CEIS TOR VERGATA Research Papers

Università di Roma Tor Vergata

ISSN 2610-931X

CEIS Tor Vergata

RESEARCH PAPER SERIES

Vol. 23, Issue 3, No. 598 – June 2025

Should I Share or Should I not? On the Sharing of Information on Past Performance in Procurement

Gian Luigi Albano, Walter Ferrarese,
Alberto Iozzi, Roberto Pezzuto

Should I share or should I not?

On the sharing of information on past performance in procurement

Gian Luigi Albano[*]
Consip S.p.A.
and LUISS Guido Carli University

Walter Ferrarese[†]
Universitat de les Illes Balears

Alberto Iozzi[‡]
Tor Vergata University of Rome
and SOAS University of London

Roberto Pezzuto[§]
Tor Vergata University of Rome

June 6, 2025

Abstract

Many real-world public-sector purchases involve a combination of verifiable and non-verifiable dimensions of quality, leading to a classical incomplete-contracting problem. This paper analyses how public buyers may use debarment lists — in essence, blacklists of under-performing contractors — to incentivize quality provision in repeated procurement tenders. A key question is whether debarment lists should be shared among multiple agencies or maintained separately. Sharing multiplies the punishment for bad performance (an under-performing firm loses access to all agencies, not just one), which might strongly deter shirking. However, this paper shows that sharing debarment lists backfires when mistakes may occur in judging quality ex-post: if one agency erroneously penalizes a cooperative contractor, that error propagates to every agency, potentially discouraging contractors from exerting high quality in the first place. By modelling repeated interactions and allowing for observational errors, we show the implicit costs stemming from a shared debarment list, and draw policy lessons for designing blacklists in public procurement.

Keywords: Public procurement; Relational contracts; Unverifiable quality; Debarment.
JEL codes: H57.

[*] gianluigi.albano@consip.it.
[†] walter.ferrarese@uib.es.
[‡] alberto.iozzi@uniroma2.it.
[§] roberto.pezzuto@uniroma2.it.

1 Introduction

In this paper we analyse how the implementation of debarment lists in public procurement tenders — whether used independently or shared across multiple buyers — may be used as a mechanism to incentivise high quality provision in repeated procurement interactions where executing the auctioned off projects entails exerting observable but non-verifiable, pay-off relevant quality.

Public procurement often involves acquiring complex goods and services that incorporate both verifiable and non-verifiable dimensions of quality. Buyers can in some cases rely on formal contracts and regulations to enforce certain basic criteria — what is sometimes referred to as "verifiable service characteristics" ([Buso et al., 2024]; [Eggleston, 2005]) — such as specific inputs, easily measured outputs, or minimum technical standards.

Even so, the intangible or subjective features of many services remain outside the scope of contract law. Such dimensions range from the proactivity of an IT service, to the appropriateness of a health care treatment, or even to the perceived taste quality of catering services (it is, alas, not yet a criminal offence to serve terrible coffee). These factors are often observable by the parties involved, yet not measurable. Procedural or outcome quality, for example, can be recognized by trained professionals or even by the final users of the procured good or service (e.g. professors using computers bought for the department, or patients spending time recovering in a hospital room), yet it may prove difficult or prohibitively expensive to verify it before an independent court.

Consequently, there is no straightforward way to enforce these aspects contractually. In practical terms, this means that while a public agency can specify, say, how many hours of consultancy it is paying for, or what types of materials must be used in a construction project, it has no reliable way of forcing the contractor to provide, for example, innovative project ideas, exemplary customer service, or maximum effort in meeting certain performance targets, at least not through contractual clauses that generate legally binding obligations. This challenge is not unique to public entities: in private transactions, sellers commonly rely on brand name, reputation, and long-term relationships to reassure buyers they will not shirk on non-verifiable performance requirements ([Bergman et al., 2012]).

When non-verifiable quality is important, procurement encounters a classic incomplete contracting problem, a topic that has been studied by a broad strand of the literature. Among many, beginning with [Hart, 1995], and then continuing with [Tirole, 1999], this literature has in particular highlighted that complex products or services often leave space for moral hazard. In our context, it means that contractors might shirk on quality attributes they know the buyer cares about but cannot verify and thus contract upon.

In the public sector, rules of transparency, accountability, and non-discrimination can limit a procurer's discretion in rewarding good suppliers and penalize bad ones for elusive, non-verifiable aspects of performance ([Kelman, 1990]; [Calzolari and Spagnolo, 2009]). Laws and regulations require governments to select suppliers through relatively inflexible processes, making it risky or even illegal to discriminate among firms based on subjective indicators like "past experience" or "quality of cooperation," especially if these are not fully documented or verifiable. Even in jurisdictions that allow for some degree of past performance evaluation (as in the United States), the exercise of buyer discretion must be explained rigorously to avoid accusations of favouritism or corruption ([Yukins and Kania, 2019]). The European Union has historically been more restrictive than the United States in letting public purchasers use non-verifiable past performance as a qualification or award criterion, precisely out of concern that such a practice might be abused to bar legitimate newcomers from public contracts ([Albano et al., 2023]).

Debarment rules have emerged as a formal and regulated response. Other than due to serious criminal offences (such as fraud, child labour, corruption, tax elusion, etc.) legislation has been developed in recent years that allow public procurement agencies to exclude firms from competitive tendering processes on more discretionary grounds, such as grave professional misconduct or deficient performance. That is, firms that deliver results judged as subpar, can be precluded from future government tenders, often for multiple periods, or indefinitely, in the strictest regimes (see [Drabkin and Yukins, 2022]; [Albano et al., 2017]). Countries such as the United States make broad use of debarment lists, partly to minimize performance risk (an underqualified firm will not perform satisfactorily) and reputational risk (the government's credibility suffers if it continues to award contracts to disreputable vendors, see [Yukins and Kania, 2019]). Over time, shared "blacklists" have come to serve a dual function: not only do they screen out known under-performers in a single agency or region, but they also disseminate warnings across agencies or jurisdictions, effectively punishing a misbehaving contractor in multiple markets simultaneously.

Despite its clear deterrent value, debarment raises complex trade-offs. On the one hand, the threat of exclusion can provide a powerful incentive for quality provision. A rational contractor that values future profits may exceed mere minimum verifiable standards if it knows that poor performance on any non-verifiable dimensions—once observed by the procurer—could lead to long-term exclusion and the forfeiture of all subsequent business ([Board, 2011]; [Calzolari and Spagnolo, 2009]). This dynamic parallels efficiency-wage theories in labour markets ([Shapiro and Stiglitz, 1984]), where high-powered incentives (like threats of firing) sustain effort under otherwise unobservable tasks. On the other hand, the extreme nature of debarment can harm competition if a market depends heavily on a small pool of qualified suppliers, particularly in specialized markets—since debarred firms may lose all avenues for future business ([Albano et al., 2017]). Second, given that actual performance is only imperfectly observed, even a cooperative contractor can be mistakenly flagged as non-compliant and wrongly excluded from future tenders. Such occurrence may arise from exogenous stochastic shocks, e.g. the changing of decision makers at the top of the procurement agency, or shifting interpretations of what acceptable quality really entails. These concerns are more compelling when multiple public agencies or separate jurisdictions share debarment lists. In principle, sharing information about "bad contractors" strengthens the deterrent effect: a firm that cheats in one market not only loses that market's future business but is also shut out elsewhere. However, being excluded "universally" may constitute an excessively large penalty in many cases, especially if the trigger for that penalty relies on subjective or otherwise noisy signals of performance. Moreover, from a purely strategic standpoint, an error-prone buyer—someone who might ban cooperating firms by accident—potentially undermines the overall efficacy of the shared debarment system, because that buyer's mistakes ripple through the entire network of procuring agencies.

In this paper, we develop a mode of infinitely repeated interaction to study how shared versus independent debarment lists interact with the risk of erroneous exclusion. We consider two buyers, each running a lowest-price, sealed-bid auction every period to procure a "project." Alongside a short-lived (local) firm that is replaced every period, both buyers face the same long-lived "large" firm that may participate in both auctions, provided it is not barred from them. After winning the contract, a firm chooses a level of non-contractible quality. If the firm fails to provide the buyer's desired standard, it is banned from future participation in that market (or in both markets, if the buyers share a single list). Crucially, we allow for the possibility that even a cooperating contractor is sometimes mislabelled as non-cooperative—a friction captured by a parameter which reflects each buyer's ability to correctly identify cooperative firms.

This setup contributes to the growing literature on relational contracts and repeated procurement. Early seminal contributions like [Bull, 1987] and [MacLeod and Malcomson, 1989] illustrate how repeated interactions can resolve moral hazard when certain actions are observable but not verifiable by courts. Subsequent works like [Levin, 2003] and [Fuchs, 2007] refines these ideas, exploring how optimal long-term agreements are structured, especially under moral hazard or adverse selection. Applications to procurement frequently highlight that a low-price auction can encourage moral hazard if quality is not contractible, yet repeated interactions—if carefully designed—may help mitigate under-provision ([Manelli and Vincent, 1995]; [Board, 2011]; [Calzolari and Spagnolo, 2009]). Nonetheless, these models typically rely on a straightforward assumption that buyers can identify "cheaters" without error once they observe the relevant outcome or signal. Our paper moves beyond such assumption by explicitly introducing the possibility that a contractor delivering the cooperative quality might still be flagged as non-cooperative, with positive probability. Drawing on the observation that many dimensions of project delivery are inherently subjective—different monitors or inspectors may hold diverse views on what the "correct" level of quality is—this friction changes both the structure of incentives and the buyer's willingness to commit to a strict debarment policy.

A key insight emerging from our analysis is that, when two markets share a single debarment list, the accuracy of each buyer's assessment interacts in a multiplicative manner. A "high-accuracy" buyer that almost never makes mistakes can actually suffer from linking its debarment decisions to those of a less reliable buyer. If the second buyer mislabels cooperating contractors with sufficiently higher probability, the large firm faces an excessive risk of a global ban each time it tries to supply that market. This can, paradoxically, reduce the potential to sustain high quality, because the firm becomes more reluctant to invest in cooperating if it fears a random ban across both markets. Our findings suggest that from the perspective of a precise procurer, sharing debarment lists with a more error-prone buyer effectively "imports" the second buyer's weaknesses and raises the discount factor threshold required for the large firm to find it worthwhile to keep cooperating. This suggests that policymakers should carefully weigh the pros and cons of unifying blacklists across multiple organizations.

The rest of this paper is organized as follows. In Section 2, we present the formal model, describing each buyer's environment, the firms' cost structures, and the stage game. Section 3 briefly addresses the one-shot version of the game to illustrate why quality provision unravels in static contexts. In Section 4 we then introduce the repeated game setting, specifying how mistakes enter into debarment decisions, and detailing how sharing a debarment list versus keeping separate lists modifies the large firm's incentives. Section 5 discusses the results and compares the two institutional settings. Section 6 briefly concludes.

2 The model

Markets and buyers. There are 2 independent markets (e.g., countries, regions, etc.), indexed with $j \in \mathcal{J} := \{A, B\}$. In each market, a buyer wishes to procure a project repeatedly over time: the procurement activity occurs at time t, where $t = 0, \ldots, \infty$. At each t (we will drop the reference to time whenever possible), buyer j derives from the project a gross utility equal to $v_j + q_j$, where $v_j > 0$ is the fixed intrinsic value of the project and $q_j \geq 0$ is the variable quality of the project. Quality is observable but unverifiable in a court of law; thus, it is not contractible. Letting p_j be the price paid for the project, buyer j's net utility is given by $U_j(v_j, q_j, p_j) = v_j + q_j - p_j$ when procuring the project, zero otherwise. Buyers have an identical discount factor, equal to $\delta \in (0, 1)$.

Firms. A finite number of firms $N_A \geq 1$ and $N_B \geq 1$ are active and can operate in their market j only. We will refer to these firms as "small" firms. In both A and B, there is also an infinite outside pool of inactive small firms ready to enter their respective markets, in the following sense: we assume that when a small firms exits either of the two markets, an identical number of firms from the outside pool will enter that market so that, in each repetition of the game described below, the number of small firms active in each market is always constant and equal to N_A and N_B. Small firms are short-lived: they only participate in the market once, then exit. In addition to the small firms, another firm is also present, with divisions in both markets A and B: we will refer to this firm as the "large" firm. This firm is long-lived: as long as it is allowed to, it is present in both markets forever. We assume that $N_A = N_B = 1$. Our hypotheses together imply that at each t, the large firm always competes with one small firm only in both markets, A and B.

Besides the restriction on operating in one single market described above, the small firms and the divisions of the large firm have an identical cost structure. Denoting with i the small firm or one of the divisions of the large firm, at each t cost is given by $C(\theta_i, q) = \theta_i + \psi(q)$, when they provide the project, zero otherwise. The fixed firm-specific cost component θ_i for each firm is a draw from a discrete random variable with support $\{\theta_L, \theta_H\}$ and with $\theta_L < \theta_H$. These draws are i.i.d. for all firms, except for the two divisions of the large firm, that always draw the same θ_i. Throughout, we assume $Pr(\theta_i = \theta_L) = \frac{1}{2}$; in other words, at each t, the probability of a firm being inefficient is identical across firms/divisions and independent of any previous realization of its cost parameter (if any). Furthermore, we assume that $v_J \geq \theta_H$, for all j, to ensure that the base level project is beneficial to the buyers under all cost realizations.

The other cost component $\psi(q)$ is the cost of producing quality q, assumed to be increasing and convex in q, with $\psi(0) = 0$. Thus, a small firm's profit is given by $\pi_i(p_j, \theta_i, q_j) = p_j - \theta_i - \psi(q_j)$ when they provide the project, and zero otherwise. Moreover, denoting with p_A and p_B the prices paid by the buyer in market A and B and with q_A and q_B the quality of the project delivered in market A and B, the large firm's profit is given by $\pi_\ell(p_A, p_B, \theta_\ell, q_A, q_B) = \pi_{\ell A}(p_A, \theta_\ell, q_A) + \pi_{\ell B}(p_B, \theta_\ell, q_B) = p_A + p_B - 2\,\theta_\ell - \psi(q_A) - \psi(q_B)$ when it provides the project in both markets, and zero otherwise (a similar expression applies in the case of the large firms providing the project in one market only).

Finally, all firms have the same discount factor δ as the buyers.

Mechanism. In each period, each buyer awards the project running a lowest-price sealed-bid auction. That is, in each auction, the project is awarded to the lowest-bid firm, which receives a payment equal to the submitted bid; in case of a tie, the project is awarded by using a fair random device. We refer to the winning firm as the contractor.

Admission to the auctions is by invitation by the buyers. That is, a firm may be restricted from participating in the auction at a given t. If such restriction holds starting at time t and for future repetitions of the game (if any), we say that that firm has been debarred from market j at time t or, equivalently, that this firm has been added at time t to the buyer j's debarment list. In reality, debarment from procurement tenders is almost always temporary, that is, in the terms of our framework, firms can be added to the debarment lists for $k < \infty$ subsequent periods. Since this would greatly complicate the analysis, in this paper we assume that once a firm enters a debarment list, it cannot get out of it: once it occurs, debarment is permanent. Our aim here is to focus on the strategic implications that arise from using debarment lists as mechanisms to incentivise quality provision. Any effect we get will just be of a lower magnitude when bans are allowed to be temporary.

Since quality is not contractible, bids are one-dimensional. Bids are assumed to be discrete numbers and, thus, are a multiple of the smallest monetary unit, ϵ, which is assumed to be arbitrarily small. Firms face no bidding costs.

Informational Structure. We consider a game of incomplete information. Each buyer has incomplete information on the costs of the firms operating in its market, as they do not observe the realisations of θ. Conversely, each firm perfectly knows its own and the cost parameters of the rivals operating in the same market when they are realised. All players observe the bids made by all firms in the same market. After completion of the project in one market, only the buyer (and the contractor) observes the quality delivered by the contractor.

Furthermore, we analyse and compare two different institutional scenarios. In the first one, we assume that the debarment list is specific to each market. This implies that if the large firm (the only firm able to operate in both markets) is debarred in one market, it still can participate in the auctions in the other market. Referring to the buyers' activity of selecting the firms participating in the auctions, we define this set-up as *no information sharing.* In a second game, we assume that a single debarment list is common to both markets. This implies that when the large firm is debarred in one market, it also loses the possibility of participating in the auction in the other market. We define this alternative setup as *information sharing.*

The game. At each t, the following stage game takes place in each market:

stage 1: (the *debarment stage)* the buyer chooses the firm(s) participating in the auction by adding to its debarment list the firm(s) it does not want to participate.

stage 2: (the *participation stage)* firms learn their own cost, and their rivals'. Each firm not on the debarment list chooses whether or not to participate in the auction. Debarred firms exit the game forever.

stage 3: (the *bidding stage*) firms participating in the auction bid. The firm placing the lowest bid wins the auction and is awarded the project. In case of a tie among m firms, the project is awarded randomly to one firm with probability $\frac{1}{m}$.

stage 4: (the *execution stage)* the contractor delivers the project, choosing quality $q \in [0, \infty)$, that is observed by the buyer. Payoffs are realised.

We analyse the infinitely repeated games resulting from an infinite repetition of this stage game. In the case of *no information sharing*, the two dynamic games taking place in the two markets are independent of one other and can be analysed in isolation. This is not the case with *information sharing*, since the large firm's actions in one market affect the game's unfolding in the other market.

We look for Perfect Bayesian equilibria in which players play undominated strategies on and off-path.

3 The stage game

Let us first analyse the static, one-shot version of the game described above. Here, at equilibrium, any firm selected as contractor will find it optimal to deliver quality $q = 0$. Thus, in the bidding stage, all participating firms bid anticipating $\phi(0) = 0$ in the cost function of all firms. In stage 2, participating is weakly dominant for all cost realizations, thus all firms participate to their respective auctions. In stage 1, each buyer always prefers to admit both the local small firm and the local division of the large firm, given that $v_j \geq \theta_H$ — the worst case scenario for the buyers. In summary, we get:

Proposition 1. *The static game admits a unique Perfect Bayesian Equilibrium where all buyers choose to admit both firms to their respective auctions. The winning contractors deliver quality equal to zero.*

Here the buyers are effectively powerless: they will never get any positive exerted quality, and will award the project with the low price θ_L only with probability 1/4: in half of the possible combinations of cost realisations, either firm will be efficient facing an inefficient opponent, bidding $\theta_H - \epsilon \approx \theta_H$, and in the remaining cases, half of the time both will be inefficient, again bidding θ_H.

4 The repeated games

In this section, we analyse the infinitely repeated game resulting from the infinite repetition of the stage game described in Section 2. We will look at the two scenarios of *no information sharing* and *information sharing*.

The infinitely repeated environment adds two important features to our games. The first one is the players' possibility to use history-dependent strategies. A second one is instead a consequence of the institutional features we want to capture in our model: debarment is not limited to a single period, but it has a long-lasting effect. Recall that we take this feature to the extreme and assume that once a firm has been debarred in a market, it is not allowed to participate in the auction in that market forever.

Our main object of analysis is the dynamic incentive to the provision of unverifiable quality by the contractor that the prospect of debarment may provide. In particular, as it will be described more formally later in this Section, we want to investigate whether the threat of being added to the debarment list by the buyer from time $t+1$ in case of a quality provision different from the cooperative level q^C at time t may provide the firm the incentive to offer quality q^C. We want to investigate whether the incentive properties of such a threat may depend on the institutional feature of the communication between buyers, that is, whether or not the debarment following a poor past performance is extended to other markets.

Here, we introduce some friction in our incentive mechanism of the debarment list by allowing for the possibility that the buyer debars a firm that delivered q^C, thus identifying it as non-cooperating. In other words, we allow for the possibility that cooperating contractors may be wrongly flagged as non-cooperating. This is meant to represent the fact that, even when contractors do cooperate, things can go awry: the buyer's evaluation of the realised project's quality, and thus her subsequent decision on whether to debar the firm or not, might differ from the actual quality that was exerted by the firm. This could be due to many reasons: different interpretations of what "q^C" actually means between buyers and contractors, or an unexpected change in the people responsible for deciding what firms to add to the debarment list, who might have different opinions regarding what q^C actually corresponds to, i.e. a buyer in a sense changing her mind ex-post. Alternatively, there could be an exogenous and stochastic shock that makes the quality intended and exerted by the firm differ from the quality observed by the buyer. Indeed, these are some of the reasons why quality, albeit observable — possibly imperfectly, is often not contractible, as it is difficult if not impossible to objectively measure it. We refer to this friction as a "mistake" made by the buyer, and define $1-\alpha_j$ as the probability that such event occurs in market j. In other words, α_j is an index that measures the buyer's institutional stability, and ability in correctly identifying cooperating firms. Note that this is a physical property of our environment: it is not something over which players may exert any control, and implies that firms may be debarred forever due to an exogenous force.

We conduct our analysis with players adopting the following grim-trigger strategies:

> *Buyer j:* at $t=0$, no firm is added to the debarment list. At any $t \geq 1$, if the quality provided by the contractor at $t-1$ is different from q_j^C, it adds the contractor to the debarment list.

Otherwise, if the quality provided by the contractor at $t-1$ is equal to q_j^C, it does not add the contractor to the list.

Large firm: at any $t \geq 0$, if not debarred, each division of the large firm chooses to participate in the auction, bids in its market as in a standard Bertrand game, where its costs are equal to $C_j(\theta_\ell, q_j^C) = \theta_\ell + \psi\left(q_j^C\right)$ and the rival's cost are $C_i(\theta_i, q_j^C) = \theta_i$. If awarded the project, it delivers quality q_j^C. Otherwise, if it is included in the debarment list, it does not bid.

Small (active) firm: at any $t \geq 0$, if not debarred, it chooses to participate in the auction, bids in its market as in a standard Bertrand game, where its costs are equal to $C_i(\theta_i, 0) = \theta_i$ and the rival's cost are $C_j(\theta_\ell, q_j^C) = \theta_\ell + \psi\left(q_j^C\right)$. If awarded the project, it delivers quality equal to 0. Otherwise, if it is included in the debarment list, it does not bid.

Let these strategies be denoted as σ_j for the buyers, σ_ℓ for the large firm, and σ_i for small firm i. They are functions that map - at each stage of each repetition of the game - the publicly observable history into the players' actions.

The above strategies describe a set of bilateral relational procurement contracts under which, in each repetition of the game after the first, buyers add each small firm to the debarment list - if previously selected as contractor and thus providing zero quality, and never add the large firm to the list. Small firms bid anticipating not to deliver quality, while the large firm bids in each market anticipating to provide the quality desired by the respective buyer. When a small firm wins the auction, it delivers quality equal to zero, it is thus debarred in the next period and substituted by another identical firm in the outside pool. When instead the large firm is selected as a contractor in a market, it delivers the quality desired by the buyer. Such an outcome is sustained by the buyer's threat to add the large firm to the debarment list when it is selected as a contractor and does not deliver quality.

This set of RPCs is self-enforcing if the players' strategies form a Perfect Bayesian Equilibrium of the repeated game. We characterise the equilibrium of the game by checking the conditions for the absence of profitable one-shot deviations (POSDs henceforth) for each player ([Mailath et al., 2006]).

Our players have several possible deviations from their candidate equilibrium strategies. First, the buyers can deviate by changing their admission strategy. In principle, every kind of such deviations is allowed, e.g.: from never debarring any firm under any contingency, to debarring a given firm only in some given repetitions, debarring only the large/small, etc. Still, it is important to keep in mind that debarring once entails debarring forever, as firms cannot come off of the debarment list. Secondly, the firms may deviate in the participation stage (e.g.: the large firm might decide to not participate in some repetitions and avoid the risk of being banned, or not participate in just one market, etc.); deviate in the bidding stage to any price-bid; deviate in the execution stage to any quality; and finally they may deviate using any combination of the above deviations. This makes the analyses of the incentives potentially very complex. As we will show in the next sections, we are able to restrict our analysis to just a few types of deviations. In the next section we will analyse, in the two scenarios outlined above, the incentives that sustain the above described set of RPCs.

4.1 The case of two independent debarment lists

In this section, we assume that the buyer in market $j = \{A, B\}$ styles a debarment list whose content is not shared with buyer $k \neq j$. In this environment, we will discuss the buyers' and the firms' incentives to deviate from the aforementioned strategies — σ_i, σ_ℓ, and σ_i, with $i = \{A, B\}$.

We begin our analysis by first defining $\rho(q_\ell)$ as the probability that the large firm wins the auction at any t, for a given profile of qualities that she is anticipating to play at any occurring execution stage. Then, we provide a preliminary result which we will use throughout the section.

Lemma 1. *In the case of two independent debarment lists, the probability that both the large and a small firm are present in the game at time $t \geq 0$ is:*

$$\xi(\rho(q_\ell), \alpha_j, t) = [1 - \rho(q_\ell)(1 - \alpha_j)]^t \tag{1}$$

The intuition to understand expression (1) is as follows. Since the small firms never disappear (they cannot all get debarred forever, as there is an infinite supply of them in the outside pool), the buyers face either one or two firms at the beginning of each repetition, but never zero. Consider now the grim trigger strategies for the buyers in this setting, that is, they debar the large seller if they observe that she delivered a quality different than the cooperative one. Recall that buyers commit a mistake with probability α_j: a cooperating contractor may be banned forever with positive probability. Let k be the number of times that the large firm has won an auction, up to a given repetition. Now, at time t, by independence, the probability of a history where the large firm has won $k \leq t$ times and has lost $t - k$ times is $\rho^k(1 - \rho)^{t-k}$. Secondly, note that there can be $\binom{t}{k}$ numbers of such a history. Thirdly, to have two firms present at the beginning of each repetition, it must be that each of the k times the large firm won, the buyer has not the erroneously debarred her. This happens with probability α_j^k. Finally, this analysis must be carried out for all $k \leq t$. Thus, the overall probability that at a repetition t, *on-path*, a buyer faces two firms is the following (dropping the buyers' index j in α and the argument of ρ for clarity's sake):

$$\sum_{k=0}^{t} \binom{t}{k} \rho^k (1 - \rho)^{t-k} \alpha^k \tag{2}$$

Let us unpack this formula: first, note that it holds for every repetition. For example, trivially, at time 0, the large firm will be present at the auction for sure, hence the probability of facing two firms should be 1. At time 1, to have two firms, either the large firm did not win at time 0, which happens with probability $(1-\rho)$, or she did win, and did not get erroneously banned, which happens with probability $\alpha \cdot \rho$, so the probability should be $(1-\rho)+\alpha\rho$. Indeed, following expression (2), we have:

$$t = 0 \rightarrow \binom{0}{0}\rho^0(1-\rho)^{0-0}\alpha^0 = 1$$
$$t = 1 \rightarrow \binom{1}{0}\rho^0(1-\rho)^{1-0}\alpha^0 + \binom{1}{1}\rho^1(1-\rho)^{1-1}\alpha^1 = (1-\rho)+\alpha\rho$$

At time 2, to have two firms, we have to consider the fact that the large firm can win both previous auctions and not get banned twice ($a^2\rho^2$), lose both ($(1-\rho)^2$), but also win the first, lose the second, and not get banned precisely when it won ($\rho\alpha(1-\rho)$), and the complementary case ($(1-\rho)\rho\alpha$). Thus the probability should be: $a^2p^2 + 2\rho(1-\rho)\alpha + (1-\rho)^2$, which indeed we also get by using our expression (2) at time $t = 2$.

$$t = 2 \rightarrow \binom{2}{0}\rho^0(1-\rho)^{2-0}\alpha^0 + \binom{2}{1}\rho^1(1-\rho)^{2-1}\alpha^1 + \binom{2}{2}\rho^2(1-\rho)^{2-2}\alpha^2 = a^2\rho^2 + 2\rho(1-\rho)\alpha + (1-\rho)^2$$

And so on and so forth. To get to equation (1) in the lemma, we just solve the sum in equation (2) and simplify.

4.1.1 The buyers' incentives

Provided that all the firms in the market under analysis follow their prescribed strategies, buyer j does not have a POSD if:

$$\Pi_j^C \geq \mathbb{E}[\pi_j^D] + \Pi_j^D \tag{3}$$

where Π_j^C is the discounted value of the stream of expected payoffs while sticking to the strategy σ_j, $\mathbb{E}[\pi_j^D]$ is the one-shot expected payoff while deviating from the cooperative strategy, and Π_j^D is the stream of expected payoffs along the punishment path.

In the left-hand side (LHS hereafter) of equation 3, the buyer may obtain a positive cooperative payoff at any time t only when both firms are present at the beginning of the repetition (since, both on and off-path, a firm participating alone in an auction can bid a price that extracts the entire buyer's valuation for the project). This happens with the probability given in expression (1). Since we will make use of it later on in the paper, let us show the multiplier of the per period expected cooperative payoff, used to calculate the discounted value of sticking to the strategy σ_j. Note that on-path the small firms play $q = 0$, and the large firm bids anticipating to exert quality q^C when selected as a contractor, thus her probability of winning an auction, $\rho(q^C)$, is a constant, which we indicate hereafter with $\bar{\rho}$, that can be fully anticipated by the players. Hence, by Lemma 1, the coefficient multiplying $\mathbb{E}[\pi_j^C]$ in the LHS of (3) is obtained as follows:

$$\sum_{t=0}^{\infty} \delta^t [1 - \bar{\rho}(1 - \alpha_j)]^t = \frac{1}{1 - \delta - \alpha_j \bar{\rho}\delta + \delta\bar{\rho}}. \tag{4}$$

Regarding the right-hand side (RHS hereafter) of equation 3, out of all the possible types of deviations, it is sufficient to consider three specific ones. First, in the deviating period, a buyer can add both firms to the debarment list. This is obviously unprofitable as it would involve getting a zero payoff forever with certainty. Similarly, it cannot be profitable to indict only the large firm since, at each subsequent period, the unique small firm will be able to bid in such a way as to extract all the surplus from the buyer. The last option is to deviate to a ban of the small firm only: in this case, the buyer gets zero in the deviation period t, in which the large firm participates in the auction alone, and goes back to the cooperative path from $t+1$. Hence, here the RHS would coincide with the LHS in (3) but with an extra discount factor in front, making this deviation disadvantageous as well. Thus, the incentive compatibility constraints of the buyers are always satisfied.

4.1.2 The firms' incentives

Small firm. Concerning the small firms, recall that the grim trigger strategy σ_i for a non-debarred small firm is to participate in the auction, bid in its market as in a standard Bertrand game anticipating to deliver zero quality, and, if awarded the project, to deliver quality equal to zero. It is straightforward to see that this is already the only strategy profile sustainable at equilibrium for the small firm. Deviating in the participation stage is weakly dominated by participating. The bidding behaviour in σ_i is already assumed to be the optimal one given the auction format. Finally, deviating to any positive quality is strictly dominated by playing $q = 0$, since small firms are short lived. Thus, the small firm incentives are always satisfied.

Large firm. Next, we look at the large firm incentives. Provided that the buyers and the small firm active in the market follow their prescribed strategies, the large firm has no POSD if:

$$\Pi_{\ell,j}^C \geq \mathbb{E}[\pi_{\ell,j}^D] + \Pi_{\ell,j}^P \quad \text{for} \;\; j = \{A, B\} \tag{5}$$

where $\Pi_{\ell,j}^C$, $\mathbb{E}[\pi_{\ell,j}^D]$, $\Pi_{\ell,j}^P$ have the same interpretation as in the previous section. Regarding the LHS of the two ICCs, the first thing to notice is that due to the presence of a small firm at any repetition, the large firm is never going to be the only potential participant. Hence, the coefficients multiplying the (two — one for each market) expected cooperative payoffs are given by equation (4). Thus, the LHS of equation

(5) becomes:

$$\frac{1}{1-\delta-\alpha_j\bar{\rho}\delta+\delta\bar{\rho}}\cdot\mathbb{E}[\pi^C_{\ell,j}], \quad \text{for } j=A,B. \tag{6}$$

Regarding the RHS of equation (5), and differently from the buyers, it is straightforward to see that the large firm does not always trivially cooperate: some deviation payoffs are not lower than the cooperating ones everywhere. Out of all the possible ones, it is sufficient to consider the best type of deviations: that is the ones where, in a period t, the large firm participates in both markets and, in each, bids anticipating $\phi(0)$, and then delivers zero quality if selected as contractor. Although the buyer is going to debar the large firm from $t+1$ onwards with certainty, the cumulative one-shot deviation payoff could offset the cooperative stream of profits.

To see why this is the only relevant type of deviation, consider the possible alternatives. First, due to fact that markets are independent, the large firm cannot be banned in market j when delivering a non-cooperative quality in market $k\neq j$. In other words, participating in one market is weakly dominated by participating in both, for any anticipated bid and exerted quality profile. This, *a fortiori*, makes not participating in both markets even less profitable. As said, the aforementioned one-shot deviation may happen in any repetition: the deviating firm could keep getting the expected cooperation payoffs up to time $t-1$, deviate in t, and go back on path from $t+1$ onwards. Thanks to the one-shot deviation principle ([Mailath et al., 2006]), we may restrict our attention to consider a deviation in the first period, $t=0$, which gives us a sufficient and necessary condition for σ_ℓ to be sustainable at equilibrium. Thus, we can summarise our discussion in the following result:

Lemma 2. *In the case of two independent debarment lists, the large firm has no POSD iff:*

$$\frac{1}{1-\delta-\alpha_j\bar{\rho}\delta+\delta\bar{\rho}}\,\mathbb{E}\,\pi^C_{\ell,j}\geq\mathbb{E}\,\pi^D_{l,j} \quad \textit{for all } j=\{A,B\} \tag{7}$$

Small firms. Concerning the small firms, recall that the grim trigger strategy σ_i for a non-debarred small firm is to participate in the auction, bid in its market as in a standard Bertrand game anticipating to deliver zero quality, and, if awarded the project, to deliver quality equal to zero. It is straightforward to see that this is already the only strategy profile sustainable at equilibrium for the small firm. Deviating in the participation stage is weakly dominated by participating. The bidding behaviour in σ_i is already assumed to be the optimal one given the auction format. Finally, deviating to any positive quality is strictly dominated by playing $q=0$, since small firms are short lived. Thus the small firm incentives are always satisfied.

It follows that equation (7) in Lemma 2 fully characterizes the conditions required for the grim strategies to hold.

4.2 The case of a common debarment list

In this section we discuss the case of a centralised debarment list that is fed by the debarment decisions of both buyers. This implies that a firm not providing quality in one market is debarred forever in both markets. Also, given the possibility that either buyers ban a firm providing the cooperative quality, even a cooperating firm that operates in both markets may be banned forever and everywhere as a result of such a mistake by one of the two (or both) buyers.

Mirroring the previous section, we begin our analysis by providing a preliminary result:

Lemma 3. *In the case of a common debarment list, the probability that both the large and a small firm are present in the game at time $t \geq 0$ is instead of :*

$$\xi'(\rho(q_A), \rho(q_B), \alpha_A, \alpha_B) = [1 - \rho(q_A)(1 - \alpha_A)]^t \times [1 - \rho(q_B)(1 - \alpha_B)]^t \tag{8}$$

Since the large firm could be debarred in either market, the probability that it is allowed to bid in market j at time t is the product of the probability of not being debarred in the same market times the same probability in the other market. In other words, two firms (one large and one small) will be operating in market A after every history that envisages two firms in market A, but also two firms in the other market. The probability that each of the two histories unfolds during the play can be found using the same argument as in the discussion following Lemma 1, for the *no information sharing* case. Thus the probability that the two firms operate in market i at time t is given by

$$\left(\sum_{k_A=0}^{t} \binom{t}{k_A} \rho(q_A)^{k_A} (1 - \rho(q_A))^{t-k_A} \alpha_A^{k_A}\right) \times \left(\sum_{k_B=0}^{t} \binom{t}{k_B} \rho(q_B)^{k_B} (1 - \rho(q_B))^{t-k_B} \alpha_B^{k_B}\right) \tag{9}$$

where k_A and k_B are the number of auctions awarded to the large firm in market A and B, respectively. Similarly to the case of *information sharing*, this reduces to 8.

4.2.1 The Buyers' incentives

The same reasoning regarding a deviation of the buyers in the *no information sharing* case applies: any possible deviation strictly decreases the buyers on-path stream of payoffs. Hence, as in the case of *no information sharing*, the ICC of the buyers are always satisfied.

4.2.2 The firms' incentives

Small firm. Concerning the ICC of each small firm, nothing changes from the two independent debarment list world: the behaviour prescribed by σ_i involves playing dominant strategies in each stage of the game.

Large firm. We now focus on the large firm. Provided that the buyers and the small firm active in the market follow their prescribed strategies, the large firm has no POSD if:

$$\Pi^C_{\ell,A} + \Pi^C_{\ell,B} \geq \Pi^D_\ell + \Pi^P_\ell, \tag{10}$$

where $\Pi^C_{\ell,j}$ has the same interpretation as in the previous section, and Π^D_ℓ, and Π^P_ℓ are the discounted streams of expected profits in case of deviation and punishment, respectively.

We begin by deriving the discount factor for the different terms in 10, by using Lemma 3. As to the sum of the cooperative payoffs, the multiplier for the firms payoff in each period is thus given by

$$\sum_{t=0}^{\infty} \delta^t [1-\bar{\rho}+\alpha_A\bar{\rho}]^t [1-\bar{\rho}+\alpha_B\bar{\rho}]^t = \frac{1}{1-\delta(1-\bar{\rho}+\alpha_A\bar{\rho})(1-\bar{\rho}+\alpha_B\bar{\rho})} \equiv \mathcal{D}^C. \tag{11}$$

Equation (11) is the multiplier that gives us the stream of the cooperation payoffs for the large firm, on the equilibrium path. We have thus obtained the LHS of the large firm's incentive, that is, the version of expression (6) for the common debarment list world:

$$\Pi^C_{\ell,A} + \Pi^C_{\ell,B} = \mathcal{D}^C \times \left[\mathbb{E}[\pi^C_{\ell,A}] + \mathbb{E}[\pi^C_{\ell,B}]\right]. \tag{12}$$

We now turn our attention to the RHS of 10. The values of Π^D_ℓ and Π^P_ℓ depend on the deviation chosen by the firm. We can restrict our attention to three relevant types of deviations at time $t = 0$:

1) *Full deviation (FD):* at the participation stage, participate in both markets and, in the execution stage, if awarded the project, set $q = 0$ in any market;

2) *Partial deviation A (PD.A):* at the participation stage, participate in market B and do not participate in market A and, in the execution stage, if awarded the project in market B, set $q = q^C$;

3) *Partial deviation B (PD.B):* as in *PD.A*, with A and B inverted.

We consider the corresponding payoffs. The total payoff associated with the **full deviation** coincides with the one-shot deviation payoffs:

$$\mathbb{E}[\pi^D_{\ell,j}], \quad \forall\, j \tag{13}$$

afterwards, the large firm is debarred forever and gets nothing. Regarding the **partial deviations**, as in section 4.1.2, such deviation profiles may occur at any time t. Again to simplify the analysis, WLOG, we can consider the deviation occurring in the first period, $t = 0$. Thus, the large firm does not participate in market B (A) and loses the corresponding expected payoff in the first repetition. However, by doing so, it has a greater chance of winning in both markets in the next periods, since it cannot be banned due to a potential mistake occurring in market B (A). Indeed, following the behaviour prescribed by this deviation strategy, the probability of being present in each repetition is given by the following:

$$\left(\sum_{k_A=0}^{t}\binom{t}{k_A}\bar{\rho}^{k_A}(1-\bar{\rho})^{t-k_A}\alpha_A^{k_A}\right)\times\left(\sum_{k_B=0}^{t-1}\binom{t-1}{k_B}\bar{\rho}^{k_B}(1-\bar{\rho})^{t-1-k_B}\alpha_B^{k_B}\right) \tag{14}$$

which is greater than equation (9), i.e. the probability of the same event occurring following the on-path strategy profile. Moreover, this type of deviation does not imply any punishment, as it does not involve playing any quality different than q^C. It follows that the RHS of the large firm's ICCs for the partial deviation B writes:

$$\mathbb{E}\,[\pi^C_{\ell,A}] + \mathcal{D}^{RHS}_{\sigma_\ell^{\to B}}(\alpha_A,\alpha_B,\bar{\rho},\delta)\,\mathbb{E}[\pi^C_{\ell,j}], \quad \forall\, j \tag{15}$$

Where

$$\mathcal{D}^{RHS}_{\sigma_\ell^{\to B}}(\alpha_A,\alpha_B,\bar{\rho},\delta) = \frac{\delta((\alpha_A-1)\bar{\rho}+1)}{(1-\delta((\alpha_A-1)\bar{\rho}+1)((\alpha_B-1)\bar{\rho}+1))}.$$

The expression in (15) has the following interpretation: in the first period the large firm fully cooperates only in market A, thus getting only $\mathbb{E}\,[\pi^C_{\ell,A}]$ and renouncing to $\mathbb{E}\,[\pi^C_{\ell,B}]$. Then, from $t = 1$ onwards, it gets back on path, and the associated (expected) discounted stream of payoffs is calculated taking into account that it is now more likely (as per equation 14) to "survive" after $t = 0$. Similarly for partial deviation A we have:

$$\mathbb{E}\,[\pi^C_{\ell,B}] + \mathcal{D}^{RHS}_{\sigma_\ell^{\to A}}(\alpha_A,\alpha_B,\bar{\rho},\delta)\,\mathbb{E}[\pi^C_{\ell,j}], \quad \forall\, j \tag{16}$$

Where

$$\mathcal{D}^{RHS}_{\sigma_\ell^{\to A}}(\alpha_A,\alpha_B,\bar{\rho},\delta) = \frac{\delta((\alpha_B-1)\bar{\rho}+1)}{(1-\delta((\alpha_A-1)\bar{\rho}+1)((\alpha_B-1)\bar{\rho}+1))}.$$

Finally, we can summarise this section with the following result, which gives us the ICCs of the large firms and thus, since they are the only ones that are not always satisfied, the necessary and sufficient conditions to have the grim trigger strategies σ_j, σ_ℓ, and σ_i hold at equilibrium. Combining equations (12), (13), (15), and (16), we get the following Lemma:

Lemma 4. *In the case of a common debarment list the large firm has no POSD iff:*

$$\mathcal{D}^{LHS}_{\sigma_\ell}(\alpha_A,\alpha_B,\bar{\rho},\delta)\,\mathbb{E}[\pi^C_{\ell,j}] \geq \mathbb{E}\,[\pi^D_{\ell,j}], \quad \forall\, j \qquad FD \tag{17}$$

$$\mathcal{D}^{LHS}_{\sigma_\ell}(\alpha_A,\alpha_B,\bar{\rho},\delta)\,\mathbb{E}[\pi^C_{\ell,j}] \geq \mathbb{E}\,[\pi^C_{\ell,A}] + \mathcal{D}^{RHS}_{\sigma_\ell^{\to B}}\,\mathbb{E}[\pi^C_{\ell,j}], \quad \forall\, j \qquad PD.B \tag{18}$$

$$\mathcal{D}^{LHS}_{\sigma_\ell}(\alpha_A,\alpha_B,\bar{\rho},\delta)\,\mathbb{E}[\pi^C_{\ell,j}] \geq \mathbb{E}\,[\pi^C_{\ell,B}] + \mathcal{D}^{RHS}_{\sigma_\ell^{\to A}}\,\mathbb{E}[\pi^C_{\ell,j}], \quad \forall\, j \qquad PD.A \tag{19}$$

5 Comparison of debarment list regimes.

In this section we assess the impact that the two debarment list regimes have on the threshold levels of the discount factor δ which guarantee that the grim strategies $\sigma_i, \sigma_\ell,$ and σ_j hold at equilibrium, as implied by Lemmas 2 and 4. That is, we assess the merits of the debarment list mechanism to incentivise quality provision, and show the potential benefits and costs of sharing a single debarment list.

First of all, note that on the equilibrium path, since for all firms $Pr(\theta_i = \theta_L) = \frac{1}{2}$, then the probability that a cooperating large firm wins the auction is $\rho(q_\ell^C) = \bar{\rho} = \frac{1}{4}$. This is because, given the behaviour prescribed by the strategies, the large firm is going to win the auction with certainty only when she is efficient facing an inefficient small firm, which happens with probability $\frac{1}{4}$. The small firms, anticipating $\phi(0)$, find it optimal to bid an amount equal to the highest cost realization. The large firm, which has to take into account the cost $\phi(q^C)$ that she is going to sustain in the execution stage, is able to slightly underbid the small firms by $\epsilon \approx 0$ and win only in the aforementioned case, netting a payoff of $\theta_H - \theta_L - \phi(q^C)$. In all other cases she loses with certainty: she is unable to match her opponent's bid (equal to the realization of the highest cost parameter), as it would incur in losses with positive probability, which is weakly dominated by bidding a higher amount and losing the auction. The remaining parameters to discuss are the two α_j, that is, the parameters that measure the buyers ability to correctly detect a cooperating firm, or alternatively the buyers' stability of judgment in evaluating their contractors' performances. We consider two different scenarios.

5.1 Buyers make no mistakes.

As a benchmark, let us first analyse the case where $\alpha_j = 1$ for all buyers. That is, the case where buyers never make mistakes and erroneously debar cooperating contractors. It is straightforward to see that this implies that now two firms are always present in each repetition: equations (1) and (3) are both equal to 1. We can now identify the threshold level that allows the grim strategies to hold.

Proposition 2. *Consider $\alpha_A = \alpha_B = 1$ and let $\bar{\delta}_{base} = \frac{\phi(q^C)}{\theta_H - \theta_L}$. Then if $\delta \geq \bar{\delta}_{base}$:*

i) the vector of strategies $(\sigma_j, \sigma_\ell, \sigma_i)$ defines a set of RPCs in which q_j^C is played whenever the large firm is selected as a contractor in market j, which happens with probability $\bar{\rho} = \frac{1}{4}$.

ii) Having one shared or two independent debarment lists is irrelevant.

Proof.

Since firms are symmetric and their expected payoffs, on and off-path, only depend on the vector of bids and thus on the expected costs, $\mathbb{E}[\pi_{\ell,j}^D] = \mathbb{E}[\pi_\ell^D] = \frac{1}{4}(\theta_H - \theta_L)$, and $\mathbb{E}[\pi_{\ell,j}^C] = \mathbb{E}[\pi_\ell^C] = \frac{1}{4}(\theta_H - \theta_L - \phi(q^C))$. Thus, equation (7) in Lemma 2 boils down to:

$$\frac{1}{1-\delta}\,\mathbb{E}[\pi_\ell^C] \geq \mathbb{E}[\pi_\ell^D]$$

$$\delta \geq 1 - \frac{\mathbb{E}[\pi_\ell^C]}{\mathbb{E}[\pi_\ell^D]} \implies \bar{\delta}_{base} = \frac{\phi(q^C)}{\theta_H - \theta_L}$$

To prove point *ii)*, it is straightforward to see that when $\alpha_A = \alpha_B = 1$ equations (18) and (19) are identities, and equation (7) and (17) coincides, meaning that the conditions to sustain $(\sigma_j, \sigma_\ell, \sigma_i)$ are identical in the common and independent debarment list regimes. ■

When the debarment policy enforcement is frictionless, Proposition 2 tells us that the threat faced by the large firm of being debarred when not providing the requested level of quality is sufficient to sustain the cooperative quality level with positive probability. Since the two markets are equal in terms of expected payoffs for the large firm, sharing or not debarment lists is irrelevant with respect to quality provision incentives. This result is dependent on the fact that we are assuming identical buyers. In the next subsection we relax this assumption.

5.2 One Buyer debars cooperating firms with positive probability.

To understand the potential costs of the debarment mechanism, let us now consider the case when enforcement is not frictionless. Since our aim here is to understand the effect of asymmetries among buyers, when it comes to the ability to correctly evaluate the quality exerted by contractors, we fix one player, say buyer A, to the previous case, and let the other have $\alpha_B < 1$.

The first effect, as we have seen in the previous sections, is that now the large firm may be debarred even when cooperating, thus creating an additional implicit cost in playing anticipating to play q^C. Indeed we have the following result.

Proposition 3. *Consider $\alpha_A = 1$, $0 < \alpha_B < 1$. Then*

i) The minimum level of δ required to ensure cooperation in all markets is higher than $\bar{\delta}_{base}$.

ii) Buyer A would prefer not to share debarment lists.

The first point is straightforward and follows directly from applying Lemmas 2 and 4. Both in the independent and shared debarment lists scenarios, the critical level of δ to sustain the grim trigger strategies are given by:

$\bar{\delta}_{indip} = \bar{\delta}_{shared} = \frac{4(\phi(q^C))}{(3+\alpha_B)(\theta_H - \theta_L)} > \bar{\delta}_{base} = \frac{\phi(q^C)}{\theta_H - \theta_L}$

This is again because, with the assumed parameters, condition 17 is binding in the shared debarment list world, and equal to condition 7 of the independent lists world.

The second point tells us that when buyer A can perfectly assess quality, never thus debarring cooperating firms, sharing the debarment list with buyer B is potentially damaging: in the independent world, the threshold level of δ to ensure cooperation, when considering only market A, is lower — coinciding with $\bar{\delta}_{base}$ — than in the shared lists world, where the riskier ICCs calculated in $j = B$ are binding and entail a higher δ, i.e. $\bar{\delta}_{shared}$. In other words, the grim trigger strategies as defined in Section 4 are tied to the δ implied by the "worst" buyer, in terms of ability to correctly identify cooperators.

6 Conclusion

In this paper we have analysed how the institutional design of debarment lists — whether managed independently by each buyer or consolidated into a single, shared one — affects the provision of non-contractible quality in repeated procurement interactions.

With our modelling strategy, we assumed away the potential benefits of sharing debarment information across multiple buyers to focus on the costs. Under a shared debarment list, a non-cooperating contractor in one market is blacklisted by all buyers, effectively forfeiting its future profits in every market. While such a universally harsh penalty could, *prima facie*, bolster relational incentives to maintain high quality, as the cost of deviating is multiplied, in our analysis we show that this device can backfire when buyers randomly misjudge exerted quality. When even cooperating suppliers face erroneous exclusion with positive probability, consolidating debarment lists magnifies the consequences of mistakes: a single misjudgment in one market immediately propagates to all others.

This finding underscores a fundamental trade-off: while a universal ban might be powerful in deterring non-cooperating behaviour, it can also undermine the intended equilibrium outcomes if buyers' evaluations of quality contain noise or if different procurers vary significantly in their ability to assess performance. A buyer with a near-flawless evaluation process may be left worse off by sharing the debarment list with another buyer who more frequently mislabels compliant contractors. Since mistakes in one market spill over into the other, this unintended cross-market "contagion" may discourage firms from delivering high quality, increasing the threshold level of discount factor required for sustaining the desired outcome. Policymakers thus face a dilemma: while universal debarment list may offer a strong deterrent, agencies that evaluate performance with high fidelity may prefer some degree of autonomy rather than fully harmonizing blacklists with less reliable peers.

Empirical work could investigate whether change in regimes from multiple to single debarment lists has had effects in contractors provision of non-verifiable quality.

References

Albano, G. L., Cesi, B., and Iozzi, A. (2017). Public procurement with unverifiable quality: The case for discriminatory competitive procedures. *Journal of Public Economics*, 145:14–26.

Albano, G. L., Cesi, B., and Iozzi, A. (2023). Teaching an old dog a new trick: Reserve price and unverifiable quality in repeated procurement. *Journal of Economics & Management Strategy*, 32(2):377–399.

Bergman, M., Lundberg, S., and Spagnolo, G. (2012). Public procurement and non-contractible quality : evidence from elderly care. Technical report, Umeå University.

Board, S. (2011). Relational Contracts and the Value of Loyalty. *American Economic Review*, 101(7):3349–3367.

Bull, C. (1987). The Existence of Self-Enforcing Implicit Contracts*. *The Quarterly Journal of Economics*, 102(1):147–159.

Buso, M., Berardino, C., Coretti, S., and Turati, G. (2024). Contracting Unverifiable Quality in Healthcare: The Importance of Political Stability for Relational Contracts. *Available at SSRN: https://ssrn.com/abstract=4763713*.

Calzolari, G. and Spagnolo, G. (2009). Relational Contracts and Competitive Screening. *CEPR Working paper no. DP7434*.

Drabkin, D. and Yukins, C. (2022). Congressionally Mandated Study on Contractor Debarments for Violations of U.S. Labor Laws. Technical report, US Department of Defense.

Eggleston, K. (2005). Multitasking and mixed systems for provider payment. *Journal of Health Economics*, 24(1):211–223.

Fuchs, W. (2007). Contracting with Repeated Moral Hazard and Private Evaluations. *American Economic Review*, 97(4):1432–1448.

Hart, O. (1995). *Firms, Contracts, and Financial Structure*. Clarendon Lectures in Economics. Oxford University Press, Oxford, New York.

Kelman, S. (1990). *Procurement and Public Management: The Fear of Discretion and the Quality of Government Performance*. AEI Press.

Levin, J. (2003). Relational Incentive Contracts. *American Economic Review*, 93(3):835–857.

MacLeod, W. B. and Malcomson, J. M. (1989). Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment. *Econometrica*, pages 447–480.

Mailath, G. J., Samuelson, L., Mailath, G. J., and Samuelson, L. (2006). *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, Oxford, New York.

Manelli, A. M. and Vincent, D. R. (1995). Optimal Procurement Mechanisms. *Econometrica*, 63(3):591–620.

Shapiro, C. and Stiglitz, J. E. (1984). Equilibrium Unemployment as a Worker Discipline Device. *The American Economic Review*, 74(3):433–444.

Tirole, J. (1999). Incomplete contracts: Where do we stand? *Econometrica*, 67(4):741–781.

Yukins, C. R. and Kania, M. (2019). Suspension and Debarment in the U.S. Government: Comparative Lessons for the EU's Next Steps in Procurement. *GWU Legal Studies Research Paper No. 2019-39*.

RECENT PUBLICATIONS BY *CEIS Tor Vergata*

When the Going Gets Tough: the Impact of Health Shocks on Divorce
Javier Adrián López Artero, Anna Sanz-de-Galdeano and Daniela Vuri
CEIS *Research Paper*, 597 April 2025

A Stochastic Volatility Approximation for a Tick-By-Tick Price Model with Mean-Field Interaction
Paolo Dai and Paolo Pigato
CEIS *Research Paper*, 596 April 2025

Credit Markets, Corporate Governance and Growth
Emanuele Brancati, Paolo E. Giordani, Maurizio Iacopetta and Raoul Minetti
CEIS *Research Paper*, 595 March 2025

Using Multiple Tools to Enhance Competition in Public Procurement
Clara Calini, Alessandra Catenazzo and Elisabetta Iossa
CEIS *Research Paper*, 594 February 2025

Rage Against the Machine or Humans?
Luca Delle Foglie, Stefano Papa and Giancarlo Spagnolo
CEIS *Research Paper*, 593 February 2025

Firms' Legality and Efficiency: Evidence from Public Procurement
Elisabetta, Chiara Latour
CEIS *Research Paper*, 592 February 2025

Green Ambiguity
Marco Carli
CEIS *Research Paper*, 591 February 2025

Who Decides Matters: Female Representation and Academic Career Advancement
Marianna Brunetti, Annalisa Fabretti and Mariangela Zoli
CEIS *Research Paper*, 590 December 2024

Multivariate Rough Volatility
Ranieri Dugo, Giacomo Giorgio and Paolo Pigato
CEIS *Research Paper*, 589 December 2024

With a Little Help from Nurseries. Childcare Services and Mothers' Employment in Italy
Chiara Puccioni and Daniela Vuri

CEIS *Research Paper*, 588 December 2024

DISTRIBUTION

Our publications are available online at www.ceistorvergata.it

DISCLAIMER

The opinions expressed in these publications are the authors' alone and therefore do not necessarily reflect the opinions of the supporters, staff, or boards of CEIS Tor Vergata.

MEDIA INQUIRIES AND INFORMATION

For media inquiries, please contact Barbara Piazzi at +39 06 72595652/01 or by e-mail at piazzi@ceis.uniroma2.it. Our web site, www.ceistorvergata.it, contains more information about Center's events, publications, and staff.

DEVELOPMENT AND SUPPORT

For information about contributing to CEIS Tor Vergata, please contact at +39 06 72595601 or by e-mail at segr.ceis@economia.uniroma2.it