted overrun that would be necessary to estimate a structural model.

We therefore estimate the model using a random-effects specification model:

$$(2) \quad S_{ic} = \alpha R_{ic} + \beta X_{ic} + \gamma Z_{ic} + \delta M_{ic} + \lambda C_{ic} + v_i + w_{ic},$$

and its equivalent excluding C_{ic} .

The dependent variable of interest is measured only when there is an overrun.²⁰ We have also estimated a Heckman [1979] sample selection model, excluding the vector of project characteristics X_{ic} in equation (2). This exclusion is justified, as we have seen that they do not predict the type of contract. We do not

19. This argument assumes that the project characteristics that we include in the regression are correlated with project risk and the difficulty of drafting a complete contract. We present later evidence that overrun is correlated with these characteristics. We also found (but to save space do not report) that the effort put into drafting the proposal (a measure of the difficulty to draft the contract) is correlated with these characteristics.

20. In cases where there is a 0 or negative overrun, we could impute the sharing of the overrun predicted by the contract since in such cases the contract is never revised (in fact, 0 overrun often means a very small positive or negative overrun). The results are similar if we use this alternative variable. However, as we are interested in how large overruns are shared, it would be misleading to include the cases where there was very little overrun, because what happened in these cases does not necessarily tell us what would have happened if the overrun had been large.

report the results, which are very similar to the random-effects specification.²¹

Results are presented in columns (3), (4), and (7) of Table III. Young firms pay 15 percent more of the overrun than old firms. Firms that work with a repeated client pay 17 percent less of the overrun than other firms. Both of these numbers are significant at the 90 percent level. ISO firms tend to pay more of the overrun than non-ISO firms. Firms pay 20 percent less of the overrun when they do internal projects. Bigger firms, firms that work with a foreign client, and firms that work for a big client (a Fortune 500 company, or an Indian group) pay less of the overrun. Conditional on the type of contract (column (7)), young firms pay more of the overrun (9 percent), but the difference is not significant. Firms that work with the client for the second time pay more of the overrun. The difference is as large as in the unconditional specification and remains significant. The fact that the age effect is small once we condition on the type of contract is not surprising. Once we control for contract type, the effect of age in our basic model comes from the fact that young firms generate more overrun (this is what we call the mechanical effect). In the more general model suggested in subsection II.5, the correlation between age and firm-generated overrun disappears. As we will see, the data confirm the absence of such a correlation. The strong effect of repeat contracts, even after we control for the type of contract, is probably due to an effect that our model does not capture: in the real world, clients that have had a good experience with a firm are more like to behave reliably since they want to retain the firm. In other words, reliability may be endogenous.

We use the fact that we have two observations for most firms to estimate a fixed-effects model. Fixed-effects wipe out both the effect of the age of the firm and the effect of whether it is an ISO firm, thus allowing us to look more closely at the effect of being in a repeated contract. Specifically, it allows us to check whether this effect is spuriously generated by the fact that better firms pay less of the overrun and are more likely to work again with the same client. It also gives us better estimates of the effect of internal contracts and other characteristics. Columns (5) and (8) present results for repeated contracts. Not surprisingly, the fixed-effects

21. We have also experimented with a double-sided tobit model, which takes into account the fact that the share of the overrun paid by the firm is bounded between 0 and 100 percent. The results are qualitatively similar; hence, we focus on the linear model.

specification is much noisier than the random-effects specification. However, the point estimate of the effect of working with a repeated client is practically unaffected (-20 percent instead of -17 percent).²² The effect of working with an internal client on the share of the overrun, shown in column (6), is *bigger* in the fixed-effects specification than in the random-effects specification.

IV.3. Overruns and Project Characteristics

An alternative interpretation of our results could be that some characteristic of the firm, unobserved by us, but known to the client (say, the quality of the firm) is correlated both with the choice of contract (e.g., good firms get time and materials contracts) and with our measures of reputation. However, as we have discussed in the theory section, this class of explanations would imply a correlation between the performance of the firm and its reputation. In contrast, our model, in its more general form, suggests that there can be a correlation between the contractual outcomes and the reputation of the firm, even in the absence of any link between the performance of the firm and its reputation.

Therefore, we estimate equations similar to equation (2) but with a measure of firm performance—such as total overrun (as a percentage of the initial estimate), overrun generated by the firm, and delay as the dependent variables. To save space, we report only random-effects estimates. Fixed-effects point estimates are similar.

Results are presented in columns (1) to (5) of Table IV. Neither client overrun nor overrun due to the firm is correlated with the reputation of the firm. They are not correlated with client and firm characteristics either, but are correlated with some of the project characteristics. Shorter projects have smaller overruns in percentage terms, as do projects where the platform is familiar to the firm and simple projects such as Y2K or CAD projects. The coefficients on reputation as well as those corresponding to firm, client, and project characteristics are similar irrespective of whether or not we control for the contract. Interestingly, in the cross section, total overrun (but not the overrun generated by the firm) seems to be bigger in fixed-cost contracts. If, perhaps as suggested by Crocker and Reynolds [1993], time and materials contracts were associated with less effort being put into the

22. The only coefficients affected are those of working with an Indian client and of doing Y2K projects, which become negative.

TABLE IV
REGRESSION RESULTS: OVERRUNS AND PROJECT CHARACTERISTICS

	Total overrun			Overrun due to the firm			Project size		Subjective complexity
	Unconditional		Conditional	Unconditional		Conditional	Median regression	Probit	OLS
	Random effect	Random effect	Random effect	Random effect	Random effect	Random effect			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
Reputation									
Young firm	-0.48 (5.0)	-3.8 (5.0)	2.5 (3.4)	1.5 (3.5)	-0.28 (0.22)	-6.4 (5.4)	0.48* (0.29)	-0.42* (0.21)	
Repeated contract	1.8 (4.9)	1.5 (4.8)	-0.92 (3.5)	-1.2 (3.5)	-0.2 (0.24)	-16.3* (5.4)	0.34 (0.30)	-0.07 (0.20)	
ISO-certified firm	15 (7.9)	16 (7.7)	5.4 (5.4)	6.1 (5.5)	0.28 (0.37)	18* (8.5)	-0.52 (0.40)	0.08 (0.21)	
Contract									
Fixed-cost contract		14* (5.4)		5.5 (3.9)					
Time and material contract		5.5 (7.5)		5.6 (5.4)					
Firm and client characteristics									
Number of employees (100)	-0.77* (0.24)	-0.38 (0.040)	0.32 (1.7)	0.38 (1.7)	0.08 (0.13)	7.2* (2.4)	0.23* (0.11)	0.11* (0.066)	
Client is big	0.88 (4.8)	0.43 (4.8)	1.8 (3.4)	1.3 (3.4)	-0.34 (0.23)	7.4 (5.3)	-0.05 (0.24)	0.11 (0.17)	
Client is Indian	4.1 (5.6)	2.8 (5.6)	-0.55 (4.0)	-0.78 (4.0)	0.07 (0.25)	-12* (6.4)	-0.34 (0.33)	0.03 (0.18)	
Project characteristics									
Estimated project size (man-months)	0.11* (0.040)	0.11* (0.040)	0.04 (0.03)	0.04 (0.03)	0.01 (0.22)				
Area is familiar	5.6 (5.1)	-5.9 (5.1)	-1.5 (3.7)	-1.6 (3.7)	0.08 (0.25)				
Platform is familiar	-20* (10)	-18* (10)	-12* (7.4)	-12* (7.5)	0.21 (0.58)				
Y2K, data manipulation, etc.	-13* (7.2)	-13* (7.1)	-9.41* (5.1)	-8.8* (5.1)	0.20 (0.32)				

Standard errors (corrected for clustering at the firm level in columns (5), (7), and (8)) are in parentheses.
An asterisk indicates that the coefficient is significant at the 90 percent level of confidence.
All regressions include the following additional variables: number of employees squared, project size squared, and indicator for whether tools are familiar to the firm.
Sample includes all external projects (167 projects in 92 firms).

precontracting stage, one would have expected the opposite to hold. However, the association disappears in the fixed-effects regression. We also find that project delay is not systematically associated with reputation.

Finally, we examine whether the choice of projects is influenced by a firm's reputation. Once we control for size, young firms do not tend to do shorter projects (column (6)). However, consistent with our model, they do more Y2K projects (column (7)), and their projects are, according to their own assessment, simpler (column (8)).²³

V. CONCLUSION

We set out in this paper to look for evidence that reputation plays an important role in determining contractual outcomes. We argue that the evidence strongly supports this view, although given that we do not directly observe people making use of reputation, and that there are important unobservables that might affect our estimates, some doubts clearly remain.

The conclusion that reputation matters is of course important, for it gives support to a range of theories that are based on limitations of contracting. Moreover, it might suggest part of the explanation as to why the Indian software industry is not much larger (Indian software exports were only worth 3.4 percent of the 1995 worldwide outsourcing business) despite its obvious labor-cost advantage²⁴ and the fact that it is a very labor-intensive industry. Reputation at the firm level is one possible reason. Most Indian firms are simply not trusted enough to be given important contracts. While our evidence cannot directly substantiate this view, the fact that reputation is important within the Indian industry suggests that it may also be important when an American client is deciding whether to go to a firm in India or to one in the United States. The industry clearly recognizes this. When we surveyed the firms, the association of software exporters and the government were actively supporting the firms to acquire ISO

23. We asked the firms to provide all the metrics that were relevant to the project: line of codes, number of modules, of screens, etc. After this we asked them the following: taking everything into account, how would you rate the complexity of the project (on a scale of 1 to 5)? The dependent variable in column (8) is this subjective measure.

24. The United States imports a very large number of Indian software professionals for short-term assignments at a cost of more than twice what they would earn in India.

certification as a way of improving their reputation. As we show, and as they have come to realize, ISO certification is not an effective source of reputation. The association of software producers has now switched to recommending a competing standard (called SEI), specific to the software industry. Our results also suggest that the industry might benefit from the establishment of a credible system for rating firms modeled on credit rating systems, by making it possible for the market to efficiently aggregate all that is known about each firm.

APPENDIX: DATA COLLECTION

The data set used in this paper was collected by one of the authors during the period of November 15, 1997, to February 15, 1998. The survey instrument was developed after extensive conversations with members of the software industry in Bangalore during the preceding summer, and pretested on ten firms before the actual survey began. The sample frame was the list of firms belonging to the government-exporting organization (the Software Technology Park). Almost all exporting firms belong to this organization. We excluded 100 percent wholly owned subsidiaries of large U. S. firms (Microsoft, AT&T, Hewlett-Packard, etc.), as well as the largest Indian firms (TCS, Wipro, Infosys). We interviewed half of the remaining firms. We picked the firms randomly (in alphabetical order, skipping over one firm at a time). We contacted the firms by telephone and arranged an appointment with the CEO, usually in the following week.

The response rate was excellent. No CEO whom we contacted refused to meet with us. In Bangalore, where we spent the longest period of time, all the interviews were arranged. In the other centers, we had to replace a few firms due to unavailability of the CEO or another senior person during the time we spent there. Although nonavailability is in principle not random, we believe that this was not a big enough problem to significantly affect the randomness of the sample. The total sample consisted of 125 firms and 230 contracts.

The interview was conducted with the CEO in most cases, or with another senior person in the firm. We collected general information about the firm, and specific information about the last two projects completed by the firm using a structured questionnaire, available on our web pages. The important limitation is that we could not collect precise information on the client, data

profits, or price charged per man-months by the firm. The project managers involved in the two projects discussed were often present as well, or contacted by telephone for classifications. The interviews lasted from an hour to an hour and a half.

The first part of the questionnaire was straightforward and was answered by the respondent without any need for verification. The project-specific information was either recalled by the respondent or checked from records. The projects had been completed very recently, and the information we collected is part of the data routinely collected by the firms in the course of the project. They therefore had no difficulty answering these questions.

To facilitate the interview and obtain comparable answers from firm to firm at each interview, as we were conducting the interview we drew a timeline representing the different phases of the project (proposal, specifications, development). We used the timeline to ask them what the initial estimate was, when it was revised (after the specifications, at what point in the project), and what the actual effort was. Several questions can be used to determine the actual amount of the overruns. First, we asked them how the actual figure compared with the estimate (making sure whether they referred to the last estimate or the initial estimate), and the reasons for these changes. Second, we asked them separately whether there was any delay due to the client, any delay due to internal difficulty, and any request for changes (and when). In the course of asking these questions, we also asked them whether there was any overrun due to each of these reasons, and we made sure that the numbers corresponded to what they had told us before. If they did not, we probed until we reached consistency (which was usually easily done).

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

REFERENCES

- Bajari, Patrick, and Steven Tadelis, "Incentives versus Transaction Costs: A Theory of Procurement Contracts," mimeo, Stanford University, 1999.
- Crocker, Keith, and Kenneth Reynolds, "The Efficiency of Incomplete Contracts: An Empirical Analysis of Air Force Engine Procurement," *RAND Journal of Economics*, XXIV (1993), 127–146.
- Diamond, Douglas W., "Reputation Acquisition in Debt Markets," *Journal of Political Economy*, XCVII (1989), 828–862.
- Fama, Eugene F., "Agency Problems and the Theory of the Firm," *Journal of Political Economy*, LXXXVIII (1980), 288–307.
- Heckman, James, "Sample Selection Bias as a Specification Error," *Econometrica*, XLVII (1979), 153–161.
- Holmstrom, Bengt, "Moral Hazard in Teams," *Bell Journal of Economics*, XIII (1982), 324–340.

- Kandori, Michihiro, "Social Norms and Community Enforcement," *Review of Economics Studies*, LIX (1992), 63–80.
- McMillan, John, and Christopher Woodruff, "Interfirm Relationships and Informal Credit in Vietnam," *Quarterly Journal of Economics*, CXIV (1999), 1285–1320.
- Pressman, Roger S., *Software Engineering, A Practitioner's Approach*, 4th edition (New York, NY: McGraw-Hill Companies, 1996).
- Tirole, Jean, "A Theory of Collective Reputations (with Applications to the Persistence of Corruption and to Firm Quality)," *Review of Economic Studies*, LXII (1996), 1–22.