

Past Performance and Procurement Outcomes - WORK IN PROGRESS -

Francesco Decarolis, Riccardo Pacini and Giancarlo Spagnolo*

September 27, 2018

Abstract

Reputational incentives may be a powerful mechanism for improving supplier performance. We analyze their role in contract awarding, exploiting an experiment run by a firm which introduced a new vendor rating system scoring suppliers' past performance and linking it to the award of future contracts. We study responses in both price and performance to the announcement of the switch from price-only to price-and-rating auctions. Across the 136 parameters scored, overall compliance improves from 25 percent to 80 percent. Improvements involve all parameters and suppliers, but are more pronounced for parameters receiving a higher weight in the announced scoring auction. Prices do not significantly change overall, but we find some evidence of lower prices right after the announcement when suppliers compete to win contracts to get scored, and of higher prices once they have established a good reputation. Even under an upper bound estimate for this price increase, however, the cost of the policy is below its lower bound benefit estimate, which derives from a reduction in fatal accidents driven by improvements on the subset of parameters involving worksite safety.

To gain a long run perspective on the experiment's effects, we are in the process of updating the paper with data received on January 2018 covering the 2010-2017 period.

*Decarolis: Bocconi University, francesco.decarolis@unibocconi.it. Pacini: Agenzia del Demanio, riccardo.pacini01@agenziademanio.it. Spagnolo: SITE-Stockholm School of Economics, EIEF and Tor Vergata, spagnologianca@gmail.com. We are grateful to Nick Bloom, Luis Cabral, Liran Einav, Rosa Ferrer, Stefano Gagliarducci, Ricard Gil, Luigi Guiso, Hugo Hopenhayn, Neale Mahoney, Enrico Moretti, Gustavo Piga, Piero Sabbatini, Stephane Saussier, Steve Tadelis, Paola Valbonesi and to the participants at the seminars at Barcelona GSE, Berkeley, FGV Rio de Janeiro, Stanford University, Stockholm University, University of Bologna and University of Paris I for their useful suggestions. We are thankful to the firm running the experiment and to the Italian Procurement Authority for providing us with the data. Decarolis thankfully acknowledges financial support from the European Research Council (ERC Starting Grant 679217). Spagnolo is grateful to the Wallander and Hedelius foundation for financial support (Handelsbanken P2013-0162).

I Introduction

The importance of reputational forces linking future business to past performance has been documented in decentralized private markets where contracting is limited by the complexity of the traded good or by the lack of legal enforcement.¹ In this paper, we exploit a firm experiment to quantify, for the first time and in considerable detail, the effects of reputational incentives on quality and prices in a centralized public procurement market where law enforcement institutions are available and traded goods are not too complex.

While opinions on the appropriateness of linking past performance to future procurement contracts could not be more divergent across the Atlantic, very little evidence has been available to inform the debate. This is partly because the idea of systematically using reputational scores based on past performance in public procurement is a relatively recent one. In private procurement, past performance indicators have always affected the selection of suppliers and their behavior because private buyers are free to act upon them, e.g. by refraining from selecting suppliers with a poor track record and favoring those with a good one. In public procurement, this type of discretionary management practice is typically limited. The need to prevent favoritism and corruption led lawmakers around the world to ensure that open and transparent auctions, where bidders have equal treatment (even when they have a very different track record), are used as often as possible. Another way by which lawmakers limit civil servants' discretion is constraining the use of non-verifiable information (e.g. observable, but non-verifiable qualities of an offer or task performed) to justify supplier selection or monetary transfers. This has often limited the type of reputational scores that public buyers could consider to those based exclusively on verifiable past performance. The reputational mechanism that we study will indeed combine the use of verifiable measures of past performance together with transparent, competitive auctions.

A well known feature of competitive auctions in the context of contract procurement, however, is that with imperfect contracts, bolstering competition at the bidding stage might come at the cost of poor ex post performance (Spulber (1990)). Balancing this price versus

¹On complex goods, see Banerjee and Duflo (2000) on the Indian software industry; for the lack of legal enforcement, see the burgeoning literature on electronic platforms recently surveyed in Tadelis (2016).

performance trade-off, and – specifically – how the use of past performance can contribute to that, is a fundamental, yet unsolved, problem of public procurement. With the Federal Acquisitions Streamlining Act of 1994, the US undertook a major policy change that saw the use of contractors’ past performance as the pillar of a new approach to procurement aimed at reducing the rigidity of the procedures built in the Federal Acquisition Regulations. It enabled public buyers to use more flexible purchasing methods similar to private sector management practices, among which is placing a strong weight on suppliers’ past performance when selecting bids.² Federal agencies thus started to record past contractors’ performance evaluations and share them through common platforms for use in future contractor selection.³ Challenges with collecting and sharing this information created obstacles to the full implementation of this reform, but recent legislation has reinvigorated the role of reputational systems for federal procurement.⁴ Interestingly, the EU follows a very different system that exemplifies well the traditional view on past performance in public procurement. Until 2014, the EU Procurement Directives, which coordinate the public procurement regulations of EU member states, essentially forbade the use of past performance as a contractor selection criterion, with the exception of extreme cases of major contractual violations sanctioned by the judiciary. This has been one of the features under broad attack during the recent consultation for the revision of the EU Directives, but only minor changes were incorporated in the new 2014 EU Directives: past performance information can now be used to limit bidders’ participation in the awarding procedures, but not to select among bidders.⁵

Quantifying the costs and benefits of running a reputational mechanism in public pro-

²For an early, influential contribution on the comparison of private versus public procurement management practices see Kelman (1990).

³The reform was pushed by Steven Kelman (see footnote 3 above) when he served as Administrator of the Office of Federal Procurement Policy in the Office of Management and Budget from 1993 to 1997, playing a lead role in the Administrations “reinventing government” effort that led, among other things, to the Federal Acquisition Streamlining Act of 1994 and the Federal Acquisition Reform Act 1995.

⁴According to a Congressional Research Services report, Manuel (2015): *Reports alleging that contractors received new contracts or orders despite poor performance under prior ones have recently prompted interest in the role that evaluations of past performance play in contracting, as well as attempts by some Members of Congress and the Obama Administration to strengthen requirements pertaining to performance evaluations. (...) the Administration updated the FAR in 2013 to standardize the factors used in evaluating contractors performance, and require that all past performance information be entered into the Contractor Performance Assessment Reporting System. [And to implement the Congress’ initiatives] P.L.112-81 and P.L. 112-239.*

⁵See Gordon and Racca (2014) and the responses to the EU consultation at http://ec.europa.eu/internal_market/consultations/docs/2011/public_procurement/synthesis_document_en.pdf.

curement is therefore of first order importance, especially given the sheer economic size of this market,⁶ but little reliable evidence is available at present. The main challenge to studying the benefits of using reputation in public procurement is that changes in the way reputation is used are rarely observed. Our study overcomes this problem by exploiting a very rich set of data related to the introduction of a past performance monitoring system in the procurement practices of a large Italian multi-utility company (“the Firm”). This company provides water and power to a vast area in central Italy and, as illustrated in Table 1, is close to some of the main US operators, like BGE (Baltimore), ComEd (Chicago), LADWP (Los Angeles) and PECO (Philadelphia), in terms of several observables such as the number of employees, the customer base, the power grid size, the total turnover and the expenditures for maintaining and upgrading the power grid. For these firms, the quality and safety standards followed in the execution of the construction works performed to maintain and update the power grid are essential not only for an effective functioning of the grid, but also to limit both the number and the severity of the workplace accidents which mark this industry.

In 2007, to understand how it could improve contractual performance in terms of both quality and safety, the Firm started an “experiment” with a reputational mechanism. First, it selected two related types of works in the electricity sector involving either public illumination works or the construction/maintenance of electrical substations. Then, for these works, it laid down a list of 136 observable parameters measuring both quality and safety features of the job carried out. Three months after the new audits had begun, the Firm made its first public statement explaining that the results of the new audit system would be converted into a numerical “reputation index” and that, after a few more months of data collection, this index would be used to award new contracts. For the Firm, this was an experiment in the sense that it introduced the new audit system only for a subset of contracts (all of those involving public illumination and electrical-substation works, but none of those related, for instance, to water delivery) and its stated goal was to learn whether the new audit system could be beneficial for its overall procurement. While this experiment does not satisfy all the characteristics of an ideal field experiment, it is nevertheless a very useful natural experiment

⁶The OECD estimates that public procurement represents 15-20 percent of GDP in developed countries.

that , to the best of our knowledge, allows us to study what is the first quantitative analysis of the effects of reputational mechanism in public procurement.⁷

This paper studies the evolution of both price and performance around the time when the new audit system was publicly announced, but before the scoring auctions incorporating reputation were first used. Figure 1 shows the timing of the experiment by separating into two lines the timing of the audit and awarding systems reforms. The audit system (top line) changed in late 2007, while the awarding system (bottom line) changed in mid 2010. For two and half years after the new audits were introduced, the Firm continued awarding contracts through lowest price auctions. During this period, the Firm recorded the performance of its contractors and held a series of five meetings with them to explain the functioning of both the new system and the future scoring auctions. In the first meeting (indicated as $t1$ in Figure 1), the Firm announced that the new audits would be used in the future to switch from price-only to price-and-rating auctions, with a linear scoring rule auction assigning 75 percent of the weight to the price offered and 25 percent to the reputation index. In the following four meetings (indicated as $t2, \dots, t5$), it updated contractors on the forthcoming award rule change and it disclosed how compliance with the 136 parameters was evolving across all the contractors audited. Our analysis focuses on how price and performance responded to these announcements, before the switch to the scoring auctions occurred.

In the first part of the paper, we analyze how compliance with the parameters monitored evolved in response to the timing of the five public announcements. Using audit data for the years 2007-2009, we find clear evidence of a substantial change in contractors' behavior: compliance in the 136 parameters increased from 25 percent before the first announcement ($t1$) to more than 80 percent after the fifth announcement ($t5$). We find that essentially all active suppliers improved their compliance in similar ways and they did so strategically, with compliance increasing relatively more for those parameters with higher weights in the

⁷For the distinction among types of firm-level experiments see Bloom et al. (2014) who state that “in personnel economics there is a tradition of exploiting changes in firm policies initiated by a CEO (a natural experiment such as Lazear (2000)) or engineered by the team of researchers (e.g., the fruit farm field experiments summarized by Bandiera, Barankay and Rasul (2011)).” In our case, the experiment was autonomously decided by the Firm’s CEO, but we had an active role in its design and implementation, in the collection and analysis of the first few years of data, as well as in the five communications to the suppliers extensively discussed below. For a general discussion of experiments in economics see List and Reiley (2008).

computation of the reputation index. While this is compatible with a strategic allocation of effort, multi-tasking likely occurred mostly within the set of parameters scored and did not cause a reduction in the effort on unmonitored tasks. In fact, according to the Firm's engineers the broad set of parameters chosen was exhaustive in terms of determining safety and quality for the chosen contract types. Moreover, cost and time overruns - two proxies that are often used for performance, but that the Firm decided not to include in the index - did not worsen. Indeed, the Firm's own evaluation of the experiment also found the increase in performance to be fully satisfactory.

The external validity of what can be learned from a firm-level experiment is a typical concern in the literature (Bloom et al. (2014)). In our case, however, it is reassuring to observe that the Firm is similar to both some other major operators active in the US, like the multi-utility companies of the four US cities mentioned above, and to the other companies providing the same services in Italy. Indeed, for this latter group of operators, we were able to verify the similarity with the Firm at the level of the individual contracts procured.

The availability of these contract-level procurement data for the other operators active in Italy is also essential for the second part of our analysis. This second part studies whether and to what extent the improvement in performance caused a change in procurement costs. This is key to evaluate the desirability of the policy change and is made feasible by the nearly ideal timing of the experiment. While the continued use of price-only auctions for more than two years after the announced switch to scoring auctions (see Figure 1) was due to legal constraints, this feature is what allows us to study the price effects linked to performance without confounding effects driven by changes in the market structure. Specifically, for the price-only auctions that we analyze, all suppliers are treated symmetrically regardless of their rating and this reduces the chances that our estimated price effects are driven by the barriers to entry that a sudden switch to price-and-rating auctions would have erected.

The empirical strategy used in this part of the analysis takes advantage of a second dataset containing information on all the contracts awarded not only by the Firm, but also by all other Italian public contracting authorities (CAs) during the period 2005-2010. We use the variation across procurers and over time to develop a difference-in-differences

estimation strategy. Hence, a key pillar of this strategy is the comparability between the Firm and other contracting authorities that did not alter their audit system. Our first set of findings shows that, if we consider the date of the first announcement to be the one characterizing the occurrence of the policy change, then there is no significant effect on the price paid by the Firm. More specifically, by looking at any symmetric window of time around the first announcement, prices remain stable on average. However, when we extend the empirical model to account for the evolution of compliance, the price response appears more nuanced. Using the results from the first part of the analysis, we partition the period after the first announcement into a first phase when compliance grows, and a second phase when it flattens out at high levels. When we extend the baseline difference-in-differences model to account for these two phases, we find that the original finding of no effect results from the combined effects of prices declining when compliance improves, but increasing after compliance stabilizes. We interpret this evidence as suggestive of a first phase in which suppliers compete harder to win contracts - only contract winners can be audited and, thus, earn or improve their reputation index. That is, winning a contract has the additional benefit of improving the chances of winning future contracts. After all contractors have earned a high reputation index, however, this benefit is outweighed by the increased cost of high compliance, and auction prices become correspondingly higher. The estimates indicate that the overall cost increase of higher quality, after the phase of competition for reputation, ranges between zero and 9 percent of the contract reserve price.

We complement this part of the analysis with a back of the envelope cost-benefit calculation. Under the worst case estimate of the price increase observed in the final part of the sample, the increase in procurement cost amounts to €2.4 million per year. By using the OECD figures for the value of a statistical life together with the same statistical model employed by the Firm's engineers to map the relationship between changes in parameter compliance and the occurrence of fatal accidents, we estimate that the benefit from increased compliance on the safety parameters ranges between €3.5 and €5.3 million per year. This is a lower bound for the benefits that does not account for both quality improvements and safety improvements associated with non-fatal injuries. Hence, the fact that it exceeds the upper bound of the cost leads us to conclude that the policy change was beneficial.

The final part of the analysis studies whether the observed effects are the result of changes in the selection of contractors bidding or in their behavior. The evidence is definitely compatible with the presence of moral hazard: suppliers that are observed bidding both before and after the new rating system is announced stop offering suspiciously low prices. These are precisely the abnormal, low ball bids often associated with poor contractual performance. On the other hand, we find only limited effects of selection based on three features in the data. First, while several suppliers leave the market, the timing of their exit is not associated with the announcements. Second, both the firms that leave the market and those that remain have similar bidding patterns. Third, along many observable characteristics, the firms leaving the Firm's auctions are no different from the suppliers that leave the auctions of another large multi-utility company that did not participate in the experiment and that we use as a benchmark. Thus, the main result from this study is that the gains from curtailing suppliers moral hazard when executing contracts may be higher than those from always bolstering price competition, and that a reputational mechanism based on objective past performance can be a powerful tool to achieve this goal.

II Literature

Our paper contributes to a number of literatures on both reputation and public procurement. As mentioned earlier, it contributes to the empirical literature trying to quantify the effects of reputation on quality, prices or contract choice. The already mentioned study of the Indian software industry by Banerjee and Duflo (2000) develops a model of reputation formation as a signaling game in which each party can propose a type of contract and there is systemic overrun due to the complexity of the goods/services supplied. Empirically, it finds that firms with a better reputation are more likely to be involved in time-and-material contracts and, in most cases, pay for a smaller share of cost overrun. It also finds less firm-generated overrun and less total overrun in contracts involving more reputable firms. This study is close to ours in spirit, but at the same time is also rather different. In contrast to that study, we have a fixed contract form and a single public buyer, and we analyze the effects of announcing the

introduction of a reputational index based on objectively measured past performance.

The analysis of the effects of the reputational mechanism on prices also connects our study to the literature on reputational mechanisms in electronic platforms recently surveyed by Tadelis (2016). A frequent finding in this literature is that reputation affects the probability of selling, but, as in our study, the effect on price is typically small. Within this literature our paper is probably closer to Klein, Lambertz and Stahl (2016) who, like us, find that an effective reputational mechanism tends to curb in particular moral hazard. Clearly, a major difference is that this literature focuses on reputation based on subjective feedback and on quality outcomes also measured through subjective feedbacks, apart from prices. Instead, in our experiment we have access to detailed, objective performance measures and a structured reputational mechanism whereby the quality of past performance affects future business directly through the scoring rule (announced) for future procurement auctions.

Our study also contributes to the literature on the impact of competition in public procurement with imperfect contracting. Previous theoretical papers have shown that, under imperfect contracting, the results on the optimality of open auctions (e.g., Bulow and Klemperer (1996, 2009)) need not to apply. In particular, Spulber (1990) shows that in the construction sector, where contracting is typically imperfect, open competition spurs adverse selection and ex post opportunism of contractors. More generally, Manelli and Vincent (1995) show that when gains from trade are mainly in non-contractible quality dimensions, price-only auctions are the worst of all conceivable allocation mechanisms. Bajari and Tadelis (2001) show that bilateral negotiations may be better than open competition for highly complex projects because of the costs of specifying ex ante all contingencies. On the empirical side, only a handful of papers have tried to quantify how procurement design affects both quality and prices. A few recent examples are: Bajari, McMillan and Tadelis (2009) which builds on the insights from Bajari and Tadelis (2001) to empirically analyze the selection of contractors and award mechanisms based on reputation; Decarolis (2014) on how the use of different awarding rules affects the price vs. performance trade-off, also depending on the effort intensity placed by procurement agencies on screening bids for their responsiveness; Liebman and Mahoney (2016) on how yearly fiscal rules on expenditures in US federal pro-

curement generate poor quality in information technology projects bought close to the end of the fiscal year; Lewis-Faupel et al. (2016) on how the introduction of electronic auctions lead to improved quality of infrastructure in India and Indonesia; and Kang and Miller (2016) on how the desire to ensure a high level of non-contractible quality might be one of the reasons for the very low levels of competition in U.S. government procurement.

Our findings are also related to a recent wave of studies highlighting the importance of adopting a dynamic framework to understand public procurement markets. On the theoretical side, that past performance and reputation may play a crucial role for the understanding of repeated public procurement under imperfect contracting was recognized in several studies (e.g. Kim (1998), Doni (2006) and Calzolari and Spagnolo (2009)). Overall, this theoretical literature concludes that when contracting is imperfect, (buyer) discretion taking into account past performance can have positive effects on public procurement outcomes. A recent paper by Chassang and Ortner (2016) confirms both theoretically and empirically that a dynamic approach to repeated procurement is indeed essential to understand the role of “minimum bid requirements” on supplier collusion and procurement outcomes. On the purely empirical side, Marion (2016) finds that a dynamic system, allowing firms active in the Iowa Department of Transportation procurement to get exemptions from an affirmative action policy requiring subcontracting to “disadvantaged enterprises” by accumulating enough “points” from previous subcontracts, is superior to a static alternative; Gil and Marion (2012) analyze the effect of repeated interaction in the subcontractors market for California’s highways and find that prices are lower when future contracts are expected to arrive, even if there was no previous joint collaboration, but that they are even lower if the contractor and subcontractor had collaborated before; and Coviello, Guglielmo and Spagnolo (2016) show that restricted auctions, where only invited suppliers can bid, may lead to at least as good outcomes as open auctions, and that, when these more discretionary procedures are used, dynamic incentives are at play: incumbents win more often if they delivered better performance in the past, and they deliver earlier and at lower cost. Our paper contributes to this literature by empirically measuring the power of dynamic incentives induced by a structured reputation mechanism that links current performance to future contract awards.

This last aspect also directly links our paper to the literature on scoring auctions in public procurement. In two important recent contributions, Lewis and Bajari (2011) and Lewis and Bajari (2014) use data on US highway construction to show how the high costs that slow highway completion inflicts on commuters can be substantially reduced by curbing moral hazard, optimizing the structure of the procurement contract and inserting time incentives in the scoring rule of the procurement auction. We also show that inserting a measure linked to quality of procurement in the scoring rule can substantially limit moral hazard, but our measure is based on past performance rather than on the performance of the tendered contract. Our evidence that better aligning suppliers' incentives to deliver a high performance need not result immediately in higher prices is also found in a recent laboratory experiment on a similar topic by Butler et al. (2014) and in a study of the French public procurement of services by Beuve and Chever (2014). The latter study is close to ours as it empirically studies price and performance responses in 102 contracts for cleaning services where a public buyer, in response to a court ruling invalidating its use of past performance to exclude a bidder, decided to invest in greater contract completeness.

The particularly detailed performance measures, from random audits by centrally managed inspectors we have access to, relate our paper to the work of Olken (2007) on corruption in Indonesia, where similarly detailed performance measures are used to evaluate the effectiveness of different systems to curb corruption. Related to the issue of corruption, Bandiera, Prat and Valletti (2009) show that, in the context of the experiment that we study, corruption concerns could be less of a priority than finding ways to increase performance, e.g. with reputational mechanisms. Their paper studies waste in the procurement of Italian goods and manages to quantify active waste linked to deliberate corruption and passive waste linked to inefficiency, incompetence, red tape, etc. It offers important, empirically grounded indications to policy makers. Our paper is a step in the same direction, as it identifies and quantifies the likely costs, benefits, and channels of a change linked to the introduction of a reputational mechanism based on objective performance measures in public procurement.

The new data obtained covering the period up to 2017 during which first the Scoring Rule was introduced and then removed to be replaced by a first price rule with increased

monitoring allows us to compare regime changes. In this we are similar to Athey, Coey and Levin (2013), Decarolis (2014), Coviello, Guglielmo and Spagnolo (2016) and Hyytinen, Lundberg and Toivanen (2018).

Finally, on the policy side, this study makes an important contribution to the above mentioned debate in Europe and the US. The slow adoption in US federal procurement of the 1994 policy reform on reputational mechanisms suggests that its implementation costs are not trivial. Our paper provides the first empirical evidence on the likely benefits of implementing a reputational mechanism in public procurement, suggesting they may be rather large. This is important in the European context too because the reliance on past performance could reduce the worry, expressed by several experts (see, for instance, Saussier and Tirole (2015)), that the ongoing shift toward a discretionary system of awarding and renegotiation procedures - under the new EU Directives of 2014 - will create distortions in the EU public procurement market.

III The Context of the Experiment

The experiment entailed the introduction of a new vendor rating system by a large procurer, “the Firm.” The Firm is one of the leading Italian multi-utility companies in the business of water and power. Concerning power, the Firm is vertically integrated, owning and operating the majority of its generation, transmission and distribution systems. From this point of view, it is very similar to some of the largest US power operators such as the Los Angeles Department of Water and Power (LADWP), ComEd (Chicago), BGE (Baltimore) and PECO (Philadelphia).⁸ Moreover LADWP shares with the Firm the integration of the water business as well.⁹

Operational efficiency is central for these firms to be successful. Indeed they all devote a significant budget to long-term power reliability programs developed to replace aging infrastructure and make permanent repairs to generation, transmission and distribution in-

⁸ComEd, BGE and PECO are all subsidiaries of the Chicago-based Exelon Corporation.

⁹Both horizontal and vertical integration offer some advantages in these sectors such as bigger customer base, reduced exposure to market volatility and access to shifting profit pools.

infrastructure that fail during outages. These programs provide for the maintenance, upgrade and replacement of transformers, poles, underground cables, underground vaults, station transformers, distribution and receiving stations, in order to guarantee a reliable and secure power system.¹⁰ Moreover, specific investment is required to integrate increasing amounts of intermittent renewable generation resources and transformational technologies such as energy storage, electric vehicles, and other aspects of the smart grid. Hence, resource planning and infrastructure asset management need to be aligned to ensure aging assets are replaced with infrastructure that is able to meet new system requirements and maintain reliability with a modern generation mix. Table 1 reports some figures for 2015 about the Firm and four US power operators. The table shows that, for instance, the Firm is identical to PECO (Philadelphia) in terms of customer base and that it sits between LADWP (Los Angeles) and BGE (Baltimore) in terms of the size to its power grid. The similarities underscored by the values in Table 1 are important to gauge the external validity of the experiment that we will describe. The next section will illustrate the comparability of the Firm with other operators in Italy at a finer level by comparing individual contracts.

As shown in Table 1, the Firm spent about US \$200 million in 2015 on procuring construction and maintenance works. Since the Firm is controlled by a majority shareholder that is a local administration, the procurement of these works must follow the awarding rules laid down in the Italian Public Procurement Code (“the Code”).¹¹ Being a multi-utility company, the Firm falls within the “special sectors” which enjoy some flexibility in applying the Code. Starting in 2007, this allowed it to begin an experiment with a new vendor rating system. The basic idea was to set an objective measure of contractual performance, with the plan of using its ratings in the awarding stage of future procurement processes. This was seen as a useful attempt to improve performance without relying on Italian courts, which are well known to be inefficiently slow (see Djankov et al. (2003)). Below, we first describe how the reputation measure was constructed and then how it was incorporated into the auctions. The latter step faced a series of legal obstacles that we describe at the end of this section and

¹⁰These are all works exposing workers to safety hazards linked to electricity-induced accidents.

¹¹The Code, Legislative Decree 163 of 12 April 2006, is the law that implemented the EU Procurement Directives 17/2004 and 18/2004 in Italy and that is relevant for the period that we study.

that crucially affected the timing of the experiment. In particular, they induced the Firm to slow down the implementation of the announced switch to price-and-reputation auctions. The timing of the experiment, as well as the time span of our analysis, is shown in Figure 1.

A. The Reputation Index

The Firm designed an experiment with a vendor rating system focusing on the procurement of works in the electricity sector. It selected two types of works, involving either public illumination or electricity distribution (mostly entailing the maintenance of electrical substations and wires), that were considered sufficiently homogenous to define a list of items key to assessing contractual performance. A total of 136 parameters were identified for this goal. As shown in Figure 1, from October 2007 all new and ongoing contracts for public illumination or electricity distribution began to be (randomly) audited for these parameters. Before then, engineers used to inspect work sites and write descriptive memos about their conditions. These memos, however, did not translate into any quantitative assessment of performance and they only served to keep track of how the work was evolving.¹²

As Table 2 reports, the set of 136 parameters is divided into 12 categories, further divided into 2 macro classes: “safety” (51 parameters; 7 categories) and “quality” (83 parameters; 5 categories). For instance, “Equipment and machinery,” the first category in Table 2, comprises 5 parameters involving the adequacy of both the formal documentation and the physical conditions of equipment and machineries.¹³ Parameters in this category are quite general and can be inspected for essentially all work sites. Other categories, instead, involve parameters specific to a subset of contracts only. For instance, the 25 parameters in “Underground works” involve features that are assessed only for jobs involving underground wires and electrical substations. In a typical audit, 33.08 parameters are scored.

[INSERT TABLE 2 APPROXIMATELY HERE]

The system works as follows. Scores are collected by teams of rotating auditors (Firm engineers) in one or more visits to the work sites, with a score assigned to each of the 136 parameters. The score is 1 if the value is “compliant,” zero if “not compliant” or “n/a” if it is impossible to inspect. Two key features of the system entail randomization: which contracts are audited and which engineers from the Firm are assigned to the team inspecting

¹²Although they could have been used to enforce penalties, penalties are rarely enforced in this market.

¹³While important for the safety of the work site, these features might also influence the quality of the work executed, thus making the distinction between the two classes of quality and safety blurry.

each work site are both determined through a process of random drawing. Thus, a single contract might be audited one or more times and by the same or different engineers.

The scores on the individual parameters are then aggregated into a unique reputation index (RI). Each parameter is associated with a weight, ranging from 2 to 10, and the RI is calculated as a weighted average mean across a predefined time span according to:

$$RI = \frac{\sum_{i=1}^m \sum_{j=1}^{136} p_{ij} u_j}{\sum_{j=1}^{136} u_j}, \quad (1)$$

with $p_{ij} \in \{0, 1\}$ indicating the score obtained in each of the 136 parameters over all the m audits considered and with $u_j \in \{2, 3, \dots, 10\}$ being the weight attached. Hence, the reputation index can range from 0 to 1 and entails no differential discounting of audits taking place within the predefined time window. In the period that we study, this window was announced as equal to one year.¹⁴

Finally, it is worth mentioning that analogous concerns with contractors' performance exist in similar companies and are addressed in related, albeit not identical, ways. For instance, the LADWP *Contractor Performance Program* states that: “A *Contractor Performance Scorecard system will be maintained on all contractors that have been identified by end users and Contract Administrators of having not met the terms of the contract/purchase order. An “infraction point” will be assessed against the contractor when contractual terms are not met and are communicated to the contractor. Each point will stay on a contractors record in Supply Chain Services for a 12 month period after infraction. If a particular contractor receives a total of 3 or more points within a 12 month period, that contractor may be debarred from bidding with the Department of Water & Power for a period of up to 5 years.*”

B. The Scoring Rule Auction

The regulations in the Code require the Firm to award its contracts through auctions. Incorporating the RI in the awarding process required a switch of the auction award criteria

¹⁴More precisely, if the tender is announced in month $n = \{1, \dots, 12\}$, then the RI uses all audits collected in the previous 12 months starting from $n - 2$ (if the announcement take place before the 10th day of the month), or $n - 1$ (if it takes place after).

from the lowest price to the most economically advantageous tender (MEAT). As shown in Figure 1, on December 2007, three months after the introduction of the new auditing system, the Firm announced to its contractors its intention to switch to MEAT by adopting a linear scoring rule auction whereby the contract is awarded to the firm with the highest score S calculated as:

$$S = w_{price} \left(1 - \frac{\text{Price offered}}{\text{Reserve price}}\right) + (1 - w_{price})RI, \quad (2)$$

where w_{price} is the weight assigned to price relative to that assigned to the RI. The firm held five meetings, marked as $t1, \dots, t5$ in Figure 1, to demonstrate the new system to its contractors:

- At the December 2007 meeting ($t1$), the contractors were given a detailed presentation of the forthcoming reform and it was explained how, from a status quo of a $w_{price} = 1$, the new system would have entailed switching to a $w_{price} = 0.75$. Numerical simulations were presented to demonstrate the benefits of accumulating a high RI. Furthermore, the Firm announced that the RI, calculated as above, would apply exclusively to those bidders audited at least 7 times in the relevant time window.¹⁵ Otherwise, a bidder would be assigned an RI equal to the average RI of the bidders in the auction. The same averaging rule was going to be used for new entrants (i.e., firms never audited).
- During each of the following four meetings held between April 2008 and January 2009 ($t2, \dots, t5$), the Firm gave updates on the functioning of the vendor rating system, confirming what was announced in $t1$ and showing how the recorded compliance was evolving. Figures and tables documenting the evolution of compliance on individual parameters across each contractor were disclosed, but without revealing publicly the true identity of contractors (past performance data are classified as confidential).

Thus, although the first scoring rule auction took place well after $t1$, contractors' incentives had likely changed already at $t1$. At this time, suppliers learned that all contracts had become dynamically linked through the RI formula. Moreover, at each of the $t2, \dots, t5$

¹⁵This requirement concerns the number of audits and not the number of contracts, as a supplier can be audited multiple times for the same contract.

announcements they also learned their competitive situation relative to the other suppliers in terms of the RI accumulated. This is suggestive of potential changes in strategic entry and bidding choices that our analysis will seek to uncover.

C. Legal Limits and the Timing of the Experiment

Several features of the experiment described above, including the types of measures entering the RI, the use of a scoring rule auction and the slow transition to the latter, are all closely linked to the legal institutions within which the Firm operates. Without entering too much into the intricacies of the legal system, it is worth mentioning that the use of reputational elements for the selection of contractors in public procurement has received widespread attention in the drafting of the EU Procurement Directives. Directive 18/2004 required that “contracting authorities shall treat economic operators equally and non-discriminatorily and shall act in a transparent way”¹⁶ and that competition is a pillar of the procurement system.¹⁷ For special sectors, however, Directive 17/2004 was less stringent, since it allowed public buyers to institute their own qualification systems or, in general, to select potential candidates to be awarded on the basis of their technical and professional skills, chosen at the discretion of the contracting authorities. The unique limit in the choice of such criteria is objectivity: “contracting authorities which select candidates for restricted or negotiated procedures shall do so according to objective rules and criteria which they have established and which are available to interested economic operators.”¹⁸ Thus, reputation indicators can be used if based on measurable parameters that are verifiable by third parties and agreed upon by contractors. These features shaped the choice of the parameters in the RI.

The problem arises in the awarding phase. Since the EU gives special prominence to the free and fair competition principle, the use of reputation as an award criteria in public procurement can constitute an unfair advantage for the incumbents and a disproportionate

¹⁶Art. 2 of Directive 2004/18/EC of 31 March 2004 on the coordination of procedures for the award of public works contracts, public supply contracts and public service contracts.

¹⁷“Contracts should be awarded on the basis of objective criteria which ensure compliance with the principles of transparency, non-discrimination and equal treatment and which guarantee that tenders are assessed in conditions of effective competition” (Recital 46, Directive 2004/18/EC). “Non-discriminatory criteria should be indicated which the contracting authorities may use when selecting Competitors and the means which economic operators may use to prove they have satisfied those criteria” (Recital 39, Directive 2004/18/EC).

¹⁸Art. 54 comma 2, Directive 17/2004/CE.

disadvantage for new entrants: a potential supplier with no past experience cannot enjoy any reputational premium with respect to pre-existing competitors. This may reduce entry and competition and violate the general principle of equal treatment.¹⁹ In the contract awarding phase, the MEAT is the method that allows criteria other than price to be considered. The EU Court of Justice, however, ruled that contracting authorities, when evaluating quality with the MEAT, should consider the object of the tender and not the bidder’s characteristics.²⁰ The Italian Procurement Authority reaffirmed the same principle.

These restrictions on the use of the MEAT had major impacts on the timing of the experiment. First of all, while the Firm was determined to use the RI to improve contract performance,²¹ the risk of violating the EU Court of Justice ruling on the MEAT was perceived as a threat. This caused delays in the switch to the scoring rule of about two and half years. Then, after only a dozen scoring rule auctions had taken place (between the second half of 2010 and the beginning of 2011), a new management team took charge of the Firm. It opted for a more conservative interpretation of the rules and returned to price-only auctions, despite maintaining the new audit system described above for monitoring performance.

IV Data

The analysis is based on two sets of data: audits data, covering the performance recorded through the new auditing system; and auction data, covering bidding and other auction-related information. The data on audits, in particular, represent a rather unique opportunity to observe contractual performance. Their major limitation is that a link with the auctions

¹⁹Clearly, how to treat new entrants is purely a designer choice (see Butler et al. (2014) for a discussion). However, the description of all reputation-based systems as a form of incumbency advantage is how the policy debate has framed the problem.

²⁰See judgments in cases C-488/01 or C-31/87.

²¹Some data and experiences show that penalties are not effective because they are not even applied; for example, a study conducted for Consip, the Italian central purchasing agency, on a sample of procurement contracts for goods and services, demonstrated that penalties were applied in just 3.7% of eligible cases. In Japan, the difficulties in ensuring a good contractual performance induced public officials to coordinate bidding cartels, a practice know as *kansei-dango*. For a quantitative assessment of collusion in Japanese procurement auctions and a quantification of its potential costs see Kawai and Nakabayashi (2016).

dataset cannot be established because of anonymity requirements imposed by the Firm.²²

A. Audits

The “Audits data” is a panel dataset that contains the outcomes of all the audits performed between the introduction of the new auditing system, on October 16, 2007, and November 19, 2009. The Firm provided us with this dataset, which contains 64,537 observations recording the scores assigned to each of the 136 parameters inspected during 1,951 audits involving 187 contracts and 45 different contractors. Table 3 reports some summary statistics, aggregating parameters at the level of the 12 categories. The table shows that there is substantial heterogeneity in how many times the parameters in each category are inspected. It also reveals that about 20 percent of the observations involve public illumination (PI) works, with the rest being electricity distribution works. The last three columns of the table report the share of compliant parameters, dividing the sample period in three phases: pre $t1$, post $t5$ and between these two phases. For nearly all categories there is an improvement over time; for instance, for “Work site safety” the average pre $t1$ compliance is 30 percent, between $t1$ and $t5$ it is 61 percent, and post $t5$ it is 81 percent. While the size of the increase differs across categories, Figure 2 confirms that average compliance grows from about 25 percent to 80 percent. The average compliance measure in this figure is calculated by averaging together the 0/1 scores assigned in all audits taking place during the month of reference and weighting them by the weights used in the RI formula. The vertical bars show the 95 percent confidence interval for the mean. Their size tends to decrease over time, due to less variance in the recorded compliance. The next section explores the timing and sources of the evidence shown in Figure 2.

[INSERT TABLE 3 and FIGURE 2 APPROXIMATELY HERE]

B. Auctions

The second dataset contains data on the awarding of public procurement auctions held be-

²²While the awarding of a public contract is considered information that must be accessible to the public to limit corruption risks, the performance of contractors is considered sensitive information. An analogous distinction is present in the US regulation as well, so that, in accordance with FAR 42.1503(4)(d), information that is stored in the Past Performance Information Retrieval System (PPIRS) is classified as Source Selection Sensitive and is not releasable unless directed by the agency who submitted the data.

tween 2005 and 2010 for the type of works involved in the Firm’s experimentation. These works belong to a well-defined contract category identified by the Italian regulation as “OG10,” which makes it feasible to select comparable projects. These data cover the Firm as well as all other Italian public contracting authorities awarding OG10 contracts, and their source is a private company, Telemat spa, which provides information on both past and prospective public tenders to firms subscribing to its services. The data include the object of the contract, the reserve price, the awarding price and date, the identity of both the procurer and the winning contractor, and various other contract-specific information. For instance, by using textual search methods we were able to separate contracts into those involving public illumination and those involving electrical substations. For a subset of auctions, we integrate the data with the information on losing bids and on the subsequent life of the contracts.²³

The key dependent variable for the price analysis is the winning discount. An important feature is that the reserve price over which such a discount is calculated is unlikely to be affected by the Firm’s reform. This is because the Firm (as well as all other CAs subject to the Code) is not in full control of the reserve price. This quantity, publicly known to suppliers at the time of bidding, is obtained by multiplying input quantities (estimated by the procurer’s engineers) by their prices and summing up these products. Crucially, input prices are not the current market prices, but the list prices set every year by the region where the Firm operates and used exclusively by contracting authorities to calculate reserve prices. Table 4 reports summary statistics for both the reserve price and the other relevant variables by dividing the auction dataset into four subsets: auctions held before or after $t1$; and held either by the Firm or by other procurers. The comparison of the top and bottom panels on the left side of Table 4 shows that the average winning discount in the Firm’s auctions declines slightly, from 22.8 percent to 20.6 percent, but not in a statistically significant way. A similar conclusion is suggested by the visual inspection of the scatter plot in Figure 4 which shows a lack of any clear pattern across the winning discounts in the Firm’s auctions.

²³For these additional data we have three sources: *a)* the Firm gave us access to the full set of bids - both winning and losing bids; *b)* for some of the other contracting authorities, we obtained the same information on losing bids through textual analysis of the official documents of the contract award; *c)* information on the subsequent life of the contracts, including time and cost renegotiations, was available from 2005 up to the first half of 2008 through the dataset of the Italian Procurement Authority, which covers the universe of all public procurement auctions with a reserve price of at least €150,000.

However, for the post- $t1$ period, the final phase of the sample contains lower discounts than those present right after $t1$. The following analysis will study this aspect in depth.

[INSERT TABLE 4 APPROXIMATELY HERE]

[INSERT Figure 3 APPROXIMATELY HERE]

The difference-in-differences analysis that we will present in the next section will broadly confirm this overall lack of clear price changes associated with $t1$. For this analysis, we will exploit the presence of control group auctions that are close to those held by the Firm along many observables: Table 4 shows that the winning discount, the contract duration and the share of public illumination contracts are quite similar across the treatment and control groups, but that the awarding price tends to be higher in the Firm's auctions.²⁴ To further evaluate the comparability of the winning discounts in the two groups before $t1$, we report in Figure 4 its evolution over the sample years: the straight line represents the treated group auctions, while the dashed line represents the control group auctions. Before $t1$ the two series have very similar trends - thus confirming the usefulness of a difference-in-differences approach - and are also similar in their levels. The figure also offers an insight into the price dynamics that the Firm experienced post $t1$: discounts increased right after $t1$, but then, roughly after $t5$, they decreased. The very different behavior in the control group suggests that this is likely due to the Firm's reform and not to changes in market conditions. The following analysis will establish these effects formally and offer an interpretation for them.

[INSERT FIGURE 4 APPROXIMATELY HERE]

V Empirical Analysis

This section analyzes the effects on price and performance of the Firm's announcements. We begin from an assessment of performance using the audits data. We then move on to

²⁴Nevertheless, it is important to stress that the main effort to ensure the comparability of the auctions was at the data collection stage, where we selected only auctions that, in terms of their object, were a close fit to the public illumination and electricity distribution contracts auctioned off by the Firm.

an analysis of prices exploiting the auctions data. We finally conclude with a cost-benefit analysis that builds upon our price and performance estimates.

A. Effect of Announcements on Performance

Figure 2 shows a marked improvement in compliance during the sample period. The three questions that we explore here are: can this increase be associated with the timing of the announcements? Is its magnitude confounded by composition effects? What does it reveal about suppliers' behavior? Regarding the first question, Table 5 reports the results of Chow and Bai-Perron tests for the presence of structural breaks in the time series of the compliance measure. As for Figure 2, the variable analyzed in the first two columns of the table is the monthly weighted average compliance across all parameters. The next two columns restrict the parameters to those in the quality class, while the last two columns use the subset of parameters in the safety class. The top panel of the table reports the results of a Chow test for the presence of one break at t_1 (odd numbered columns) and five breaks, at t_1, \dots, t_5 (even numbered columns). For all six tests, we reject the null of no breaks in favor of the alternative of breaks at the specified dates. More interesting, however, are the outcomes of the Bai-Perron tests, where we do not specify the dates of the breaks and instead let the test determine them, either without specifying how many breaks there are (odd numbered columns) or specifying that there are 5 breaks at unknown dates (even numbered columns). The test results are a clear indication that t_1 is a breakpoint. This confirms that the first meeting held by the Firm with the suppliers to explain the RI system had a significant impact on their contractual performance. As regards the other break dates, all tests allowing for an unspecified number of breaks identify a break near $t_5 + 1$.²⁵ This is also quite revealing since, by the fifth meeting, suppliers had found out that average compliance had reached a fairly high level across all active suppliers and parameters. As discussed below, this likely changed the strategic environment in the auctions, through a change in the perceived value from further improvements in compliance.

[INSERT TABLE 5 APPROXIMATELY HERE]

²⁵Either exactly at $t_5 + 1$ in the case of the overall compliance, and at $t_5 + 2$ for the quality parameters or at t_5 for the safety parameters.

The second question that we explore is whether the higher compliance observed is the result of different sets of parameters or contractors audited. We begin by looking separately at parameters in the quality and safety classes. Figure 5 reports for each month the total weight (averaged across all audits in the month) of parameters relating to these two classes. Safety parameters always carry a higher total weight, but their proportion relative to the quality parameters remains rather stable over time. Indeed, the evolution of the monthly average weighted parameters in these two classes reported in Figure 6 confirms a clear upward trend for both of them. As the latter four columns of Table 5 show, breaks in both series occur at $t1$, but the dates of the other breaks are not all identical. This is also related to the speed of adjustments in compliance, as we will discuss below. Before that, we complete the graphical analysis of the composition issue by taking an even more disaggregated view of the performance measures through their grouping into categories. As Figure 7 shows, the number of parameters audited per month is quite heterogenous, but Figure 8 reassures us that the compliance increase over time is quite homogenous across categories.²⁶ Similarly, while there is heterogeneity in how many audits each contractor receives,²⁷ performance increases are rather homogenous across contractors. Figure 9 reports the monthly weighted average compliance by separating contractors into four groups on the basis of the quartile of the distribution of the number of audits they receive and reveals that compliance improves for all of them. This graphical evidence is also confirmed by the results in Table 6 where we report the estimates of linear regressions of the average monthly compliance by contract and supplier on dummy variables for the four break dates detected by the Bai-Perron test (see column (1), panel (b) of Table 5) and other controls (the share of safety parameters among those audited and whether the contract is for public illumination). Column (1) confirms the significance of all four break dates. However, more interestingly, when we gradually augment the set of regressors to include fixed effects for suppliers, contracts and months, we find that the dummy for $t1$ preserves its statistical significance and large magnitude. We take this as a further indication of the saliency of the first announcement $t1$.

²⁶To make the figure easier to interpret, we reported only the 4 most audited categories, but the increase is present essentially in all 12 categories, as also revealed by the summary statistics in Table 3.

²⁷Ranging from nearly 200 audits for the most audited contractor to zero audits for a few contractors.

[INSERT Table 6 and FIGURES 5, 6, 7, 8 and 9 HERE]

A different way to look at the heterogeneity across firms in the announcement response is to exploit the audit randomness. Since at every point in time the choice of which contract to audit is random, we can compare whether the compliance in the very first audit that a firm receives post $t1$ is different between firms audited and those not audited prior to $t1$. In the data, 33 firms receive at least 1 audit post $t1$, with 26 of these audited only post $t1$ and 7 also having already been audited before then. Table 7 reports the results of linear regressions of the share of compliant parameters in the first audit (post $t1$) on an indicator of whether the firm was already audited prior to $t1$ and other controls.²⁸ Both when compliance is unweighted (first three columns) and when it is weighted through the parameters' weight in the RI formula (last three columns), the results indicate a lack of any statistically significant difference between the two supplier groups. Although the small sample size requires interpreting this evidence with caution, this result is also in line with the saliency of the first announcement highlighted by the earlier findings.

[INSERT TABLE 7 APPROXIMATELY HERE]

The final question concerning performance is to what extent suppliers responded to the nuances of the new system. In particular, as sometimes observed in experiments, the mere change in the environment might trigger some response, or, more specifically, as a form of *Hawthorne effect* (or observer effect), suppliers might improve performance once they are aware of the new monitoring.²⁹ Alternatively, behavior might not be improving at all and what we observe may reflect contractors beginning to collude with their monitors. To address these concerns, we resort to a series of probit regressions performed at the level of each individual audited parameter. We estimate the following probit model for the probability of

²⁸The regressions include controls for the share of safety parameters among those audited (in models (2),(3),(5) and (6)), a dummy for whether the contract is for public illumination (in models (2),(3),(5) and (6)), fixed effects for the quarter of the year (in models (2),(3),(5) and (6)) or dummy variables for being past each of the breaks $t2, \dots, t5$ (in models (3) and (6)) and to account for the increasing compliance over time (in models (2),(3),(5) and (6)).

²⁹The *Hawthorne effect* is a change, typically an improvement, in some aspects of behavior in response to the awareness of being observed, see Levitt and List (2011).

the score being 1 (i.e., compliant) on features of parameters, contracts and suppliers:

$$Pr(\textit{compliant}) = \Phi[t + f + \alpha \textit{weight} + \theta \textit{quick} + \gamma_j \sum_{j=2}^{12} \textit{category}_j], \quad (3)$$

where Φ is the normal cdf, *compliant* is the score (0 or 1) taken by the parameter audited, t and f are fixed effects for the year and contractor, *weight* is the weight associated with the parameter, *quick* is a dummy for whether the parameter can be adjusted within one month at a small cost and *category_j* are dummies for the category to which the parameter belongs.

[INSERT TABLE 8 APPROXIMATELY HERE]

We are particularly interested in the coefficient on *weight* as this has the potential to reveal the strategic nature of the suppliers’ responses. Table 8 shows the probit marginal effects for two separate samples: audits held in the period before $t1$ (first four columns), and audits held after then (last four columns). We find that the sign of the coefficient on *weight* changes from negative to positive. Thus, after $t1$, suppliers become more compliant in those parameters with the strongest potential to bolster their RI. This switch in the coefficient sign is evident across all specifications, as we move from a baseline model, controlling only for *weight*, and we expand the model to incorporate parameter, contract and firm features.³⁰

Regarding the other coefficients in Table 8, the one on *quick* is useful to assess the potential for collusion between suppliers and monitors. Indeed, performance might be improving because the repeated interaction allows the parties to learn how to collude under the new system. However, this interpretation of the data would seem less plausible if the improvements were concentrated on those parameters that should be faster to effectively adjust. With the help of expert engineers, we created a dummy variable, *quick*, that is equal to 1 if the transition from a score of not compliant to one of compliant can be reasonably achieved within a one month time frame without incurring extraordinary costs. For instance, examples of parameters with *quick* equal to 1 are those involving the adequacy of “personal

³⁰All estimates in Table 8 are based on the subset of parameters that are audited at least once both before and after $t1$. The results remain qualitatively the same for the post- $t1$ sample if all audits are included.

protection tools” (mostly helmets) or the presence of signs warning of ongoing works nearby. The adequacy of the machineries, instead, is an example of a parameter with *quick* equal to zero. While clearly arbitrary, this dummy variable is helpful to test the reasonableness of the performance response observed in our data. Indeed, the finding that the coefficient on *quick* is positive (and that its significance increases post $t1$) is suggestive of suppliers effectively changing their behavior. This interpretation is further strengthened by what we report below with regard to the behavior in the auctions. However, it is relevant here that while it is impossible to fully rule out the possibility of collusion/corruption, the system of random rotation of auditors and of random selection of the sites to inspect was explicitly meant to curtail these types of risks. Indeed, the Firm never expressed to us concerns about episodes of corruption or collusion during the period our data cover.

B. Announcements Effect on Price

The effect of the announcement on prices is analyzed through a difference-in-differences (DD) strategy. The unit of analysis are the auctions held by the Firm (treated group) and by other CAs (control group). To identify the causal effect of the Firm’s announcement at $t1$ on prices, we estimate the following regression model:

$$D_{ist}^w = a_s + b_t + cX_{ist} + \beta_1(Treatment) + \epsilon_{ist}, \quad (4)$$

where D^w is the winning discount (over the reserve price) and the index i indicates the auction, s the entity awarding the contract and t the year. *Treatment* is a dummy variable equal to one for the contracts awarded by the Firm from $t1$ onward and zero otherwise. The coefficient of interest is β_1 , the effect of the announcement on the winning discount, conditional on fixed effects for the entity awarding the contract (a_s) and time (b_t), and on other covariates (X). The latter set includes characteristics of both the contract - four dummy variables for value of the reserve price and two dummy variables for the work type - and of the contracting authority - its region and whether it is a local authority (municipality, county or region) or not.

In addition to the break at $t1$, we also analyze a second model that incorporates the other

break systematically detected by the Bai-Perron test at $t5 + 1$. This allows us to account for the two differential phases of RI accumulation and stabilization. Thus, we extend the previous model to include a dummy for auctions held from $t5 + 1$ onward, $D_{t>t5+1}$:

$$D_{ist}^w = a_s + b_t + \beta_1 Treatment_{st} + \beta_2 Treatment_{st} * D_{t>t5+1} + D_{t>t5+1} + \gamma X_{ist} + \varepsilon_{ist}, \quad (5)$$

Under this second model, β_1 now measures the effect on the Firm’s awarding discounts past $t1$, but before $t5 + 1$, while β_2 measures the same effect for being after $t5 + 1$, relative to the $t1$ to $t5 + 1$ period. Hence, the effect of the RI accumulation phase is captured by β_1 , while that of the RI stabilization phase is captured by β_2 .

The identification of the key parameters in the two models above crucially hinges on the validity of the auctions in the control group to capture price variations that would have affected the Firm’s auctions absent its reform. Specifically, while Figure 3 suggests that discounts did not change around $t1$, this might be due to simultaneous changes in market conditions. In Figure 4, the comparison of the evolution of winning discounts between the treated and control groups is indeed reassuring of the fact that the similar pre- $t1$ dynamics in the treatment and control auctions make the parallel trends assumption likely to hold. We proceed by first presenting our baseline DD estimates and then exploring their robustness to both identification and inference concerns.

Table 9 presents these baseline estimates for the models in equation (4) (first three columns) and in equation (5) (last three columns). For each model, estimates for three specifications differing on the set of covariates, X , are presented: we first include in X only the constant, then add procurer characteristics, and – finally – also add contract characteristics. We present results for two control groups: the estimates in panel (a) use all contracting authorities, while those in panel (b) use only procurers in central regions. By restricting attention to contracts held in the same geographical area in which the Firm operates, the latter control group has greater potential to capture price variations due to local conditions. However, it also has greater potential for contamination, as we explore below.

[INSERT TABLE 9 APPROXIMATELY HERE]

The results in the top panel of Table 9 show the lack of any price effect when the post- $t1$ period is considered altogether (first three columns). In addition to not being statistically significant, the estimated coefficients are relatively small in magnitude, implying a 4 percent decline in discounts, when compared to the major shift in performance documented above. Interestingly, the estimates change, revealing a rich price dynamic if the post- $t1$ period is divided into a phase pre and post $t5 + 1$. The estimates in the last three columns confirm the visual evidence of Figure 4: discounts initially increase, by about 6 percent of the reserve price, and subsequently decline, by about 15 percent of the reserve price. All estimates are highly statistically significant. This also implies that the discounts after $t5 + 1$ are 9 percent lower than pre $t1$. This effect results from the difference between β_2 and β_1 . The bottom panel of Table 9 shows that, despite the smaller sample, very similar results – in terms of size and significance – are obtained with the second control group of auctions held in central Italy.

The remaining part of this section explores the robustness of Table 9 to three types of concerns. First, the experimentation begun by the Firm might feed back to auctions held by other contracting authorities. This contamination of the control group auctions might occur through changes in either contractors costs or market structure. For instance, since some of the parameters scored involve durable equipment and machineries, the investments made by the Firm’s contractors might also alter their costs in the auctions held by other CAs. Market structure might also change if some of the Firm’s contractors respond by participating more or less in auctions held by other CAs. Since the contamination effect, if present, is likely driven by the presence of common contractors between the Firm and the control group auctions, we address this concern by excluding auctions that are more likely to face this problem. In particular, the top two panels of Table 10 replicate the earlier baseline estimates using two different subsets of control group auctions: in panel (a) we exclude all auctions held in central regions, and in panel (b) we exclude all auctions won by any firm that ever participated in one of the Firm’s auctions. The estimates in both panels are similar to those in the baseline estimates in terms of sign, magnitude and significance.

The second set of robustness checks involves the awarding procedures used. As argued

above, all the auctions considered share many characteristics, but there are nevertheless subtleties of the regulation that might affect outcomes. All procurements in the sample occur via sealed bid, price-only auctions. But across auctions, differences in both auction procedures and awarding methods exist. Auctions where a restricted set of bidders is invited to bid can be used under certain conditions, and indeed this method is used for 87 out of the 330 auctions held by the Firm.³¹ Panel (c) reports estimates excluding these 87 auctions. The results are qualitatively identical to those in the baseline estimates. Regarding the awarding criterion, 42 out of the 330 auctions are awarded via modifications of the lowest price rule. All modifications entail penalizing discounts considered “too good to be true.”³² The awarding criterion is always specified in the call for tenders, so bidders knew these 42 auctions were different and this might have altered their bidding. Indeed, when we exclude these 42 auctions from the sample, we discover some changes relative to the baseline estimates: in the first three columns of panel (d), the sign of the coefficient switches from negative to positive, although it is still not statistically significant and is still relatively small in magnitude (about 3 percent of the reserve price). The source of this change is evident from the estimates in the last three columns: the magnitude of the β_2 estimate more than halves, declining from -15 percent to -6 percent. This suggests that part of the decline in winning discounts observed for the Firm’s auctions after $t5 + 1$ is due to the presence of auctions held under awarding rules that soften price competition. In absolute terms, the magnitude of the β_1 and β_2 estimates is nearly identical, suggesting that the price increase observed after $t5 + 1$ cancels out the initial price saving observed right after $t1$. Overall, these results, as well as the overall lack of significance reported in the first three columns, indicate that the bolstering of performance did not come at the cost of major price increases.

[INSERT TABLE 10 APPROXIMATELY HERE]

³¹The Code refers to these auctions based on invitations as “negotiated procedures.” They are studied in Coviello, Guglielmo and Spagnolo (2016).

³²The Firm used the flexibility given to it by the Code to experiment with three alternatives to the lowest price rule. One method entailed awarding the contract to the contractor with the discount closest to the average discount offered, increased by 20 percent. A second method entailed using a trim mean disregarding the 10 percent of highest and lowest discounts (see Decarolis (2014)) for a description). The third method entailed randomly deciding - after the bids were submitted - whether the criterion to be used was the highest discount or the discount closest to an average of the submitted discounts (either their simple average or a trim mean mentioned above).

A similar conclusion is derived from our third robustness check. In Table 11, we evaluate potential problems with inference by using alternative methods for standard errors. The four columns report 95 percent confidence interval estimates corresponding to models (2), (3), (5) and (6) of Table 9. The rows indicating “PA-Year” report estimates where the clustering is at the year and CA level, as in Table 9. The other rows in the table present two alternatives. The rows “CA” use clustering at the CA level only. As is well known from Bertrand, Duflo and Mullainathan (2004), this can serve to correct for overestimation of the significance of the treatment effect driven by autocorrelation in the data. The table reveals that this correction has no qualitative implications for our results: relative to the baseline estimates, for all models involving both β_1 and β_2 there are no changes, while significance increases for models involving β_1 only. The latter models indicate a negative and significant effect at the 95 percent level. Since the clustering at CA and year is preferable to account for time variation, however, we prefer to rely on our more conservative baseline estimates. The second concern regarding inference is the fact that the level at which the treatment effectively takes place is that of the procurer and we observe only one procurer, the Firm, receiving the treatment. Hence, any shock hitting the Firm at t_1 biases the estimate of β_1 . As argued by Conley and Taber (2011), if the shocks potentially hitting the Firm and the control CAs belong to the same distribution, and if a sufficiently large number of control CAs are observed, valid inference can still be conducted by adjusting the standard errors. Since we have many control CAs, we use the Conley and Taber (2011) method to assess how significance changes relative to our baseline estimates. The “Conley-Taber” rows are indeed different from the baseline ones: in columns (3) and (4), β_1 loses significance, while β_2 loses significance in model (4) when the largest set of control CAs is used. Overall, this indicates that we should be cautious in interpreting the findings in Table 9 about significant and opposite signs of β_1 and β_2 . Hence, as before, a more conservative interpretation is that there are no statistically significant price changes throughout the sample.

[INSERT TABLE 11 APPROXIMATELY HERE]

C. Cost-benefit Analysis

We conclude this section with a back of the envelope cost-benefit analysis of the Firm’s reform.³³ An exhaustive welfare evaluation would require assigning a monetary value to the increased compliance on all quality and safety parameters. Our strategy, instead, is to calculate a lower bound for the potential benefit by focusing on the safety improvements - specifically, on how they reduced the probability of fatal accidents - and to compare that to an upper bound of the reform costs. Construction and maintenance jobs for electricity generation and transmission are among those with the highest incidence of workplace accidents, including deadly accidents. Electricity is widely recognized as a serious workplace hazard, exposing employees to electric shocks, burns, fires, and explosions.³⁴ The occurrence of such accidents has costs for both society and the Firm, and the public ownership of the Firm only increased its management’s concern about these safety risks.

To map the relationship between changes in parameter compliance and the occurrence of fatal accidents, we use the same statistical model used by the Firm’s engineers. This framework, known as *Heinrich’s pyramid*, is often used by practitioners in the context of industrial systems and is based on correlations between accidents of different intensity, estimated from very large (industry-wide) datasets.³⁵ The pyramid entails the following ratios: 1 fatal accident to 10 major accidents, to 30 minor accidents, to 600 material damages, and – finally – to 200,000-300,000 small deviations from safe behaviors. If we assume that each case of non-compliance in the safety parameters audited by the Firm corresponds to a small deviation in the pyramid, we can estimate a lower bound for the policy benefit of €3-5 million per year. This is calculated as follows: in a typical audit, 33.08 parameters are assessed, 85.3 percent of which are related to safety (see Table 3). There are on average 43 contracts a year, with an average duration of 250 working days (see Table 4). Suppose that the same rate of compliance observed across audits applies to every working day, then the

³³We are grateful to Steve Tadelis for encouraging us to complement the paper with this calculation.

³⁴A search among local newspapers revealed that 4 workers had died in the last 15 years while performing jobs procured by the Firm. In the U.S., the Bureau of Labor Statistics recorded a total of 5,587 fatal electrical injuries between 1992 and 2013, an average of 254 fatal electrical injuries each year. Death was due either to electrocution or to fires caused by electricity, see Campbell and Dini (2015).

³⁵See Heinrich (1931), Bird and Germain (1986) and Goodman (2012). See also its usage by modern safety apps: <http://safesiteapp.com/blog/safety/the-safety-triangle-explained/>.

55 percent improvement in parameter compliance discussed in sub-section A above implies a reduction in about 163,000 small deviations per year. Using the 200,000-300,000 figure from the pyramid, this maps into a reduction in the probability of a fatal accident of 0.54-0.82 per year. Finally, considering an average of 4 workers on the worksite per day and taking the lowest bound of the OECD (2012) estimates of the “value of a statistical life” of €1.62 million per life saved,³⁶ the estimated benefit ranges between 3.5 and 5.3 million euro per year.³⁷

Regarding the cost, the worst case scenario estimate from sub-section B above implies that, after $t_5 + 1$, the winning price becomes 9.07 percent higher relative to the pre- t_1 period (see column 5, panel (a) of Table 9). Thus, considering an average contract value of €0.53 million and 43 contracts per year, the cost of the policy would be €2.4 million. Hence, we can conclude that the benefits from adoption exceed the costs. We shall also remark that our approach is quite conservative as for costs we have used an upper bound/worst case estimate, while for benefits we have employed the lowest OECD estimate of the value of a statistical life. Using the upper bound of the OECD estimate (€5.3 million), the benefits would be in the range of €11.55 -17.33. Furthermore, our benefit calculation excludes all the additional savings accruing from both reductions in non-fatal accidents associated with better safety practices and all improvements in quality associated with increased compliance on the quality parameters.

VI Discussion: Adverse Selection and Moral Hazard

The three motives offered by the literature to understand why competitive procedures such as lowest price auctions induce a trade-off between price and performance are adverse selection, moral hazard and the winner’s curse fallacy. The latter refers to bidders’ inability to properly assess project costs at the time of bidding. In our setting, this non-equilibrium phenomenon

³⁶The number of workers present on the worksite was estimated for us by the same expert engineers who estimated the variable *quick* described earlier. The OECD (2012) values are converted to 2007 nominal euro.

³⁷This range is not our interval estimate, but the result of using the two bounds of 200,000 and 300,000 small deviations.

is unlikely to play a major role, as we observe experienced contractors repeatedly bidding in auctions for relatively simple contracts. It is interesting, however, to understand the extent to which selection and moral hazard can explain our findings, as they can have different implications for how best to design systems to integrate past performance in procurement.³⁸

The Firm’s experiment might have induced responses in both of them: performance improvements can derive from either more effort in the execution stage by contractors, or by a better selection of contractors, or by a combination of both. Indeed, Figure 10 presents evidence consistent with both by showing how the cdf of winning bids in the Firm’s auctions evolves between those held before $t1$, after $t5 + 1$, and in between these two periods. The noteworthy aspect is the disappearance post $t1$ of right tail discounts, representing discounts of one third or more relative to the reserve price. It is precisely this type of abnormally high discounts that procurers worry will be associated with poor performance. Since replicating Figure 10 for those firms bidding both before and after $t1$ leads to a similar finding of a disappearing right tail after $t1$, we can conclude that the altered bidding behavior of these suppliers is compatible with the presence of moral hazard in contractual performance.

[INSERT FIGURE 10 APPROXIMATELY HERE]

More specifically, it is useful to describe how a stylized model of a lowest price, sealed bid procurement auction with moral hazard can rationalize our findings. Equilibrium bids should depend on two elements: production costs, $C(e)$, which are an increasing function of the effort e that the bidder plans to put in place in the execution stage; and a strategic markup, $M(n)$, which is inversely proportional to the number of competitors, n . Prior to $t1$, each auction exists in isolation - the outcome of an auction does not matter for future auctions.³⁹ Thus, bidders will choose low effort levels to reduce their cost and maximize their profits. From $t1$ onward, however, even if the awarding rule remains the lowest price,

³⁸For instance, consider the length of the memory of the RI (i.e, how far back should the RI look). This likely needs to be long, possibly infinite, if screening is the concern, but short if moral hazard is predominant; see Elul and Gottardi (2015), and also Kovbasyuk and Spagnolo (2016) who show that optimal memory might differ for positive and negative ratings.

³⁹This assumes that there are no links through, for instance, capacity constraints. This is likely to be a good approximation, since the institutional environment allows for an extensive use of subcontracts that can relax capacity constraint. See Branzoli and Decarolis (2015) for subcontracting and capacity constraints.

the game played by the contractors becomes dynamic: winning an auction can imply being audited and, hence, an opportunity to modify one's own RI while reducing one's rivals' chances of improving their RI. This likely implies changes to both components of the bid relative to the pre- $t1$ case: if better compliance requires more effort, then the optimal $C(e)$ will likely be higher. Moreover, the strategic markup now depends not only on n , but also on the distribution of RI across bidders.⁴⁰ Finally, the bid now also incorporates a third element: the continuation value associated with the evolution of the RI. Indeed, winning today and earning a good RI is expected to produce savings in the stream of future auctions, once the scoring rule auction is introduced. This continuation value increases the value of winning today and, hence, balances increases in production costs. Clearly, the relevance of the continuation value depends on how many auctions suppliers perceive they will be able to use their good RI for.

It is not a priori obvious how increasing the RI weight in the scoring auction affects the outcomes. An increase in this weight helps with the moral hazard problem as it bolsters the benefits of more effort. However, the effect on bidding during the phase before the introduction of the scoring rule is ambiguous. There are two effects, which in a sense correspond to a marginal and an inframarginal effect (alternatively, the intensive and extensive margin). First, winning lowest price auctions gives bidders the opportunity to prove themselves and thus increase their RI (marginal effect). Second, if the implementation of the scoring rule auction is sufficiently delayed that all bidders have the potential to earn a good RI, then symmetric competition in the (future) scoring rule auction will imply that many of the rents from a good RI will be competed out, to the point that winning in the (present) lowest price auctions becomes less attractive (the inframarginal effect).

Thus, an explanation for the patterns observed in the data is that, right after $t1$, the increase in $C(e)$ was dominated by the changes in the strategic markup and the continuation value. After contractors accumulated a good RI, however, the value of winning an auction in the pre-scoring rule period declines as obtaining positive audit reports cannot offer a

⁴⁰That is, even if before $t1$ the environment could be characterized as a symmetric auction with bidders being ex ante identical in terms of costs, after $t1$ firms became asymmetric in terms of their RI stock. This asymmetry can potentially cause changes in the size of the equilibrium markups, for instance by making bidders with lower (or no) RI more willing to shade less their true cost.

competitive edge over rivals. Thus, in this phase the increased production cost dominates.⁴¹

In the data, however, the effects of the new system concerned not only bidding and performance, but also participation choices. Indeed, while the summary statistics show that the number of bids submitted remains stable and approximately equal to 10 both before and after $t1$, the set of bidders changed in the Firm's auctions: while there are 34 suppliers placing at least one bid both before and after $t1$, there are other 36 suppliers who place at least one bid before $t1$, but no bid afterwards. We refer to the latter group of firms as "*exitors*" and to the former as "*stayers*." There are also 3 new entrants placing bids only after $t1$, but never before then. This implies that the average number of bids placed per bidder doubles; from 0.14 (i.e., 10/70) to 0.27 (i.e., 10/37). This increased participation is due to the *stayers*, not to unusually high bidding frequencies for the 3 new entrants, and is likely driven by the same incentive to earn RI that we discussed when analyzing the evidence on winning bids. As regards *exitors* and new entrants, however, their mere presence potentially indicates that the experiment might have also triggered some selection effects.

If we focus on *exitors*, however, the data provides only weak evidence of possible selection effects.⁴² In particular, Figure 11 shows the timing of the exits does not seem clearly linked to $t1$. This figure reports the last date at which each of the *exitors* (represented by the numerical identifiers on the vertical axis) placed a bid. The smooth path of exits indicates more of a gradual process than a sharp drop at $t1$.

[INSERT FIGURE 11 APPROXIMATELY HERE]

Furthermore, as illustrated by Figure 12, if we compare the cdf of winning bids by

⁴¹The intuition for this latter effect is that higher effort pre-scoring rule improves a bidder's expected payoff once the scoring rule becomes effective. However, in a symmetric equilibrium, all bidders win the same number of lowest price auctions and assign the same value to effort. This implies that the equilibrium payoff once the scoring rule is implemented is independent of the weight it assigns to the RI relative to price; the only effect of increasing this weight is thus to increase effort early on. But an increase in this effort decreases the expected payoff from winning an auction pre scoring rule. Finally, this leads bidders uniformly to bid less aggressively in the pre scoring rule period. This result bears some resemblance to models where the strategic effect of an exogenous change (in this case, the expected change in the weight assigned to the RI from zero to 25 percent) more than outweighs any positive direct effect, to the point that equilibrium payoffs are decreasing in the RI weight (see Cabral and Villas-Boas (2005)). We are grateful to Luis Cabral for helping us to elucidate this unintuitive and important element of the strategic environment.

⁴²For the 3 new entrants, the type of analysis performed below for the *exitors* cannot be replicated as only one of them could be matched to the firm registry described below.

both *exitors* and *stayers* (in the pre- t_1 auctions), we do not observe significant differences. Finally, even in terms of characteristics, *exitors* do not seem to be substantially different from *stayers*. Table 12 reports summary statistics for the subset of *exitors* and *stayers* that we could match to the Infocamere database, the Italian firm registry.⁴³ Statistics for the *exitors* are reported in the first four columns of panel (a), followed by statistics for the *stayers* in the following four columns. Along most dimensions, *exitors* are smaller than *stayers*; this is the case for revenues, profits and capital. The average number of employees is also lower, but in this case the median is nearly identical. For both groups, the wide variation in characteristics among firms means that the differences in the averages are not statistically significant and it is not obvious to how to interpret the results. Thus, to benchmark them we present in panel (b) the analogous statistics obtained for the suppliers active in the auctions of the multi-utility company of the city of Turin. This is the multi-utility company that awards most contracts within the DD control group. Analogously to what was done for the Firm, we partition its suppliers into those bidding both before and after t_1 (*stayers*) and those bidding only before t_1 (*exitors*). The comparison of the two groups leads to similar conclusions to those found for the Firm’s suppliers: the average revenues, profits and capitals are higher among *stayers*. But the data are again characterized by many extreme observations and the result is reversed for revenues and profits when looking at the median.

[INSERT FIGURE 12 APPROXIMATELY HERE]

[INSERT TABLE 12 APPROXIMATELY HERE]

We conclude that, overall, there is no strong evidence that the pool of *exitors* in the Firm’s auctions is selected in any particular way relative to the typical exit behavior in the market. Thus, the effects that we uncovered in the earlier section are likely driven to a large extent by changes in the behavior of the Firm’s contractors. Perhaps, this is not surprising given that the slow switch to scoring rule auctions allowed most contractors to undertake the steps needed to improve their contractual performance. In turn, this might also be part of the reason why competition remained high in the Firm’s auctions, thus contributing to limit

⁴³The registry covers nearly all Italian firms, for a description see Conley and Decarolis (2016).

price increases. In the light of this conclusion, it is interesting to note that most of the policy debate referenced earlier in this paper is centered around issues of supplier selection and most often ignoring the potential disciplining effects of reputational mechanisms on moral hazard.

VII Conclusions

This paper has studied the merits of using past performance to spur greater efforts from contractors when executing public works. The evaluation of the evidence from an experiment undertaken by a large multi-utility company has shown strong improvements in performance after the Firm announced its intention to use past performance scores to award future contracts. To some extent this may resemble the well-known *Hawthorne effect*. However, contrary to the *Hawthorne effect*, the improvement was not short-lived, even if we consider that the contractors could have stopped trusting the Firm over the delayed implementation of the new awarding rule, and that it was easier for contractors to improve their score when the starting point was lower rather than later, when the marginal cost to improve became higher. Regarding prices, we find some evidence of an initial fall in prices followed by a moderate price increase. The overall price increase appears rather small when compared to the substantial improvement in performance.

Although the empirical evidence in this paper is suggestive of the benefits from implementing reputation mechanisms in public procurement, an extension of our analysis would involve exploring welfare implications. In particular, the back of the envelope calculations that we performed could be complemented by more accurate estimates of the savings from both reductions in non-fatal accidents and improvements in quality. Furthermore, although several different mechanisms might explain how suppliers increased their quality and safety performance, it is interesting to note that the explanation offered by the management of the Firm is that most of the gains came from improvements in management practices within contractors. Thanks to new data on management practices collected in the last ten years through the World Management Survey,⁴⁴ there has been increased attention on the role of

⁴⁴See: <http://worldmanagementsurvey.org>.

management in explaining productivity differences (Bloom et al. (2014)). In this respect, exploring the details of the managerial changes implemented by the suppliers would be useful to understand how (the announcement of) new procurement rules triggered an improvement in management. Regarding this, it is also important to highlight that, while we have stressed the public procurement implications of our analysis, our findings are also relevant to private procurement practices, where the use of vendor rating systems is widespread but little is known about their effectiveness.

Finally, once the merits of this kind of reputation mechanism in improving contractor performance are proven, many aspects remain open and give room for future research; for example, how to optimize the parameters' weight, how to discipline the rating for new entrants, how to structure the weights in the awarding criteria, and how to choose the optimal “memory” of the indicator (i.e. how long the window of time over which the RI is calculated should be and how heavily older information should be discounted). Even the ideal speed at which the switch to a reputation system should occur is an interesting, but little studied problem.

References

- Athey, Susan, Dominic Coey, and Jonathan Levin.** 2013. “Set-Asides and Subsidies in Auctions.” *American Economic Journal: Microeconomics*, 5(1): 1–27.
- Bajari, Patrick, and Steven Tadelis.** 2001. “Incentives versus Transaction Costs: A Theory of Procurement Contracts.” *RAND Journal of Economics*, 32(3): 387–407.
- Bajari, Patrick, Robert S. McMillan, and Steven Tadelis.** 2009. “Auctions Versus Negotiations in Procurement: An Empirical Analysis.” *Journal of Law, Economics and Organization*, 25(2): 372–399.
- Bandiera, Oriana, Andrea Prat, and Tommaso Valletti.** 2009. “Active and Passive Waste in Government Spending: Evidence from a Policy Experiment.” *The American Economic Review*, 99(4): 1278–1308.

- Bandiera, Oriana, Iwan Barankay, and Imran Rasul.** 2011. “Field Experiments with Firms.” *Journal of Economic Perspectives*, 25(3): 63–82.
- Banerjee, Abhijit V., and Esther Duflo.** 2000. “Remedying Education: Evidence from Two Randomized Experiments in India.” *The Quarterly Journal of Economics*, 122(3): 1235–1264.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan.** 2004. “How Much Should We Trust Differences-in-Differences Estimates?” *The Quarterly Journal of Economics*, 119(1): 249–275.
- Beuve, Jean, and Lisa Chever.** 2014. “Quality of Outsourced Services, Opportunism and Contract Design. Evidence from Cleaning Contracts.” *mimeo*.
- Bird, Frank E, and George L Germain.** 1986. *Practical loss control leadership*. Loganville, Ga. Institute Publications.
- Bloom, Nicholas, Renata Lemos, Raffaella Sadun, Daniela Scur, and John Van Reenen.** 2014. “The New Empirical Economics Of Management.” *Journal of the European Economic Association*, 12(4): 835–876.
- Branzoli, Nicola, and Francesco Decarolis.** 2015. “Entry and Subcontracting in Public Procurement Auctions.” *Management Science*, 61(12): 2945 – 2962.
- Bulow, Jeremy, and Paul Klemperer.** 1996. “Auctions Versus Negotiations.” *American Economic Review*, 86(1): 180–194.
- Bulow, Jeremy, and Paul Klemperer.** 2009. “Why Do Sellers (Usually) Prefer Auctions?” *American Economic Review*, 99(4): 1544–1575.
- Butler, Jeffrey V., Enrica Carbone, Pierluigi Conzo, and Giancarlo Spagnolo.** 2014. “Reputation and Entry in Procurement.” *Working paper*.
- Cabral, Luis M. B., and Miguel Villas-Boas.** 2005. “Bertrand Supertraps.” *Management Science*, 51(4): 599–613.

- Calzolari, Giacomo, and Giancarlo Spagnolo.** 2009. "Relational Contracts and Competitive Screening." CEPR Discussion Papers 7434.
- Campbell, Richard B., and David A. Dini.** 2015. "Occupational Injuries from Electrical Shock and Arc Flash Events." Fire Protection Research Foundation Final Report 1.
- Chassang, Sylvain, and Juan Ortner.** 2016. "Collusion in Auctions with Constrained Bids: Theory and Evidence from Public Procurement." *Working Paper*.
- Conley, Timothy G., and Christopher R. Taber.** 2011. "Inference with Difference in Differences with a Small Number of Policy Changes." *The Review of Economics and Statistics*, 93(1): 113–125.
- Conley, Timothy G., and Francesco Decarolis.** 2016. "Detecting Bidders Groups in Collusive Auctions." *American Economic Journal: Microeconomics*, 8(2): 1–38.
- Coviello, Decio, Andrea Guglielmo, and Giancarlo Spagnolo.** 2016. "The Effect of Discretion on Procurement Performance." *Management Science*, Forthcoming.
- Decarolis, Francesco.** 2014. "Awarding Price, Contract Performance and Bids Screening: Evidence from Procurement Auctions." *American Economic Journal: Applied Economics*, 6(1): 108–132.
- Djankov, Simeon, Rafael La Porta, Florencio Lopez-De-Silanes, and Andrei Shleifer.** 2003. "Courts." *The Quarterly Journal of Economics*, 118(2): 453–517.
- Doni, Nicola.** 2006. "The Importance Of Reputation In Awarding Public Contracts." *Annals of Public and Cooperative Economics*, 77(4): 401–429.
- Elul, Ronel, and Piero Gottardi.** 2015. "Bankruptcy: Is It Enough to Forgive or Must We Also Forget?" *American Economic Journal: Microeconomics*, 7(4): 294–338.
- Gil, Ricard, and Justin Marion.** 2012. "Self-Enforcing Agreements and Relational Contracting: Evidence from California Highway Procurement." *Journal of Law, Economics, and Organization*.

- Goodman, William M.** 2012. “Measuring Changes in the Distribution of Incident-Outcome Severities: A Tool for Safety Management.” *Case Studies in Business, Industry and Government Statistics*, 5(1): 32–43.
- Gordon, Daniel I., and Gabriella M. Racca.** 2014. “Integrity Challenges in the EU and U.S. Procurement Systems.” *Integrity and Efficiency in Sustainable Public Contracts. Corruption, Conflict of Interest, Favoritism and Inclusion of Non-Economic Criteria in Public Contracts*, Gabriella M. Racca & Christopher R. Yukins, eds., Bruylant.
- Heinrich, Herbert William.** 1931. *Industrial accident prevention: a scientific approach*. McGraw-Hill.
- Hyytinen, Ari, Sofia Lundberg, and Otto Toivanen.** 2018. “Design of public procurement auctions: evidence from cleaning contracts.” *RAND Journal of Economics*, 49(2): 398–426.
- Kang, Karam, and Robert A. Miller.** 2016. “Winning by Default: Why is There So Little Competition in Government Procurement?” Mimeo.
- Kawai, Kei, and Jun Nakabayashi.** 2016. “Detecting Large-Scale Collusion in Procurement Auctions.” Mimeo.
- Kelman, Steven.** 1990. *Procurement and Public Management: The Fear of Discretion and the Quality of Government Performance*. AEI Studies.
- Kim, In-Gyu.** 1998. “A model of selective tendering: Does bidding competition deter opportunism by contractors?” *The Quarterly Review of Economics and Finance*, 38(4).
- Klein, Tobias J., Christian Lambertz, and Konrad O. Stahl.** 2016. “Market Transparency, Adverse Selection, and Moral Hazard.” *Journal of Political Economy*, Forthcom.
- Kovbasyuk, Sergei, and Giancarlo Spagnolo.** 2016. “Memory and Markets.” *mimeo*.
- Lazear, Edward P.** 2000. “Performance Pay and Productivity.” *American Economic Review*, 90(5): 1346–1361.

- Levitt, Steven D., and John A. List.** 2011. “Was There Really a Hawthorne Effect at the Hawthorne Plant? An Analysis of the Original Illumination Experiments.” *American Economic Journal: Applied Economics*, 3: 224–238.
- Lewis-Faupel, Sean, Yusuf Neggers, Benjamin A. Olken, and Rohini Pande.** 2016. “Can Electronic Procurement Improve Infrastructure Provision? Evidence from Public Works in India and Indonesia.” *American Economic Journal: Economic Policy*, Forthcoming.
- Lewis, Gregory, and Patrick Bajari.** 2011. “Procurement Contracting with Time Incentives: Theory and Evidence.” *The Quarterly Journal of Economics*, 126(3): 1173–1211.
- Lewis, Gregory, and Patrick Bajari.** 2014. “Moral Hazard, Incentive Contracts, and Risk: Evidence from Procurement.” *The Review of Economic Studies*, forthcoming.
- Liebman, Jeffrey B., and Neale Mahoney.** 2016. “Do Expiring Budgets Lead to Wasteful Year-End Spending? Evidence from Federal Procurement.” *Mimeo*.
- List, John A., and David Reiley.** 2008. “The New Palgrave Dictionary of Economics,.” In *Field Experiments*, ed. Steven N. Durlauf and Lawrence E. Blume. Palgrave Macmillan.
- Manelli, Alejandro M, and Daniel R Vincent.** 1995. “Optimal Procurement Mechanisms.” *Econometrica*, 63(3): 591–620.
- Manuel, Kate M.** 2015. “Evaluating the “Past Performance” of Federal Contractors: Legal Requirements and Issues.” *Congressional Research Service*, R41562.
- Marion, Justin.** 2016. “Affirmative Action Exemptions and Capacity Constrained Firms.” *Mimeo*.
- OECD.** 2012. *Mortality Risk Valuation in Environment, Health and Transport Policies*. OECD Publishing.
- Olken, Benjamin A.** 2007. “Monitoring Corruption: Evidence from a Field Experiment in Indonesia.” *Journal of Political Economy*, 115(2).

- Saussier, Stephane, and Jean Tirole.** 2015. “Strengthening the Efficiency of Public Procurement.” *Les notes du conseil d’analyse économique*, 22.
- Spulber, Daniel F.** 1990. “Auctions and Contract Enforcement.” *Journal of Law, Economics and Organization*, 6(2): 325–44.
- Tadelis, Steven.** 2016. “The Economics of Reputation and Feedback Systems in E-Commerce Marketplaces.” *IEEE Internet Computing*, 20(1): 12–19.

Tables and Figures

Table 1: Comparison with U.S. Multi-Utility Providers

	Firm	LADWP	ComEd	BGE	PECO
Total Employees (000)	5.0	9.4	6.8	3.3	2.6
Power Customers (mln)	1.6	1.4	3.8	1.3	1.6
Power Grid (000/miles)	19	14	90	26	14
Total Turnover (bln/\$)	3.2	4.4	4.9	3.1	3.0
	(2.1)	(3.3)			
Power Supply (TWh)	11	26	86*	29*	36*
Power Grid	206	318	2,400	500	475
Maintenance/Upgrade (mln/\$)					

The Firm and LADWP figures on employees and turnover include the water business too. BGE and PECO figures on employees and turnover include the gas business too. All values are for 2015. Values with a * symbol are estimates: the supply is estimated proportionally to the customers out of the total supply of all Exelon subsidiaries (195TWh). For the total Turnover (bln/\$), the values in par thesis refer to power only.

Table 2: Reputation Index Components

Class	Category	Parameters	
		Number	Avg. Weight
Safety	Equipment and machinery	5	8.4
	Documentation	9	6.9
	Works execution	8	8.8
	Personnel	4	9.3
	Works site regularity	10	8.2
	Works site safety	10	9.4
	H.T. works site controls	5	8.8
Quality	Works on joints	19	5.7
	Customer relationship mgnt	3	7.3
	Air works	25	6.7
	Underground works	25	6.0
	Works on transformer station	13	6.2

The table reports the two classes and 12 categories in which the 136 parameters are subdivided. For each category, the first number reported is the number of parameters in that category, while the second is the average RI weight across these parameters.

Table 3: Summary Statistics - Audit Data

Class	Category	Number of observations	PI Share	Share Compliant Pre t1	t1-5	Post t5
Safety						
	Equipment and machinery	6,761	.238	.701	.899	.942
	Documentation	9,407	.212	.333	.541	.721
	Works execution	6,311	.231	.187	.753	.88
	Personnel	3,875	.228	.319	.602	.684
	Works site regularity	13,615	.204	.103	.449	.686
	Works site safety	14,730	.231	.306	.614	.813
	H.T. works site controls	351	0	-	.62	.801
Quality						
	Works on joints	1,471	.206	1	.939	.976
	Customer relationship mgnt	69	0	1	.75	1
	Air works	107	.009	-	.963	1
	Underground works	7,607	.112	.398	.634	.715
	Works on transformer station	233	0	1	.993	1
Sample size:		64,537	13,274	1,997	31,478	31,062

This table reports summary statistics for the Audits data. The 64,537 scores assigned to each parameter audited were produced through 1,951 individual audits that took place on 187 distinct contracts.

Table 4: Summary Statistics - Auctions Cross Section

<i>Pre t1</i>						
	(1)			(2)		
	Treated			Control		
	Mean	SD	N	Mean	SD	N
Winning Discount	22.25	10.39	190	21.34	10.22	2039
Winning Bid	528.1	466.0	190	446.7	521.6	2039
Length(days)	399.1	175.6	174	327.4	339.7	1798
Num. Bids	.	.	0	.	.	0
Public Illumination	0.179	0.384	190	0.264	0.441	2039
Central Region	1	0	190	0.205	0.404	2039
Municipal Firm	1	0	190	0.393	0.489	2039

<i>Post t1</i>						
	(1)			(2)		
	Treated			Control		
	Mean	SD	N	Mean	SD	N
Winning Discount	23.88	9.735	296	24.81	12.04	4417
Winning Bid	712.8	530.6	296	406.4	493.1	4417
Length(days)	404.5	167.7	257	395.4	920.0	3238
Num. Bids	12.43	4.786	174	.	.	0
Public Illumination	0.240	0.428	296	0.261	0.439	4417
Central Region	1	0	296	0.195	0.397	4417
Municipal Firm	1	0	296	0.410	0.492	4417

This table reports summary statistics for the Auctions data. The “Treated” sample consists of the Firm auctions, while the “Control” sample consists of auctions held by all other CAs. The upper panel contains auctions held before $t1$, while the lower panel contains auctions held at or after $t1$. The definition of the variables is as follows: *Winning Discount* is the discount (over the reserve price) offered by the winning supplier, *Winning Bid* is the price bid by the winning supplier, *Length* is the contractual duration of the contract in days (a contractual duration of 1 year corresponds to 250 working days), *Num. Bids* is the number of bids submitted, *Public Illumination* is a dummy equal to 1 if the contract type is classified by the Firm as public illumination and zero if it is classified as work on electrical substations, *Central Region* is a dummy equal to 1 if the CA is located in one of Italy’s Center regions and zero otherwise and *Municipal Firm* is a dummy equal to 1 if the CA is a multi-utility company that is (at least partially) owned by the municipality in which it operates.

Table 5: Chow and Bai-Perron Tests

(a) Chow Tests						
	<i>Weighted Compliance</i>		<i>Quality</i>		<i>Safety</i>	
	1 break at t1	5 breaks at t1-5	1 break at t1	5 breaks at t1-5	1 break at t1	5 breaks at t1-5
F-statistic	27.80 (0.0000)	117.20 (0.000)	68.48 (0.000)	21.77 (0.000)	25.61 (0.000)	148.84 (0.000)

(b) Bai-Perron Tests						
	<i>Weighted Compliance</i>		<i>Quality</i>		<i>Safety</i>	
	F-stat breaks	5 unknown breaks	F-stat breaks	5 unknown breaks	F-stat breaks	5 unknown breaks
Number of breaks	4	5	2	5	4	5
Dates of the brakes:						
Date 1	t1	t1	t1	t1	t1	t1
Date 2	t2	t2	t3+2	t3+2	t2	t2
Date 3	t3+1	t3+1	-	t4+2	t3+1	t3+1
Date 4	t5+1	t5+1	-	t5+2	t5+7	t5
Date 5	-	t5+7	-	t5+5	-	t5+7

The table reports the results of Chow (top panel) and Bai-Perron (bottom panel) tests. The variable is the monthly weighted average compliance, measured on all audited parameters (first two columns) or on the subset of quality parameters (next two columns) or safety parameters (latter two columns). For the Bai-Perron test, the criterion used is that of sequential F-statistic determined breaks. Identical results are obtained when using the criterion of the significant F-statistic largest breaks.

Table 6: Average Compliance by Firm, Contract and Month

	(1)	(2)	(3)	(4)
t1	0.200*** (0.044)	0.195*** (0.042)	0.172*** (0.049)	0.456** (0.200)
t2	0.065** (0.031)	0.060** (0.030)	0.058** (0.029)	0.082 (0.132)
t3+1	0.122*** (0.024)	0.115*** (0.023)	0.138*** (0.023)	-0.107 (0.109)
t5+1	0.082*** (0.018)	0.067*** (0.017)	0.055*** (0.020)	0.013 (0.068)
Safety share	0.164*** (0.046)	0.125*** (0.046)	0.038 (0.048)	0.042 (0.049)
PI-type	-0.051*** (0.017)	-0.054** (0.024)	-0.039 (0.063)	-0.027 (0.063)
Firm Fixed Effects	No	Yes	Yes	Yes
Contract Fixed Effects	No	No	Yes	Yes
Month Fixed Effects	No	No	No	Yes
N	963	963	963	963

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The sample is constructed from the Audit data. The dependent variable is the average compliance (weighted with the RI parameters' weights) for each firm-contract-month triplet. Regarding the regressors, $t1$ is a dummy variable equal to 1 from $t1$ onward and zero before then. $t2, t3 + 1, t5 + 1$ are constructed analogously. *Safety share* is the (weighted) average share of safety parameters and *PI-type* is the proportion of contracts classified as public illumination, both calculated among those parameters audited in the firm-contract-month triplet.

Table 7: Compliance at the First Audit

	Unweighted			Weighted		
	(1)	(2)	(3)	(4)	(5)	(6)
Audited before t1	-0.205 (0.138)	-0.220 (0.135)	-0.219 (0.133)	-0.208 (0.141)	-0.223 (0.137)	-0.222 (0.135)
Safety share		0.411 (0.356)	0.464 (0.335)		0.461 (0.362)	0.508 (0.339)
PI-type		-0.160 (0.124)	-0.155 (0.124)		-0.161 (0.126)	-0.155 (0.125)
N	33	33	33	33	33	33
Quarter FE	Yes	Yes	No	Yes	Yes	No
Breaks	No	No	Yes	No	No	Yes

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The sample consists of the 33 firms audited at least once post $t1$. Out of these 33 firms, 26 firms were never audited pre $t1$ and 7 firms had already been audited before then. The table reports OLS coefficients for regressions of the share of complaint parameters during the first audit on a dummy for whether the firm was already audited pre $t1$, and other controls. The set of controls changes across columns and includes combinations of: quarter fixed effects, the share of safety parameters in the first audit and dummy variables for both the dates of the breaks and whether the contract audited is for public illumination. In the first three columns, the compliance measure is unweighted, while in the latter three it is weighted by the parameters weights in the RI formula.

Table 8: Probability of Compliant Parameter

	Pre-announcement				Post-announcement			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Weight	-0.026*** (0.005)	-0.024*** (0.007)	-0.024*** (0.007)	-0.025*** (0.007)	0.011*** (0.001)	0.013*** (0.001)	0.013*** (0.001)	0.013*** (0.001)
Quick		0.077* (0.036)	0.077* (0.036)	0.074* (0.036)		0.066*** (0.006)	0.065*** (0.006)	0.066*** (0.006)
C2-Documentation		-0.412*** (0.053)	-0.412*** (0.053)	-0.440*** (0.055)		-0.284*** (0.010)	-0.268*** (0.010)	-0.270*** (0.010)
C3-Works Execution		-0.518*** (0.062)	-0.518*** (0.062)	-0.523*** (0.064)		-0.189*** (0.010)	-0.189*** (0.010)	-0.192*** (0.010)
C7-Underground works		-0.302*** (0.051)	-0.302*** (0.051)	-0.296*** (0.052)		-0.291*** (0.009)	-0.288*** (0.009)	-0.286*** (0.009)
C9-Personnel		-0.308*** (0.069)	-0.308*** (0.069)	-0.332*** (0.069)		-0.349*** (0.011)	-0.359*** (0.011)	-0.365*** (0.011)
C10-Works site regularity		-0.673*** (0.054)	-0.673*** (0.054)	-0.680*** (0.056)		-0.449*** (0.009)	-0.443*** (0.009)	-0.441*** (0.009)
C11-Works site safety		-0.381*** (0.056)	-0.381*** (0.056)	-0.405*** (0.057)		-0.272*** (0.010)	-0.272*** (0.010)	-0.275*** (0.010)
Year Fixed Effects	No	No	Yes	Yes	No	No	Yes	Yes
Firm Fixed Effects	No	No	No	Yes	No	No	No	Yes
N	1,702	1,374	1,374	1,374	56,085	44,653	44,653	44,653

This table reports the marginal effects of probit regressions. The dependent variable is the score on the parameter: 1 if compliant and 0 if not compliant. The first four columns regard the subsample of scores assigned in the audits held before $t1$, while the latter four columns regard audits that occurred after $t1$.

Table 9: Baseline Estimates

Panel (a): All Contracting Authorities						
	(1)	(2)	(3)	(4)	(5)	(6)
β_1	-2.48 (3.22)	-2.49 (3.22)	-2.66 (3.14)	4.78** (1.81)	4.78** (1.81)	4.35* (1.71)
β_2				-8.68** (2.67)	-8.70** (2.67)	-8.38** (2.59)
N	6942	6942	6942	6942	6942	6942
R ²	0.37	0.37	0.38	0.38	0.38	0.38
Panel (b): Contracting Authorities in Central Regions						
	(1)	(2)	(3)	(4)	(5)	(6)
β_1	-1.65 (2.92)	-1.66 (2.92)	-2.02 (2.87)	5.42*** (1.46)	5.42*** (1.46)	4.55** (1.47)
β_2				-8.89*** (2.68)	-8.89*** (2.69)	-8.24** (2.57)
N	1767	1767	1767	1767	1767	1767
R ²	0.25	0.25	0.26	0.26	0.26	0.27
PA-type FE	No	Yes	Yes	No	Yes	Yes
Contract Chars FE	No	No	Yes	No	No	Yes

This table contains the baseline DD estimates. Standard errors clusters by year and CA are reported in parentheses. The two panels differ for the control group used: panel (a) uses all available auctions, while panel (b) uses only auctions held by contracting authorities in central regions. In each panel, the first three columns report estimates for the model in equation (4), while the last three columns report estimates for the model in equation (5). For each model, estimates for three specifications differing on the set of covariates, X , are presented: we first include in X only the constant, then add procurer characteristics and, finally, also add contract characteristics.

Table 10: Robustness Checks: Contamination and Sample Effects

Panel (a): No Contracting Authorities in Central Regions						
	(1)	(2)	(3)	(4)	(5)	(6)
β_1	-2.52 (3.05)	-2.55 (3.05)	-2.75 (2.98)	4.84** (1.62)	4.82** (1.63)	4.38** (1.53)
β_2				-8.82*** (2.57)	-8.83*** (2.57)	-8.54*** (2.49)
N	5661	5661	5661	5661	5661	5661
R ²	0.40	0.40	0.41	0.40	0.40	0.41
Panel (b): No Auctions with Common Winners						
	(1)	(2)	(3)	(4)	(5)	(6)
β_1	-4.06 (4.51)	-4.13 (4.51)	-4.01 (4.31)	5.69*** (1.00)	5.63*** (1.02)	5.33*** (1.14)
β_2				-14.79*** (3.47)	-14.80*** (3.46)	-14.15*** (3.31)
N	3585	3585	3585	3585	3585	3585
R ²	0.41	0.42	0.42	0.43	0.43	0.43
Panel (c): No Restrictions to Open Competition						
	(1)	(2)	(3)	(4)	(5)	(6)
β_1	-4.02 (4.80)	-4.04 (4.80)	-4.55 (4.72)	4.86*** (1.42)	4.85*** (1.42)	4.69*** (1.40)
β_2				-12.63** (4.30)	-12.64** (4.30)	-13.18*** (3.94)
N	6609	6609	6609	6609	6609	6609
R ²	0.38	0.38	0.39	0.39	0.39	0.39
Panel (d): No Variations to the Lowest Price Criterion						
	(1)	(2)	(3)	(4)	(5)	(6)
β_1	3.70 (2.51)	3.62 (2.53)	3.43 (2.38)	6.52*** (1.36)	6.46*** (1.38)	6.02*** (1.39)
β_2				-6.02*** (0.78)	-6.04*** (0.78)	-5.50*** (0.80)
N	3531	3531	3531	3531	3531	3531
R ²	0.43	0.43	0.43	0.43	0.43	0.44
CA-type FE	No	Yes	Yes	No	Yes	Yes
Contract Chars FE	No	No	Yes	No	No	Yes

This table contains results to evaluate the robustness of the baseline DD estimates in Table 9 with respect to control group contamination (top two panels) and features of the awarding methods (bottom two panels).

Table 11: Robustness: Inference

Panel (a): All Contracting Authorities				
VARIABLES	(1)	(2)	(3)	(4)
	W.Discount	W.Discount	W.Discount	W.Discount
PA-Year	(-14.6;5.3)	(-14.4;5.1)	(1.8;8.6)	(1.6;8.2)
PA	(-5.7;-3.5)	(-5.7;-3.6)	(4.3;6.1)	(4.0;5.8)
Conley-Taber	(-8.9;6.4)	(-8.9;6.6)	(-0.3;15.5)	(-0.6;15.4)
PA-Year			(-22.4;-6.7)	(-22.0;-6.2)
PA			(-15.8;-13.3)	(-15.4;-12.8)
Conley-Taber			(-36.5;-3.0)	(-35.8;-2.3)

Panel (b): Contracting Authorities in Central Regions				
VARIABLES	(1)	(2)	(3)	(4)
	W.Discount	W.Discount	W.Discount	W.Discount
PA-Year	(-12.9;4.1)	(-13.0;3.6)	(3.6;9.5)	(2.8;8.8)
PA	(-7.4;-1.5)	(-7.7;-1.8)	(4.3;8.7)	(3.6;7.9)
Conley-Taber	(-9.0;-6.4)	(-8.5;-6.7)	(1.5;4.1)	(1.7;3.5)
PA-Year			(-23.4;-10.2)	(-22.7;-9.3)
PA			(-19.7;-13.9)	(-18.9;-13.1)
Conley-Taber			(-19.2;-16.6)	(-18.2;-16.4)

The table reports 95 percent confidence interval estimates for the same regression models presented in columns (2), (3), (5) and (6) of Table 9. The estimates in the three rows use different methods to compute standard errors: the top row uses clustering at the year and CA level and is thus identical to the point estimates in Table 9. The second row uses clustering at the CA level to account for autocorrelation. The third row uses the Conley-Taber adjustment for a small number of treatment units.

Table 12: Summary stats: Exiting and Incumbent firms

Panel (a): Contractors Entering the Firm's Auctions

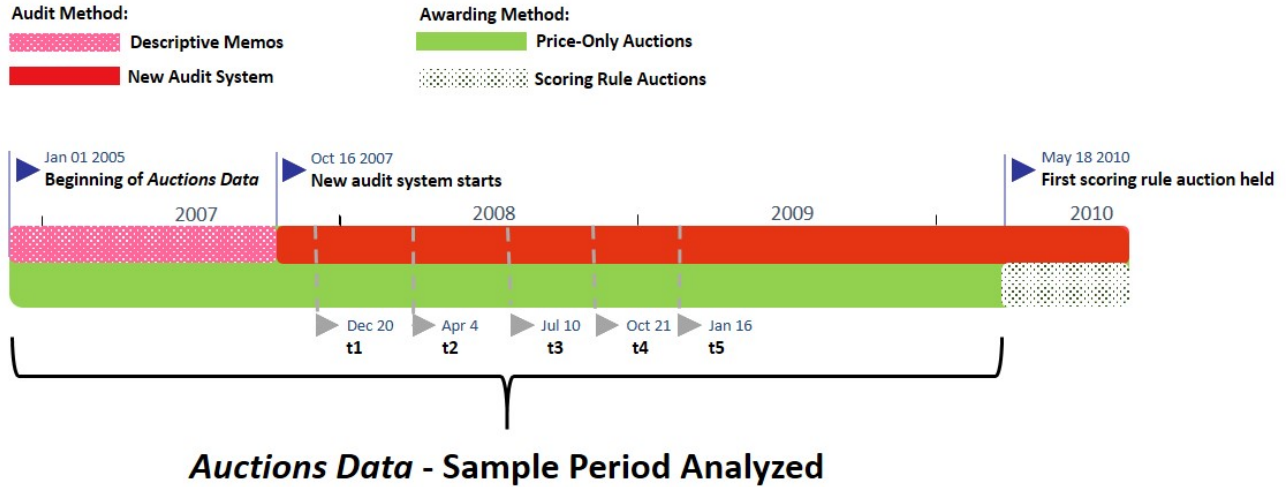
	<i>Exiters</i>				<i>Stayers</i>			
	(1) Mean	(2) p50	(3) SD	(4) N	(5) Mean	(6) p50	(7) SD	(8) N
Revenues	8,283	2,458	14,615	24	8,934	5,660	9,401	16
Profits	-21	6	697	24	32	5	73	16
Capital	391	36	788	24	998	47	2699	16
Number of Employees	10.3	5	11.1	24	51.7	4.50	180.4	16
Number of Managers	4.96	2	7.57	24	3.38	2	2.55	16
Proportion Female Managers	0.07	0	0.11	24	0.12	0	0.26	16
Public Company	0.96	1	0.21	23	0.88	1	0.34	16

Panel (b): Contractors Entering Turin's IRIDE Auctions

	<i>Exiters</i>				<i>Stayers</i>			
	(1) Mean	(2) p50	(3) SD	(4) N	(5) Mean	(6) p50	(7) SD	(8) N
Revenues	7,121	4,795	7,127	18	50,860	2,645	152,410	15
Profits	30	15	256	18	736	9.69	2,283	15
Capital	298	40	505	26	10,319	40	43,370	19
Number of Employees	9.04	9.50	5.53	26	15.1	8	15.8	19
Number of Managers	4.35	3	2.96	23	8.11	5	9.45	19
Proportion Female Managers	0.03	0	0.06	26	0.09	0	0.15	19
Public Company	0.71	1	0.46	24	0.72	1	0.46	18


The table reports summary statistics for the contractors bidding in the auctions. Panel (a) refers to the contractors active in the Firm's auctions, while panel (b) refers to the contractors bidding in the auctions of Turin's multi-utility company (IRIDE). Across all multi-utilities in the DD control group, this is the one for which we observe most contracts during the sample period. For both the Firm and IRIDE, we indicate as *exiters* those contractors observed bidding at least once before t_1 , but never after then, and as *stayers* those bidding at least once both before and after t_1 . For each of the 4 sets, the columns Mean, p50 and SD report the average, median and standard deviation taken across all firms in the set. Column N reports the number of firms considered. The firm characteristics considered are averaged over the years 2006-2010. They are: revenues, profits and capital (all expressed in €1,000), the number of all dependent workers (Number of Employees and Number of Managers), the fraction of female managers over all managers (Proportion of Female Managers) and the share of public companies.

Figure 1: Timeline




The chart illustrates the time span of the data and the timing of the Firm’s announcements. The *Auctions* dataset covers the period between 1/1/2005 and 4/1/2010. All auctions held in this period are price-only auctions. The Firm’s five announcements, marked t_1, \dots, t_5 , were held after the new auditing system started to inform suppliers about the functioning and the intended usage of this system. The implementation of the awarding rule based on both price and past performance is outside the analysis sample.

Figure 2: Average Compliance




Source: Audits data. The line shows the monthly average compliance calculated on all parameters inspected in the month of reference, weighting each parameter by its weight in the RI. The vertical lines identify each announcement date.

Figure 3: Winning Discounts (Dec 2004 - May 2010)




Source: Auctions data. The scatter plot reports all the winning discounts placed in the Firm's auction during the sample period. Note that the sample period is longer relative to that of the Audits data.

Figure 4: Evolution of Winning Bids




Sample: Auctions data. Evolution over time of the average winning discount for both treated (the Firm) and control (other CAs) procurers.

Figure 5: Safety and Quality: Average Weights across Audits




Source: Audits data. The plot represents the total weight, by audit, of parameters relating to Quality dimensions (grey bar) and Safety dimensions (black bar).

Figure 6: Safety and Quality: Evolution of Compliance over Time



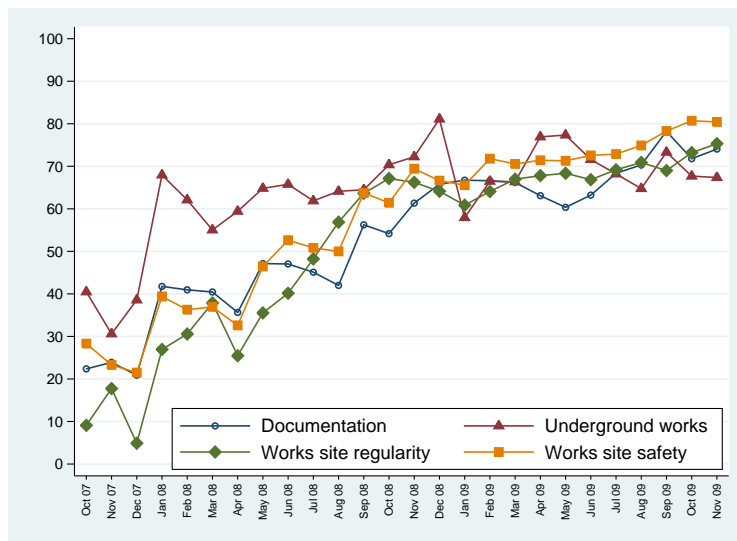
Source: Audits data. Monthly average compliance calculated separately for Safety and Quality on all parameters inspected in the month of reference, weighting each parameter by its weight in the RI. The vertical lines identify each announcement date.

Figure 7: Number of Parameters Audited



Source: Audits data. The bars represent the total number of parameters checked throughout the month of reference, distinguishing the compliant parameters (in black) from the not compliant ones (in grey).

Figure 8: Parameters Audited: Evolution of Compliance over Time




Source: Audits data. The lines show the progress of the reputation index calculated on a monthly basis for each of the four most audited Safety and Quality dimensions.

Figure 9: Contractors: Evolution of Compliance over Time




Source: Audits data. The four lines show the progress of the reputation index, calculated on a monthly basis, for 4 different groups of firms. The groups are formed on the basis of the firm’s success in concluding contracts. The line with circle markers represents the “most awarded” firms, triangles are for the “often awarded” group, diamonds are for “the less awarded group” and squares are for “the rarely awarded group.”

Figure 10: Discount CDF Pre-t1, t1-to-t5 and Post-t5



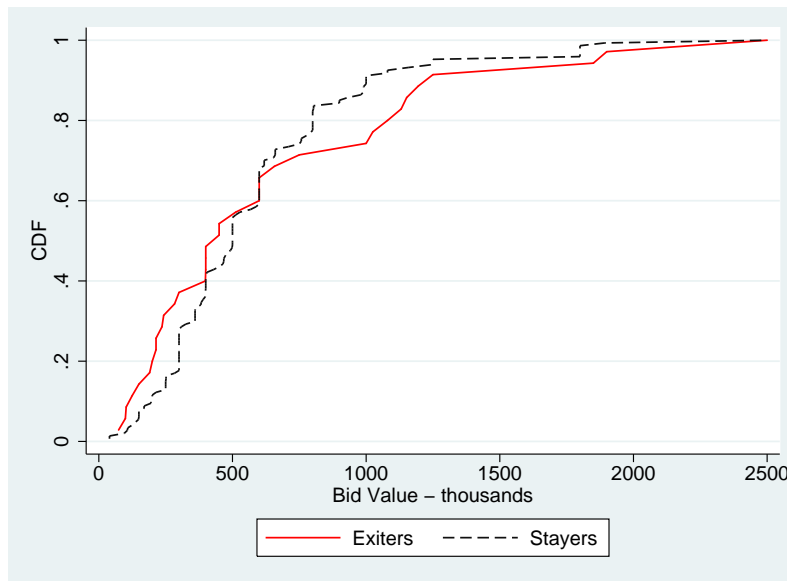
Source: Auctions data. The plot represents the cdf of the winning bid, dividing bids depending on the timing of the auction relative to t_1 and $t_5 + 1$.

Figure 11: Last auction date participated (*exitters*)



Source: Auctions data. Each bar represents the time until the supplier last bids. The figure is drawn for the sample of *exitters* only.

Figure 12: Bid CDF for exiting and incumbent firms



Source: Auctions data. The plot represents the cdf of winning bids for both *exitters* and *stayers* (in the pre- $t-1$ auctions).