IDB WORKING PAPER SERIES Nº IDB-WP-01364

When can lotteries improve public procurement processes?

Antonio Estache
Renaud Foucart
Tomas Serebrisky

Inter-American Development Bank
Infrastructure and Energy Sector (INE)

October 2022

When can lotteries improve public procurement processes?

Antonio Estache*
Renaud Foucart**
Tomas Serebrisky***

* Université Libre de Bruxelles
** Lancaster University Management School
*** Inter-American Development Bank

October 2022

Cataloging-in-Publication data provided by the
Inter-American Development Bank
Felipe Herrera Library

Estache, Antonio.
When can lotteries improve public procurement processes? / Antonio Estache, Renaud Foucart, Tomas Serebrisky.
p. cm. — (IDB Working Paper Series ; 1364)
Includes bibliographic references.
1. Government purchasing-Latin America. 2. Goverment purchasing-Caribbean Area. 3. Government auctions-Latin America. 4. Goverment auctions-Caribbean Area. 5. Administrative procedure-Latin America. 6. Administrative procedure-Caribbean Area. I. Foucart, Renaud, 1984- II. Serebrisky, Tomás. III. Inter-American Development Bank. Infrastructure and Energy Sector. IV. Title. V. Series.
IDB-WP-1364

http://www.iadb.org

When can lotteries improve public procurement processes?*

Antonio Estache[†] Renaud Foucart[‡] Tomas Serebrisky[§]

October 25, 2022

Abstract

We study the potential benefits of adding a lottery component to cut the main risks associated with standard negotiated and rule-based auction procurement procedures. We show that adopting a two stage approach in which bureaucrats first negotiate with a small number of bidders to assess their eligibility and, next, rely on a lottery to award the contract reduces corruption risks often observed in negotiated procedures. For rule-based procedures, we show that a "third-price lottery" in which the two highest bidders are selected with equal probability and the project is contracted at a price corresponding to the third highest bid can reduce limited liability, renegotiation, bid rigging and collusion risks.

Keywords: rules, discretion, procurement, lotteries, corruption, auctions **JEL-Code:** D44, D73, H57

*We thank Patrick Legros, Anthony Priolo, David Rietzke, Ian Walker, Anastasiya Yarygina Udovenko as well as participants to seminars and workshop in Brussels, Bergen and Lancaster for useful discussions and comments. Any mistake or misinterpretation is our responsibility only and should not be attributed to any of the institutions we are affiliated with.

[†]Université libre de Bruxelles, Antonio.Estache@ulb.be

[‡]Lancaster University Management School, r.foucart@lancaster.ac.uk

[§]Inter-American Development Bank, tserebrisky@iadb.org

1 Introduction

For Procurement Agencies (PA), identifying and selecting the right contractor to deliver the goods and services expected to meet needs ranging from basic (surgical masks, roads, water pumps,...) to sophisticated (medication, planes, nuclear plants...) keeps being more of a challenge in practice than procurement theory may suggest. The long record of related empirical literature provides detailed evidence on recurring outcome weaknesses such as cost overruns or mistargeting of quality. There is thus a case for improving common (and popular) procurement procedures.

Simplifying somewhat, PA around the world usually pick between two main approaches: (i) discrete and (ii) rule-based procedures. Direct negotiations between a selected firm and a local bureaucrat are an example of the former. In practice, procurement officers such as local bureaucrats and politicians use their specific knowledge of the region and project characteristics to select the most suitable bidder and agree on a way to cover their costs. Auctions in which firms bid to offer discounts over a reserve price to deliver a project are a prime example of the latter.[1]

Somewhere, from the design to the implementation stages, the selection processes under either approach often seem to fail, as evidenced by the number of dysfunctional contract renegotiations (Guasch, 2004; Guasch *et al.*, 2016; Beuve *et al.*, 2018), and the large number of cases in which cost overruns tend to be the norm rather than the exception (Flyvbjerg *et al.*, 2018). The debates on the extent to which these outcomes are the results of project design defects, incompetence of the implementation agencies or simply corruption are recurring topics in academic research and in international agencies (Bandiera *et al.*, 2009; Estache *et al.*, 2009; OECD, 2016; Estache and Foucart, 2018; Fazekas *et al.*, 2021).[2]

This paper suggests that, in environments characterized by various types of governance weaknesses and when the potential pool of contractors to pick from is large enough, adding a lottery component in the selection processes may help related issues including corruption and default risks. These benefits involve a trade-off however. They always come at the cost of allocating the contract to a firm that is, on expectation, less cost-efficient.

The novelty of our approach is to focus on simple procedures, combining features from standard processes and lotteries. In discrete procedures, our paper is to the best of our knowledge the first to introduce lotteries. In rule-based procedures, the main trade-off between complete randomization and auctions was established by Chillemi *et al.* (2009) and Decarolis (2018). The result that the optimal mechanism involves some form of

[1] See for instance Bajari *et al.* (2009) for a detailed comparison

[2] The negative image associated with renegotiation has recently been challenged by Beuve and Saussier (2021) in their study of the French parking sector.

pooling of types can be found in Burguet *et al.* (2012). Our contribution to that literature is thus mostly to provide an implementable procedure featuring such pooling, the third price lottery, and to show that it can perform better than the polar cases of standard auctions and lotteries. We also show that randomization can help against bid rigging, a result that we have not found in previous papers.

The main advantage of discrete procedures is to leverage the knowledge of procurement officer. In a negotiation, they are ideally qualified to identify the best suppliers because they may have superior information over the actual cost and characteristics of a project and/or a better understanding of the firms bidding on the market. They are therefore often more capable to directly select the best firm and commit to reimbursing their costs with a mark-up (the so-called "cost-plus contracts"), without taking the risk of the firm defaulting. Moreover, since empirical evidence suggests that discretion in some cases empowers bureaucrats (Bandiera *et al.*, 2021), the approach can lead these actors to devote more effort to their work, ultimately reducing the cost of the delivered projects.

The main risk with a discretionary procedure is corruption: the more direct power the law puts into the hands of procurement officer in a weak governance and accountability context, the higher the returns for a private firm corrupting them. When this is a concern, adding a final lottery stage to the procedure allows reducing the marginal benefit from corruption. Intuitively, this would work as follows in the simplest situation: a procurement officer selects two firms, and a lottery picks which of the two gets the contract. The bureaucrat's knowledge is reflected in the first stage but is ignored in the second one. In other words, adding a lottery comes at the costs of giving up on the ability to use all the information available to the procurement officer.

Thus, whether the lottery is desirable or not depends on the assessment of the relative importance of the risks of corruption and those associated with the information loss. The challenge is to be able to inform the drivers of this trade-off between efficiency costs and corruption costs. In Section 2, we show that a key parameter to minimize the efficiency costs is the number of bidders: if the pool of potential applicants is high, the efficiency cost of possibly ending up with the second best offer to avoid corruption is likely to be less important than if the pool is of limited size. We also show that the PA can use external audits strategically to induce the procurement officer to make the best choice between a discrete procedure and a discrete lottery.

The main advantages of rule-based procedures are transparency and information revelation: corruption is less likely when the selection criteria are clear, and an auction process leads bidders to reveal some information about their true cost. An important drawback (see for instance Decarolis, 2014) stems from the risk of default from the winner, either due to the "winner's curse" (by offering the highest bid, the winning firm discovers it overestimated the value of the project) or limited liability (the winning firm bids too aggressively, knowing that it will not deliver if the costs are too high). In that context, adding a lottery stage allows screening away some firms more likely to default or renegotiate, as those are bidding more aggressively and are thus more likely – all other things held equal – to offer the best bid. However, this gain comes at the cost of losing some of the information benefits from running an auction.

In Section 3, we show conceptually that lotteries are desirable when the risk of default or renegotiation by the winning firm is high. They offer an alternative to the widespread basic first price auctions observed in procurement processes involving large public goods and service, including those that can be unbundled into smaller projects (Hortaçsu and Perrigne, 2021). Intuitively, lotteries reduce the benefits from being the highest bidder because it does not guarantee winning the contract. As the firms who bear the lowest cost from renegotiating or even defaulting are also more likely to bid aggressively, a lottery allows the PA to award fewer contracts to them. Lotteries are also more useful when the PA puts a sufficiently high weight on the producer surplus. One advantage of auctions over lotteries is to reduce prices and reduce the profit margin of the firms. A PA who wants to support the development of local SMEs in a context in which they compete with more experienced large international suppliers should be less worried about it than one focusing on getting the lowest possible price.

Depending on how important the cost heterogeneity among firms is, the lottery can be among all firms passing a pre-qualification stage for the project, or among a subset of the highest bidders. As an example of the latter, we introduce a "third-price lottery" auction: one of the two top bidders is selected at random and delivers the projects at the discount level corresponding to the third highest bid. In this auction, all firms bid their true willingness to bid. We then use this result to compare it to a standard second-price (or English) auction. We also show that the third price lottery can be used to deter bid rigging and collusion.

The tradeoffs we identify relate to the recent literature on rules vs. discretion in procurement (Coviello *et al.*, 2018; Bosio *et al.*, 2022). These authors find that while discretion would be preferable in a country with high-quality institutions, the risk of corruption makes the use of rules preferable in parts of the world with a lower quality of institutions. One reason that makes lottery procedures attractive is that they alleviate some of the main drawbacks of each procedure (always at a cost), and may increase the scope for using each of them.

The history of explicit lotteries in procurement design is very limited. Uruguay is the only one country we are aware of having carried a large-scale effort to implement a random allocation procedure: the *menor cuantia*. It applies to the procurement of relatively small public works, and all firms passing a pre-qualification phase has the same probability of being selected by a centralized algorithm. Fadic (2020) studies this project, focusing on its main stated objective of promoting the growth of local SME.

This goal is related to our finding in this paper that a complete lottery should only be a tool for a government putting a high value on producer surplus. In this specific case study, the author finds that, indeed, lotteries provide a short-term gain for the winning firms. However, this gain does not translate into a measurable long run benefit such as higher growth revenue or current assets. One thing that might have hindered that potential long run benefit is that all firms had the same probability of winning the contract, regardless of their costs. In the procedures we suggest, the lottery is combined with an easily implementable selection process to ensure that the firm getting the contract is amongst the most cost efficient.

While we are not aware of other explicit lottery procedures, a type of auction in which all firms are selected at random and pay the reserve price - the Average Bid Auction (ABA) – comes close. Indeed, while, in practice, ABAs are equivalent to a lottery, this property is never stated explicitly. The ABA is or was present in public procurement procedures in Chile, China, Colombia, Italy, Japan, Peru, Switzerland or Taiwan among others. In the US, it has been used in the past by the Florida DoT and the New-York State procurement agency (Decarolis, 2018).

One of the variants of the ABA (there are many) could be an easy substitute to the final stage of our third-price lottery. In this ABA procedure, all firms bid on discounts over a reserve price. Then, the firm whose bid is closest to the average is selected and delivers the project using the price defined by her own bid. As shown by Decarolis (2018), such procedures are in practice random. They also lead to very high prices: most of the ABA procedures and in particular the one used in Italy, offer firms incentives to bid at the reserve price.[3] Decarolis (2018) shows, using data from Italy, that this is what firms actually do. The results we find in Section 3 relate closely to one of the main tradeoffs identified by Decarolis (2018): the ABA can be desirable to the extent that it limits renegotiation and the risk of default. However, it comes at a high cost in terms of allocative efficiency (all firms win with the same probability) and consumer surplus (all firms charge the reserve price). We show that our third-price lottery procedure allows finding an intermediary path into the two extremes of auctions with a high risk

[3]To see the logic, assume a reserve price V and $N > 2$ firms. Assume all firms bid a discount of exactly 0. In that case, all firms are selected with equal probability $1/N$ and deliver the project at the highest possible price V. Now, imagine that a firm i acts differently and offers a discount equal to $b_i > 0$. This firm is selected with probability zero since its bid will then be the furthest away from the average, the selection criteria under an ABA procedure. Hence, there is no incentive for anyone to bid more than the minimum.

of default and lotteries or ABAs with a low surplus for the consumer. In practice a third-price lottery with ABA would consist of a first stage of ascending auction, with the two remaining firms submitting an ultimate sealed bid on an additional discount and the winner determined by an ABA.

Section 2 explains how lotteries could maintain the attractive characteristics of discrete procedures while minimizing the risks of corruption. Section 3 does the same in the case of auctions (rule-based selection), in which lotteries help minimizing the costs stemming from the winner's curse and limited liability. Section 4 discusses the main practical implications of the conceptual modelling of sections 2 and 3. We conclude in Section 5.

2 Discretion

In this section, we present a simple model comparing a discrete procedure, a lottery, and an intermediary procedure combining both. We assume all our procedures to apply only to those firms passing the pre-qualification stage. We do not formally model this stage and make the assumption to allow reducing the pool of applicant only to those able to technically deliver the project.

Our discrete procedure works as follows. The PA delegates to a procurement officer with private knowledge of the firms the power to select the most efficient one and reimburse their costs using a "cost-plus" contract where we assume for simplicity the firm mark-up to be zero. We use the term "procurement officer" to describe any local bureaucrat or politician in charge of implementing the procedures we study. There is no risk of default: the true cost is known and observed by the procurement officer and the selected firm. The main drawback is that there is a moral hazard problem between the procurement officer and the PA. Discretion leaves some room for corruption: the procurement officer and the firm could agree on a contract allowing them to report a higher cost than what is necessary and share the surplus.

As we assume cost-plus contracts, all legally acquired surplus of the firm is constant regardless of the cost (and, in our setup, equal to zero). We make the further assumption that the PA does not put any value on the surplus stemming from corruption, be it ultimately extracted by the firm or by the procurement officer. Hence, any weight the PA puts on legally acquired producer surplus is irrelevant to the ranking of procedures. The objective of the PA in our discrete procedures is thus simply to minimize the cost it pays to the contracting firm.

There are N firms who passed a pre-selection phase and are competing to procure a project of value V to the consumers. A firm i has a private cost of delivering the project c_i taken from a continuous log-concave distribution with density $f(x)$ over the interval $[l, h]$.

This cost corresponds to what happens in the good state of the world. However, there is a probability θ that the state of the world is bad, in which case the firm bears an additional cost $D > 0$. This could correspond for instance to bad luck or any external circumstance the firm could not have foreseen. The assumption of log-concavity is satisfied by most commonly used density functions (see Caplin and Nalebuff, 1991; Anderson and Renault, 1999), and we need it to ensure that the presence of more participants to the auction decreases the expected difference between the lowest and the second-lowest cost amongst the participating firms.

2.1 The discrete procedure

The first procedure is fully discretionary. The procurement officer selects the firms with the lowest costs in the good state, denoted $c_{(1,N)}$. With probability $1 - \sigma$, there is no possibility of corruption, and the bureaucrat reimburses this cost plus the additional cost in case the state of the world is bad, an expected reimbursement of $c_{(1,N)} + \theta D$. Moreover, with an exogenous probability σ, the procurement officer is corrupt and enters a corruption pact with the firm in which costs are overestimated to always correspond to the worst-case scenario, so that D is repaid even when the state of the world is good, an expected reimbursement of $c_{(1,N)} + D$. We discuss the incentives of corrupt procurement officers when looking at endogenous procedure selection in Section 2.5.

This form of corruption corresponds to the idea studied in Estache and Foucart (2018) that it is often difficult for an outsider to distinguish corruption from bad luck and incompetence, and that procurement officers and private firms can exploit this information asymmetry to extract a rent. Hence, it is possible for the bureaucrat and the firm to agree to report a high cost in the good state of the world, and to pocket the difference. As the bureaucrat has the power to select a firm, she could for instance make this selection conditional to a bribe corresponding to a share of this extra cost.

A first key assumption we make is that a corruption pact is agreed at the selection stage, when the cost-plus contract is written. After the firm is selected, we do not allow for further corruption. Relaxing this assumption would mean that corruption happens in all configurations, regardless of how the firm is selected. A second key assumption is that the procurement officer always selects the firm with the lowest cost. The bureaucrat maximizing consumer welfare, conditional on extracting rents from corruption, is consistent with this assumption. It could also correspond to the fact that it is easier to conceal higher costs due to corruption when the baseline is low. Relaxing that assumption would imply that the benefits from using the knowledge of a procurement officer disappear, and hence the extent to which a discrete procedure can be desirable in the first place. In such a discrete procedure relying on a simple costs-plus contract, the expected surplus for the

PA, $S_{d,s}$, is therefore equal to

$$S_{d,s} = V - E(c_{(1,N)}) - \theta D - \sigma(1-\theta)D, \tag{1}$$

where V is the value of the project to the PA. The term $E(c_{(1,N)})$ corresponds to the expected value of the first order statistic of the distribution of costs, where the $k-th$ order statistics means the $k-th$ smallest value in the sample. The term D corresponds to the extra cost to the firm in case the state of the world is bad, with probability θ. The last term, $\sigma(1-\theta)D$, corresponds to the extra cost of corruption for the taxpayer. It is increasing in: how often a corruption pact is possible σ and how important the additional cost is in case the state of the world is bad, D. It is also decreasing in how likely it is the high cost would have happened even in the absence of corruption pact θ.

The main tradeoff is thus that a discrete procedure allows selecting the best possible firm, at the risk of this efficiency gain being captured by the firm and procurement officer in the form of corruption.

2.2 The Standard Lottery

A first alternative to the discrete procedure is a standard lottery. In this approach, all the firms who passed the pre-qualification phase are selected with equal probability by a random draw to deliver the project under a costs-plus contract. In that case, the principal bypasses the procurement officer, and reimburses the actual cost incurred by the selected firm. Hence, since the procedure does not select the most cost efficient firm, the expected surplus of the PA, $S_{d,l}$, is equal to

$$S_{d,l} = V - E(c_i) - \theta D. \tag{2}$$

The trade-off between discretion and lottery is thus that the latter solves the moral hazard problem of corruption observed in the discrete procedure. Avoiding corruption comes at the cost of increasing the adverse selection problem due to the firms having private knowledge of their expected cost, and the PA being unable to select the most efficient. It is easy to see that the PA prefers a standard lottery over a discrete procedure whenever the cost of corruption $\sigma(1-\theta)D$ is higher than the efficiency loss from allocating at random $E(c_i) - E(c_{(1,N)})$.

2.3 The Discrete Lottery

We suggest the following alternative to the standard lottery: the rule becomes that the PA asks the procurement officer to select the two most cost-efficient firms to benefit from his/her knowledge of the project. Then, a third party is asked to operate a lottery to select the winning bid out of those two. We label this a "discrete lottery". We argue that this procedure should be sufficient to avoid the kind of corruption described in the basic discrete procedure presented earlier. Indeed, we study corrupt practices in which a firm pays or commit to pay a procurement officer in exchange for getting the contract. Such a corruption pact relies on making sure that all participants benefit from it. It would be too risky for the procurement officer to receive a bribe or commit to receiving a bribe from a firm who, after not being awarded the contract in the lottery stage, would have all the incentives to denounce corruption. Hence, corruption is not feasible. In that procedure, the expected surplus of the PA is therefore equal to

$$S_{d,dl} = V - \frac{E(c_{(1,N)}) + E(c_{(2,N)})}{2} - \theta D, \tag{3}$$

Where

$$\frac{E(c_{(1,N)}) + E(c_{(2,N)})}{2} \tag{4}$$

is the expected value of the average between the first and second order statistic of the distribution of costs amongst N participants.

2.4 Comparing procedures

Ranking the last two procedures in terms of the surplus they generate, we immediately see that $S_{d,dl} > S_{d,l}$ for all values of the parameters. As neither the standard lottery nor the discrete lottery involves corruption, it is indeed never in the interest of the PA to run a full lottery and loose the benefit from the (partial) selection made by the procurement officer as they is nothing to gain in terms of reducing corruption. Whether a discrete lottery yields higher surplus than a discrete procedure depends on whether the efficiency gain from awarding the contract to the firm with the lowest cost with certainty is higher than the cost of corruption.

Proposition 1 *The expected surplus of the PA is higher under a discrete lottery than under a discrete procedure if and only if:*

$$\sigma(1-\theta)D > \frac{E(c_{(2,N)}) - E(c_{(1,N)})}{2}, \tag{5}$$

where the right-hand side decreases with the number of bidders N.

The formal proof is in Appendix A.1. This expression simplifies in the uniform case with $l = 0$ and $h = 1$ to

$$\sigma(1-\theta)D > \frac{1}{2(N+1)}. \tag{6}$$

On the left-hand side is the additional cost from the moral hazard problem in the presence of corruption. On the right-hand side, the additional cost from the asymmetric information problem when the local knowledge of the bureaucrat is not fully exploited. A key parameter for the latter is thus the number of pre-qualified bidders N: if the PA manages to attract a sufficiently large pool of applicants, the risk of having to work with the firm with the second lowest cost instead of the lowest one is low. This risk is however much higher if the number of competing firms is limited.

Another risk from corruption, not captured in our model, is the one of misallocation. If in order to extract rents corrupt procurement officers not only steal money but also allocate the market a less efficient firm - for instance because only a share of the firms is willing to engage in a corruption pact - the case for the discrete lottery becomes stronger.

2.5 Random audit with endogenous procedure choice

The procurement officer may want in some case in good faith to select only one "truly better" firm. In other cases, she may be fine to use a lottery to pick between more than two almost equivalent projects. Indeed, when indifferent, the procurement officer may actually sometimes prefer to randomize, and this is a widespread psychological trait (Agranov and Ortoleva, 2017; Dwenger *et al.*, 2018). The somewhat rigid discrete lottery does not allow for such flexibility. To incentivize further the procurement officer to choose between fewer firms only when it is beneficial to do so, the PA could combine the discrete lottery with a system of semi-random audits adapted from the Brazilian experience (Ferraz and Finan, 2008). In that case, the probability of being audited varies (such as in the study of Zamboni and Litschig, 2018), and decreases with the number of bidders subject to randomization.

Assume that procurement officers are given the choice between selecting a single firm or picking two of them in a discrete lottery. As there is no risk of corruption in the discrete lottery, we look at the case where procurement officers are audited with probability $q > 0$ only when they choose a single firm. Denote by $p \in (\frac{1}{2}, 1)$ the probability that the result of the audit is correct, and by $1 - p$ the probability that it is incorrect. Finally, $F > 0$ is the utility cost for the procurement officers of being found corrupt by the audit. In the Brazilian experience, this cost corresponds to a lower probability for politicians of being re-elected. In the case of bureaucrats, we would similarly expect the result of an audit finding corruption to reduce the probability of being promoted.

We show in Appendix B.1 that if auditing costs are not too low, a policy of random audit with endogenous procedure choice can outperform the discrete lottery. We assume that a share of "bad" procurement officers only care about the possibility to make money, and the risk of being found corrupt. The "good" ones care about the cost of procurement, and also about their reputation. The intuition is that there exists an audit probability such that bad procurement officers always prefer the safety of the discrete lottery and do not actually engage in corruption. Only good procurement officers sometimes pick a single firm, when the cost advantage is sufficiently important.

There are two important caveats with this simplified model. First, there must be a commitment to audit from the PA, perhaps by delegating the task to independent auditors. Indeed, as only honest procurement officers self-select in picking a single firm whenever the PA chooses the optimal auditing probability $q = q^*$, there is a time-inconsistency problem: a bayesian PA would not want to spend money auditing honest procurement officers only. Besides delegation, an additional reason to carry audits on honest procurement officers only would be to deter incompetence (Estache and Foucart, 2018). Second, we have looked at a simplified version of the audit in which there is only a cost when found corrupt, but no benefit of being found not-guilty (as documented by Ferraz and Finan, 2008). Adding such benefits would strengthen the case for endogenous procedure choice as more honest procurement officers would want to pick the lowest cost firm. In the polar case where the benefit of being found non-corrupt is equivalent to the cost of being found corrupt, as $p > \frac{1}{2}$ all honest procurement officers would pick a single firm.

3 Rules (Auctions)

In this section, we compare a second-price auction, a lottery, and a novel intermediary procedure combining both, the "third price lottery." As shown by Burguet *et al.* (2012) in a relating setting, any optimal procedure involves some randomization and pooling of types. The goal of this section is provide some results on a practical way to achieve this objective.

3.1 Preliminaries

Most procurement auctions either use first-price sealed bid auctions or "English" ascending auctions in which sellers are free to increase their bids – expressed as discounts on a reserve price, or ascending auctions with open exit, where the discount increases incrementally until a single firm is left in the room. Such auctions often come in the form of e-procurement in which the procedure happens on an online platform. In our context, these procedures are equivalent, as each seller knows her own cost (Milgrom and Weber, 1982). We use the second-price setting for expositional clarity, but all our results translate directly into the English auction setting. A procurement agency may sometimes prefer sealed-bid auctions for practical reasons, as they involve a single bid by all sellers. One concern with such auctions however is that sellers need to trust the procurement agency. An English auction does not have that problem, as bidders do not need to reveal the maximum they were willing to pay.

The trade-off we identify between the lottery and the auction procedure is similar to Decarolis (2018), who compares first-price and average-bid auctions. As our novel procedure combines features of the lottery and of the second-price auction, we focus our attention to those two polar cases and refer the reader interested in first-price auction results to Decarolis (2018). As in the previous Section, we assume all our procedures apply only to those firms passing the pre-qualification stage.

As before, firm i has a private cost of delivering the project c_i taken from a continuous log-concave distribution with density $f(x)$ over the interval $[l, h]$. With probability θ, the firm bears an additional cost D. The difference with the discrete procedure is that the contract is awarded at a given price, and the costs are only privately observed by the firm. We thus need to look at the incentives to deliver the project whenever the costs are higher than the contracted price, or whether firms are willing to pay the cost of default or the reputation cost of a renegotiation.

We follow Decarolis (2018) and assume a firm is characterised by its type $\omega \in \{L, H\}$. Each firm has a private cost of bankruptcy (reputation, moral, legal), that is equal to either zero (type L) or $\tau > 0$ (type H). A share μ of the firms is of type H. Only the firm knows its own type. There is also a cost for society if the project is not delivered because the selected firm has defaulted, needs a bailout or renegotiation. To keep the model tractable, we treat all events in which the winner needs to renegotiate as "defaults/bankruptcy" and assume a social cost equal to T for the consumers in that event. We assume that all firms are risk neutral and maximize their expected surplus.

The key to solve our auction procedures is to compute the maximum willingness to bid of a firm. We first compute the expected surplus of a firm winning the auction and paying a bid b_i (expressed as a discount from the maximum price V, where V is the value of the project for the consumers) in different cases. The expected surplus of a firm, of either type, who would never default, is:

$$\pi_{H,nd} = \pi_{L,nd} = V - b_i - c_i - \theta D \tag{7}$$

If a firm is of type H, it chooses not to default selectively when the state of the world is bad if the cost of default is sufficiently high,

$$\tau \geq D + b_i + c_i - V. \tag{8}$$

We follow Decarolis (2018) and assume this condition is always satisfied: the reputation cost τ for firms of type H is sufficiently high to ensure they always deliver. We however do not assume that a firm of type L always defaults, an assumption made by Decarolis (2018) in the context of a First Price Auction. The reason is that we want to keep the possibility for a firm with no reputation concern to deliver when, even in the bad state, its cost remains below the contracted price. Note that we only look at the bad state of the world when considering a possible default. The reason is that if a firm defaults in the good state of the world, it also has an incentive to do so in the bad state of the world and would therefore make no profit regardless of the bid.

Lemma 1 *1. The maximum willingness to bid of a firm j of type H and cost c_j is $\hat{b}_{H,j} = V - c_j - \theta D$.*

2. The maximum willingness to bid of a firm i of type L and cost c_i is $\hat{b}_{L,i} = V - c_i$.

The formal proof is in Appendix A.2. The lower willingness to bid of the firm of type H comes from the fact that they internalize the risk of being in a bad state of the world, while the type L firm knows that in the worst case they always keep the possibility of defaulting and getting zero profit. Given that this setup allows for the possibility of sellers receiving a rent that does not come from corruption, we also need to look at the objective function of the PA. We assume a function linear in consumer surplus and firm profit (Baron and Myerson, 1982),

$$W = CS + \alpha PS \tag{9}$$

with a weight 1 on consumer surplus and $\alpha \in (0, 1)$ on producer surplus. We consider three possible and easily comparable procedures to illustrate the pros and cons of a lottery in this setting: a second-price auction, a third-price lottery combining an auction with a random procedure, and a standard lottery. We represent these three possibilities in Figure 1. The first column shows the maximum willingness to bid of four firms, A, B, C and D, with A the firm willing to offer the highest discount over the reserve price. We do not consider at this point the type of the firms and simply take these valuations as given.

3.2 Ascending auction with open exit or second-price sealed bid auction

Our first rule-based procedure is a standard Second Price ("Vickrey") Auction. Firms submit written bids, without knowing what bids the other firms make. The highest bidder wins but the price paid is the second-highest bid. This auction is equivalent in our setting to the perhaps more familiar "open-exit" sealed bid auction, in which discounts increase and firms not willing to offer such a discount leave the procedure until there

Figure 1: Three different rule-based procedures

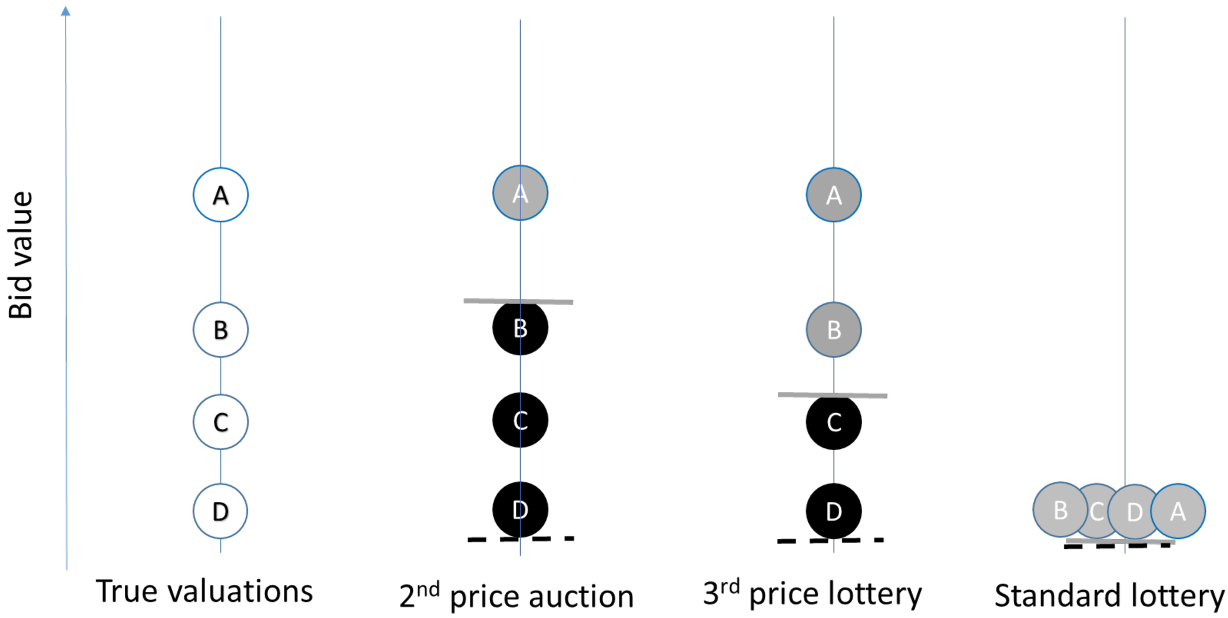

is only one firm left. Another similar procedure in our case is the "English" auction in which firms increase their bid discount by small amount until no one is willing to bid above the highest bid.

We illustrate the procedure in the second column of Figure 1. The horizontal (dashed) line at the bottom represents the lowest bid V the PA would accept, corresponding to a discount of zero. The horizontal grey line corresponds to the equilibrium bid at which a contract is agreed. In this case, it is well known that bidding one's maximum willingness to bid is a weakly dominant strategy, so that in the unique symmetric Bayesian Nash equilibrium the winner is the firm with the highest willingness to bid (in grey), and the contracted price corresponds to the second highest willingness to bid. To see this, remember that the winner of the second price auction is the highest bidder, but she only has to pay the second highest bid. It is thus an equilibrium strategy for everyone to state his or her true valuation. For instance, firm A has nothing to gain from increasing or lowering its bid as long as it is above B. And it would lose out by bidding below B and not being selected. B has nothing to gain from increasing her bid below A, as she would still lose, and would lose out by bidding above A, as it would have to deliver the project for with a discount higher than her highest willingness to bid. To measure the expected consumer surplus and profit of the selected firm, we need not only to look at the distribution of the costs, but also at the distribution of the bids and at the likelihood of default. The existence of the more aggressive bidders of the low-reputation cost type L has the impact of making it more likely that the project does not happen. In such a case, the social cost for the consumer is equal to T. However, the existence of aggressive bidders also increases the expected value of the second highest bid, to be paid by the winner. Aggressive bidders guarantee a higher contracted discount over the reserve price, but increase the probability that consumers bear the cost of default. In consequence, second-price auctions are more of a problem when T is high.

3.3 The Third Price Lottery

We introduce a novel procedure we denote as the "third price lottery." It is similar to a second price auction, except for the fact that the two highest bidders (in grey in Figure 1) are selected with equal probability by a lottery. The contract is then agreed at a discount corresponding to the third highest bidder. A useful property of this procedure is that the firms' bids are identical to the second-price auction. The third-price lottery is also equivalent to an open exit ascending or an "English" auction in which the price is raised continuously until all but two bidders have left. The winner is then selected at random among these two, and the discount corresponds to the discount at which the third highest bidder left. It is also reminiscent of the anglo-dutch auction (Klemperer, 2002): an ascending auction followed by a final sealed bid auction amongst the last two bidders. In our procedure, the first stage is similar but the second stage is a lottery.

Lemma 2 *In the third price lottery, in the unique symmetric perfect Bayesian equilibrium, each firm i bids the highest discount $\hat{b}_i$ they are willing to offer.*

The formal proof is in Appendix A.3. The result holds for reasons similar to the second price auction. We illustrate our reasoning in the third column of Figure 1. Neither A nor B could do better by bidding above or below their valuation (but above C), and both would lose out by bidding below C and being selected with probability zero. Both C and D would lose out from bidding above B. As we see from the grey line, this procedure yields a lower contracted discount than the auction without a lottery. This lower discount yields the following trade-off: On the one hand, a third price lottery implies that the project is in expectation contracted with a lower discount than in the second-price one. Hence, conditional on the firm actually delivering the project, the third-price lottery decreases consumer surplus. On the other hand, the fact that the contracted discount is lower in the third price lottery makes it more likely that a firm with high reputation concerns is included in the lottery (as those bid less aggressively). It also makes it more likely that a firm with no reputation concern ends up in position of nonetheless delivering in the bad state, simply because the difference between its bid and the third highest one is sufficiently high to compensate for the extra production cost.

3.4 The Standard Lottery

The last procedure is a lottery among all participants. We assume that the highest possible cost amongst those firms who passed the pre-selection phase is sufficiently low for their maximum willingness to bid to be positive (else, they would have no interest in joining the qualification phase in the first place). The selected firm signs a contract based on her stated bid. As for the Average Bid Auction (ABA), the equilibrium involves all firms bidding for a zero discount, and all of them selected with equal probability. The equilibrium discount is therefore lower than in any of the two other procedures. The benefit is that the likelihood of a firm with reputation concerns (of type H) being selected is equivalent to its share in the population, while all other procedures gave an advantage to the more aggressive bidders of type L. Moreover, the zero discount ensures that as long as the cost is not too high, $V \geq D + c_i$, even firms of type L do not default when the state of the world is bad.

3.5 Comparing outcomes

There are two reasons why the standard lottery yields the lowest contracted discounts over the reserve price, and the third-price lottery yields lower discounts than the second-price auction. The first is bad news: in a lottery, the PA loses at least part of the information benefit from the standard auction setting. In a second-price auction, the expected rent of the winning firm is the difference between her maximum willingness to bid and the second highest bid. In a full lottery, it is equal to her maximum willingness to bid. The second is good news: our two lottery procedures discard some unrealistically high bids. In terms of surplus generated for the consumer, the contracted bid is indeed not the entire story. If the selected firm has no reputation concern (type L) and therefore defaults selectively, the project is only completed with probability $1 - \theta$. With probability θ, it is not completed, and society incurs a cost T, unless the contracted bid is a sufficiently low discount to ensure even a firm of type L delivers in the bad state of the world. This means that the highest price we observe in lotteries in simply more realistic: it reflects the cost of actually delivering the project, in contrast to the often-empty promises of the winner of an auction.

We first provide some general results:

Proposition 2 *1. The discount at which the project is contracted is highest in the second-price auction, then in the third-price lottery, then in the standard lottery.*

2. The probability that the firm winning the contract has no reputation concerns (type L) is highest in the second-price auction, then in the third-price lottery, then in the *standard lottery.*

The proof is in Appendix A.4 and follows directly from the equilibrium strategy under the different procedures. Intuitively, the lottery is attractive in cases where the cost of default for society T is very high, and so is the risk of default. The latter can come directly from a high probability of being in the bad state of the world $1-\theta$ or a low share of firms with reputation concerns H. It could also come indirectly from the size of the rent naturally occurring in the other procedures. If the rent of the winner is sufficiently high in an auction setting, be it because of a smaller number of bidders or a larger dispersion of the costs, the risk of default is lower.

To illustrate those ideas, we report in Appendix C.1 the outcome of 10 simulations of each time 1000 procurement outcomes with six bidders.[4] We start by assuming costs are drawn from a uniform distribution, and use the following parameter values: $l = 0, h = V = T = 1, \theta = \mu = 1/2, D = 2/5$ and τ sufficiently high. This means that default has an additional cost equal to the value of the project, that half of the firms have reputation concerns, and that a cost overrun worth 40% of the value of the project happens with probability $\frac{1}{2}$. We find that for those values, the second-price auction yields the highest contracted discount, but the third price lottery is the procedure delivering the highest consumer surplus. The standard lottery offers the highest producer surplus. We also look at the case with four bidders, keeping all parameters the same as in the main simulations. This case is more favourable to the second-price auction, for two reasons. As compared to the standard lottery, the difference in the probability of default is lower than in the presence of six bidders. The expected information rent is indeed higher with fewer bidders, and it therefore becomes more likely that a bidder with no reputation concerns delivers in the bad state of the world because it can still make positive profit. One of the advantages of the third-price lottery also becomes less important: the third highest discount is expected to be very low, so that this procedure becomes much closer to a standard lottery. Looking at the total welfare, in this example, the second-price auction is preferred when the PA only cares about consumer surplus ($\alpha = 0$), the third price lottery is preferred for intermediate values of α, and the standard lottery when consumer surplus is valued the same as producer surplus ($\alpha = 1$).

More generally, which procedure is socially desirable depends on the value of the parameters and on how much the PA values producer and consumer surplus. We summarize in Appendix C.2 the results of the simulations for different values of the extra production cost to the firm in the bad state of the world (D), of the social cost for the consumer of the winning firm defaulting or renegotiating (T), and of the weight the PA puts on producer surplus α. As we would expect, when the damage for the firm and

[4]The ten separate simulations yield the same ranking for the different procedures on each criterion.

for the society of being in the bad state of the world are low, the standard second price auction becomes more desirable. The standard lottery is preferred for the largest weight on producer surplus, in particular when D and T are high. The third-price lottery is preferred in all the intermediary cases, in particular when the PA puts a strictly positive but moderate weight on producer surplus.

3.6 Collusion and bid rigging

An important limitation of our study of rule-based procedures is that we make the assumption that there is no collusion amongst bidders. Such cartels are however prevalent in practice (see for instance Kawai and Nakabayashi, 2022), and may limit the benefits from an auction procedure.

In Appendix B.2, we solve for the incentives to collude in the second price auction and the third price lottery in an infinitely repeated game. We assume that firms publicly observe each other costs, and that those costs are drawn at random in every period. Firms maximize their expected surplus in collusion by taking turns and letting the most cost-efficient win the contract at the smallest possible discount. Collusion can be sustained with the threat of punishment: any deviation from colluding behaviour leads to competition in perpetuity.

One advantage of a third-price lottery in deterring collusion is similar to a result found by Chassang and Ortner (2019) in the case of a PA setting a minimum price in auctions: by guaranteeing a higher surplus to bidders, the third price lottery diminishes the possibilities for the cartel to punish the deviator. Thus, third price lotteries could also be a useful tool to deter bid rigging cartels. However, one inconvenient of the third-price lottery in terms of deterring collusion is that a deviator outbidding other cartel members is not guaranteed to win the contract.

4 Discussion and alternative procedures

In this section, we review some risks, limitations and possible objections to lottery-based procedures, as well as some suggestions to mitigate these concerns.

4.1 Efficiency and subcontracting

We have shown that one of the main drawbacks of lotteries is that they fail to select the firm with the lowest cost. Branzoli and Decarolis (2015) however show that a mechanism failing to pick the most efficient firm does not imply the delivery of the project will not be efficient. As lotteries influence the nature of the firms selected ex-ante, they also affect the incentives to use sub-contracting ex-post. Assume the lottery selects a high cost firm: this firm could choose to subcontract all parts of the projects to the most efficient firm and pocket the difference. There is no loss of efficiency as compared to the case in which the procedure selects the best firm, but there is a clear question of redistribution. If subcontracting is perfectly efficient and the PA has no problem with the distributive outcomes it generates, standard lotteries are always optimal. Else, our intermediate third price lottery and discrete lottery procedures reduce this redistribution problem.

4.2 The value of control

One possible risk for the implementation of lotteries is the decreased benefits from the empowerment of procurement officers in discrete procedures Bandiera *et al.* (2021): If instead of selecting the winning firm, they end up selecting two or more, bureaucrats may feel that their authority carries a lower weight. This feeling can be in part alleviated by giving bureaucrats formal control on all the stages of the procedure. This implies control of the ex-post negotiations - after the lottery has selected the winning firm, but also some form of control on running the public lottery allocating the project to one firm. Experimental research has shown that individuals have an intrinsic preference for control, even if it is in practice meaningless (Bartling *et al.*, 2014; Owens *et al.*, 2014), and this preference extends to control over a lottery (Bouacida and Foucart, 2022).

4.3 The social acceptability of lotteries

Another possible difficulty in the implementation of lottery-based procedures is the reluctance of individuals to see their fate decided explicitly randomly (Bouacida and Foucart, 2022). Experimental evidence shows that such concerns can be alleviated by making the lottery less explicit, and letting it follow the rituals of reason (Elster, 1989). This kind of preference could explain why lotteries are so rare in practice, while the equivalent Average Bid Auction is so widespread: the latter does not look like a lottery, even if it is in practice equivalent to one.

If a PA aiming at implementing lotteries faces such concerns, it could look at alternatives that would be equivalent in practice to a lottery, without involving a formal randomization. One possible such alternative would be the use of an external, anonymous, engineer assessing the projects. In the case of auctions, this means providing the kind of “official estimate” of what a sustainable discount would look like typically used to screen out unrealistic bids. The procedures would then select the bid closest to this evaluation. In the case of discretion, the PA would select a very small subset of anonymized projects and let the final selection to an anonymous expert, perhaps from an independent supra national institution. Evidence from the allocation of research grants by expert panels shows indeed that that they tend to rank projects who meet a certain threshold (similar to pre-qualification in our setting) in a way that is undistinguishable from a lottery (Graves *et al.*, 2011; Pier *et al.*, 2018).

4.4 An hybrid procedure: scoring rules

Finally, our dichotomy between discrete and rule-based procedures does not consider the many contracts using features of both. In cases in which the PA does not have enough information to directly quantify important parameters, such as dimensions of quality unobservable ex-ante or the reputation of bidders for keeping their promises, many agencies rely on scoring auctions . With this procedure, the firm chosen to procure the good or the service is selected through a scoring of specific price and quality dimensions, usually weighted to reflect their relative importance in the overall assessment (Asker and Cantillon, 2008, 2010). The final score used to identify the desirable supplier is simply the weighted sum of the score assigned to each criteria (the quality criteria having often already been part of the technical pre-qualification phase). The difficulty is that the scores and the weights assigned to each criteria can be quite subjective. This is why procedures based on scoring rules are a hybrid between discretionary procedures and rule based procedures. If corruption is a risk, a firm could try to influence the weight put on each criterion.

In that case, the PA could build lottery procedure similar in spirit to our discrete lottery to randomize this weight within certain limits. An additional benefit of such a randomization is to encourage firms to offer the best project without focusing too much on the always-imperfect way its value is measured. If limited liability is a concern, randomizing amongst the two top bidders would have an effect similar to our third price lottery.

In the transport sector for instance, rule-based procedure that do not directly aim at selecting the highest or lowest bidder have been followed since the 1990s in a wide range of countries (Asian Development Bank, 2018).

A popular example among practitioners concerns routine rehabilitation and maintenance projects. CREMA contracts (Contratos de Reabilitacao e Manutencao in Brazil or Contratos para Rehabilitacion y Mantenimiento in the rest of Latin America, i.e. Rehabilitation and Maintenance Contracts) adopted since the late 1990s in Latin America have been delivering very positive outcomes in terms of efficiency and other performance measures. They require firms to maintain relatively short stretches (30 to 50km) of roads usually for a period of 3 to 5 years.[5] What makes them stand out in the procurement processes of the sector is that contractors selection can be based both on both price and non-price criteria (Stankevich *et al.*, 2009). Bidders can be either pre-qualified or post-qualified and short-listed based on the joint evaluation of technical and cost dimensions. In practice, the pragmatic "best value" approach in selecting a winner has often delivered better and more sustainable road maintenance than standard "low bid" approaches because it did a better job at accounting for multiple dimensions, including various market and management characteristics of the sector. The scores associated with these criteria can be weighted to reflect policy preferences.

The challenge is, however, still to minimize rather than simply reducing the efficiency cost of the institutional weaknesses characterizing each country and each market in which the procurement process is implemented. Despite their impressive achievements, this is a dimension that CREMAs have not always fully internalized. And this is where the lotteries discussed here could make a difference.

5 Conclusions

This paper identifies situations in which standard procurement procedures could be improved by the addition of a lottery component. The partial lotteries we suggest have the advantage of reducing the main risk from discrete and rule-based procedures, while keeping some of their informational advantage. This is in contrast with standard lottery procedures, in which the PA needs to give the entire surplus to the producers.

Putting the results of the paper in perspective with the recent literature on the choice between discrete and rule-based procedures offers us a rough guide of the circumstances in which some form of procurement lotteries is desirable. Table 1 summarize those insights in the case of discrete procedure. Based on the result of Proposition 1, we know that the two key elements are the number of potential bidders and the quality of institutions: how corruptible the procurement officers are.

When the pool of firms is small and corruption is low, a standard discrete procedure is more valuable, as it allows extracting the local knowledge of the bureaucrat without too much risk. With a larger pool, the information cost from asking the bureaucrat to provide the name of two firms instead of a single one gets smaller, so that a discrete lottery starts becoming an interesting option. The numbers of firms taking part to a procedure is not exogenous: one of the advantages of the lottery procedures is precisely to make bidding attractive to more firms. The case of a large pool of bidders with high risk of corruption is the most obvious example where a discrete lottery is the optimal choice. Finally, the

[5] Lancelot (2010), offers a detailed description of the Brazilian experience

Table 1: Discrete procedures

	Low corruption	**High corruption**
Small Pool	Discrete P	Consider rule-based
Large Pool	Discrete P / Discrete L	Discrete L

Table 2: Rule-based procedures

	Low risk of default	**High risk of default**
Small Pool	2nd price A	Consider discretion
Large Pool	2nd price A /3rd price L	3rd price L/Standard L

combination of a small pool of bidders and high risk of corruption makes both options less appealing, and a PA may consider instead switching to a rule-based procedure such as an auction.

Table 2 looks at the choice of a rule-based procedure. We know from Proposition 2 that the main trade-off is between selecting the most cost-efficient bidder, and making sure that the selected bidder actually delivers even when the state of the world is bad. When the risk of default – the cost overrun in the bad state of the world and the social cost of the project not being delivered or being delayed - is low, the latter risk is small even in an auction procedure. When the risk of default is higher, the lottery becomes more appealing. When the pool of bidders is small, the information gain from a second-price auction is maximal. Hence, a classic second-price auction is the most desirable in this situation. A larger pool of bidders starts making the case for a third price lottery, as the difference between the second and the third highest willingness to bid starts decreasing. With a large pool, for intermediate costs of default, the third price lottery is desirable. When this cost is even higher, in particular if the PA puts enough weight on the producer surplus, a standard lottery may prove being the optimal choice. Finally, the case with a small pool of bidders and a high risk of default is the one where rule-based procedures are the least appealing, and the PA may consider switching to a discrete procedure instead.

References

Agranov, M. and Ortoleva, P. (2017). Stochastic choice and preferences for randomization. *Journal of Political Economy*, **125** (1), 40–68.

Anderson, S. P. and Renault, R. (1999). Pricing, product diversity, and search costs: A bertrand-chamberlin-diamond model. *The RAND Journal of Economics*, pp. 719–735.

Asian Development Bank, . (2018). *Guide to performance-based road maintenance contracts.* Asian Development Bank.

Asker, J. and Cantillon, E. (2008). Properties of scoring auctions. *The RAND Journal of Economics*, **39** (1), 69–85.

— and — (2010). Procurement when price and quality matter. *The Rand journal of economics*, **41** (1), 1–34.

Bajari, P., McMillan, R. and Tadelis, S. (2009). Auctions versus negotiations in procurement: an empirical analysis. *The Journal of Law, Economics, & Organization*, **25** (2), 372–399.

Bandiera, O., Best, M. C., Khan, A. Q. and Prat, A. (2021). The allocation of authority in organizations: A field experiment with bureaucrats. *The Quarterly Journal of Economics*, **136** (4), 2195–2242.

—, Prat, A. and Valletti, T. (2009). Active and passive waste in government spending: evidence from a policy experiment. *American Economic Review*, **99** (4), 1278–1308.

Baron, D. P. and Myerson, R. B. (1982). Regulating a monopolist with unknown costs. *Econometrica: Journal of the Econometric Society*, pp. 911–930.

Bartling, B., Fehr, E. and Herz, H. (2014). The intrinsic value of decision rights. *Econometrica*, **82** (6), 2005–2039.

Beuve, J., Lannier, A. L. and Squeren, Z. L. (2018). Renegotiating ppp contracts: Opportunities and pitfalls. In *The Economics of Public-Private Partnerships*, Springer, pp. 135–162.

— and Saussier, S. (2021). Renegotiations and renewals of public contracts. *Review of Industrial Organization*, **59** (3), 461–482.

Bosio, E., Djankov, S., Glaeser, E. and Shleifer, A. (2022). Public procurement in law and practice. *American Economic Review*, **112** (4), 1091–1117.

Bouacida, E. and Foucart, R. (2022). Rituals of reason: A choice-based approach to the acceptability of lotteries in allocation problems.

Branzoli, N. and Decarolis, F. (2015). Entry and subcontracting in public procurement auctions. *Management Science*, **61** (12), 2945–2962.

Burguet, R., Ganuza, J.-J. and Hauk, E. (2012). Limited liability and mechanism design in procurement. *Games and Economic Behavior*, **76** (1), 15–25.

Caplin, A. and Nalebuff, B. (1991). Aggregation and imperfect competition: On the existence of equilibrium. *Econometrica: Journal of the Econometric Society*, pp. 25–59.

Chandra, N. K. and Roy, D. (2001). Some results on reversed hazard rate. *Probability in the Engineering and informational Sciences*, **15** (1), 95–102.

Chassang, S. and Ortner, J. (2019). Collusion in auctions with constrained bids: Theory and evidence from public procurement. *Journal of Political Economy*, **127** (5), 2269–2300.

Chillemi, O., Mezzetti, C. *et al.* (2009). Procurement under default risk: auctions or lotteries? *Unpublished paper). https://warwick. ac. uk/fac/soc/economics/staff/cmezzetti/papers/occm15decsub. pdf.*

Coviello, D., Guglielmo, A. and Spagnolo, G. (2018). The effect of discretion on procurement performance. *Management Science*, **64** (2), 715–738.

Decarolis, F. (2014). Awarding price, contract performance, and bids screening: Evidence from procurement auctions. *American Economic Journal: Applied Economics*, **6** (1), 108–32.

— (2018). Comparing public procurement auctions. *International Economic Review*, **59** (2), 391–419.

Dwenger, N., Kübler, D. and Weizsäcker, G. (2018). Flipping a coin: Evidence from university applications. *Journal of Public Economics*, **167**, 240–250.

Elster, J. (1989). *Solomonic judgements: Studies in the limitation of rationality.* Cambridge University Press.

Estache, A. and Foucart, R. (2018). The scope and limits of accounting and judicial courts intervention in inefficient public procurement. *Journal of Public Economics*, **157**, 95–106.

—, Iimi, A. and Ruzzier, C. A. (2009). Procurement in infrastructure: What does theory tell us? *World Bank Policy Research Working Paper*, (4994).

Fadic, M. (2020). Letting luck decide: Government procurement and the growth of small firms. *The Journal of Development Studies*, **56** (7), 1263–1276.

Fazekas, M., Blum, J. R. *et al.* (2021). Improving public procurement outcomes: Review of tools and the state of the evidence base. *Policy Research Working Paper Series*, (9690).

Ferraz, C. and Finan, F. (2008). Exposing corrupt politicians: the effects of brazil's publicly released audits on electoral outcomes. *The Quarterly journal of economics*, **123** (2), 703–745.

Flyvbjerg, B., Ansar, A., Budzier, A., Buhl, S., Cantarelli, C., Garbuio, M., Glenting, C., Holm, M. S., Lovallo, D., Lunn, D. *et al.* (2018). Five things you should know about cost overrun. *Transportation Research Part A: Policy and Practice*, **118**, 174–190.

Graves, N., Barnett, A. G. and Clarke, P. (2011). Funding grant proposals for scientific research: retrospective analysis of scores by members of grant review panel. *Bmj*, **343**.

Guasch, J. L. (2004). *Granting and renegotiating infrastructure concessions: doing it right*. World Bank Publications.

—, Benitez, D., Portabales, I. and Flor, L. (2016). The renegotiation of public private partnerships contracts (ppp): An overview of its recent evolution in latin america. *Revista Chilena de Economía y Sociedad*, **10** (1), 42–63.

Hortaçsu, A. and Perrigne, I. (2021). Empirical perspectives on auctions. In *Handbook of Industrial Organization*, vol. 5, Elsevier, pp. 81–175.

Kawai, K. and Nakabayashi, J. (2022). Detecting large-scale collusion in procurement auctions. *Journal of Political Economy*, **130** (5), 1364–1411.

Klemperer, P. (2002). What really matters in auction design. *Journal of economic perspectives*, **16** (1), 169–189.

Lancelot, E. (2010). Performance based contracts in the road sector: towards improved efficiency in the management of maintenance and rehabilitation-brazil's experience.

Li, X. (2005). A note on expected rent in auction theory. *Operations Research Letters*, **33** (5), 531–534.

Milgrom, P. R. and Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, pp. 1089–1122.

OECD (2016). Preventing corruption in public procurement.

Owens, D., Grossman, Z. and Fackler, R. (2014). The control premium: A preference for payoff autonomy. *American Economic Journal: Microeconomics*, **6** (4), 138–61.

Paul, A. and Gutierrez, G. (2004). Mean sample spacings, sample size and variability in an auction-theoretic framework. *Operations Research Letters*, **32** (2), 103–108.

Pier, E. L., Brauer, M., Filut, A., Kaatz, A., Raclaw, J., Nathan, M. J., Ford, C. E. and Carnes, M. (2018). Low agreement among reviewers evaluating the same nih grant applications. *Proceedings of the National Academy of Sciences*, **115** (12), 2952–2957.

Stankevich, N., Qureshi, N. and Queiroz, C. (2009). *Performance-based contracting for preservation and improvement of road assets*. Tech. rep., The World Bank.

Zamboni, Y. and Litschig, S. (2018). Audit risk and rent extraction: Evidence from a randomized evaluation in brazil. *Journal of Development Economics*, **134**, 133–149.

Appendix

A Proofs

A.1 Proof of Proposition 1

The condition $\sigma(1-\theta)D > \frac{E(c_{(2,N)})-E(c_{(1,N)})}{2}$ is straightforward from taking the expressions of$S_{d,dl}$ and $S_{d,d}$. We see immediately that the left-hand side increases with σ and D and decreases with θ. To see that the right-hand side is decreasing in N, we use the fact that

$$E(c_{(2,N)}) - E(c_{(1,N)}) = N\int_l^h F(x)(1-F(x))^{N-1}dx,$$

Using the density function of the first and second order statistics (see Paul and Gutierrez, 2004, p.105). By Proposition 2.3 in Li (2005), we know that a condition for an increase the number of participants to reduce the cost difference,

$$E(c_{(2,N)}) - E(c_{(1,N)}) > E(c_{(2,N+1)}) - E(c_{(1,N+1)}),$$

is that the distribution of the costs has a decreasing reversed hazard rate (DRHR), a property shared by all log-concave distributions (see for instance Result 2.2 in Chandra and Roy, 2001).

A.2 Proof of Lemma 1

The maximum willingness to bid of a firm j of type H and cost c_j is the amount that yields zero profit on expectation. It therefore has to consider both states, and is a discount equal to

$$\bar{b}_{H,j} = V - c_j - \theta D. \tag{10}$$

If a firm is of type L, it chooses not to default selectively when the state of the world is bad only if the additional cost, given the contracted bid, yields a positive profit

$$0 \geq D + b_i + c_i - V. \tag{11}$$

Denote by $\hat{b} = V - c_i - D$ the value of b_i such that for all $b_i > \hat{b}$ the condition is not satisfied, so that a firm of type L defaults selectively. Denote by $\pi_{L,d} = (1-\theta)(V - b_i - c_i)$ the expected surplus in case of partial default, the expected surplus of a firm of type L is thus,

$$\pi_L = \begin{cases} \pi_{L,nd} & \text{for } b_i \leq \hat{b} \\ \pi_{L,d} & \text{for } b_i > \hat{b} \end{cases} \tag{12}$$

Replacing $\hat{b}$ by its value, we see that for all $b_i \leq \hat{b}$, $\pi_{L,nd} > 0$, so that there is no bid below the threshold that yields zero or a strictly negative profit. Our only candidate for a maximum willingness to bid thus corresponds to the case $b_i > \hat{b}$ and is $\bar{b}_{L,i} = V - c_i > \bar{b} = V - c_i - D$.

A.3 Proof of Lemma 2

We establish that truth-telling – bidding one's maximum willingness to bid - is a weakly dominant strategy. Consider a firm i bidding b_i.

1. Assume first b_i is not part of the two highest bids. The firm is then selected with probability zero and makes zero profit. Any lower bid does not affect the probability of winning. Any higher bid could mean the firm wins if it offers a bid higher than the current second highest bid. In that case, the current second highest bid b' would become the third highest and therefore the contracted discount under the procedure. If $b' > \bar{b}_i$, any discount higher than b' yields strictly negative profit, and it is not in the interest of the firm to do so. If $b' \leq \bar{b}_i$, all bids above b' yields identical positive profit. This includes the maximum willingness to bid $\bar{b}_i$.

2. Assume now b_i is part of the two highest bids. If the third highest bid is $b' > \bar{b}_i$, firm i makes strictly negative expected profit and could get a higher surplus by bidding strictly less than the third highest bid. All $b_i < b'$ then yield identical zero profit, including the maximum willingness to bid $\bar{b}_i$. If the third highest bid is $b' \leq \bar{b}_i$, all bids above b' yields identical positive profit. This includes the maximum willingness to bid $\bar{b}_i$.

3. Similar to the standard second-price auction, this strategy thus constitutes the unique symmetric Perfect Bayesian Equilibrium. If all firms are expected to bid either less or more than their maximum willingness to bid with strictly positive probability, we know by the above reasoning that truth-telling becomes a strictly profitable deviation. There is no strictly profitable deviation if all firms are expected to bid exactly $\bar{b}_i$.

A.4 Proof of Proposition 2

1. In the second price auction and in the third price lottery, firms bid $\hat{b}_i$. As the underlying cost structure is the same, and as the discount contracted in the third price lottery is the third highest bid, while it is the second highest in the second-price auction, the contracted discount is always higher in the latter for a given

realization of the costs. The result for the lottery is straightforward as the discount is always zero.

2. Bidders of type L bid more aggressively. Hence, we can always expect the probability that the highest bidder is of type L to be higher than the probability that the second highest bidder is of type L, and the latter to be higher than the third highest bidder is of type L. The same logic applies until the lowest bid.

B Extensions

B.1 Random audit with endogenous procedure choice

There are two types of procurement officers. A share $1-\sigma$ are honest, motivated by the outcome of the procurement process, but also by the personal cost of being possibly found corrupt by an audit. Hence, an honest procurement officer prefers to pick the lowest-cost firm whenever the risk of being found corrupt by mistake is low enough to compensate the social benefit of choosing the best firm,

$$c_{(1,N)} + q(1-p)F \leq c_{(2,N)} \Leftrightarrow q \leq \frac{c_{(2,N)} - c_{(1,N)}}{F(1-p)}. \tag{13}$$

A share σ of corrupt procurement officers cares about the possible gains from corruption and the personal cost of being found guilty, but not about the cost of procurement. These procurement officers only pick a single firm if they can make enough money from corruption: they do not care about the procurement itself and would therefore never benefit from picking the lowest cost firm and bearing the risk of being found corrupt by mistake. A corrupt procurement officer therefore picks the discrete lottery whenever the benefits of corruption $(1-\theta)D$ are lower than the cost of being found corrupt,

$$(1-\theta)D \leq qpF \Leftrightarrow q \geq \frac{(1-\theta)D}{Fp} = q^*. \tag{14}$$

If the main decision the PA can take is the probability of auditing q, there exists a value q^* such that corrupt procurement officers systematically choose to pick the discrete lottery. Only honest procurement officers pick a single firm, whenever

$$c_{(2,N)} - c_{(1,N)} \geq q^* F(1-p) = \frac{(1-\theta)D(1-p)}{p}. \tag{15}$$

Denoting by $g(x)$ the distribution of $x = c_{(2,N)} - c_{(1,N)}$, the social cost, as compared to a first-best situation in which all procurement officers would be non-corrupt and pick the lowest-cost firm is equal to

$$\sigma \frac{E(c_{(2,N)}) - E(c_{(1,N)})}{2} + (1-\sigma) \int_0^{\frac{(1-\theta)D(1-p)}{p}} \frac{x}{2} g(x)dx. \tag{16}$$

When comparing with a compulsory discrete lottery, the outcome is similar for dishonest procurement officers, but the misallocation is lower for honest procurement officers, and it only happens precisely when the difference between the lowest and the second-lowest cost is small. Denoting by γ the cost of auditing, a random audit with endogenous procedure choice performs better than a compulsory discrete lottery whenever

$$\int_{\frac{(1-\theta)D(1-p)}{p}}^{\infty} \frac{x}{2} g(x)dx \geq (1 - G(\frac{(1-\theta)D(1-p)}{p}))\gamma q^*. \tag{17}$$

The left-hand side is the allocative gain: all the cases in which the lowest cost firm is sufficiently better than the second-best to incentivize the honest procurement officer to pick this firm only. The right-hand side is the cost of auditing with probability q^* procurement officers picking a single firm, which happens with probability $1 - G(\frac{(1-\theta)D(1-p)}{p})$ where G is the cdf of g. Hence, if auditing costs are not too low, a policy of random audit with endogenous procedure choice can outperform the discrete lottery.

B.2 Collusion and bid rigging

Consider the simple case in which N firms participate in the procurement process and choose whether or not to collude. Assume that no direct transfer is possible between firms. In a cartel procedure, in the second-price auction, the total surplus of cartel members is higher when all firms agree that the one with the lowest cost $c_{(1,N)}$ bids the smallest possible rebate $b = \epsilon > 0$ while all the others bid $b = 0$. For simplicity, assume that there is no bad state of the world $\theta = 0$. With common discount factor δ, the expected profit from future collusion in infinite time is thus

$$S_{collusion} = \delta \frac{V - E(c_{(1,N)})}{(1-\delta)N}, \tag{18}$$

at each period a firm has the lowest cost with probability $1/N$ and receives the difference between the reserve price and the cost.

The firm with the highest incentive to deviate is the one with the second highest cost $c_{(2,N)}$. In the second-price auction, it could gain $V - c_{(2,N)}$ by doing so. Finally, upon observing a deviation, firms could decide to end collusion and start competing, with expected profit

$$S_{competition} = \delta \frac{E(c_{(2,N)} - c_{(1,N)})}{(1-\delta)N}, \tag{19}$$

corresponding to winning with probability $1/N$ at each period the difference between the highest and second-highest cost. Hence, collusion can be sustained with the threat of punishment as long as $V - c_{(2,N)} < S_{collusion} - S_{competition}$.

In the third price lottery, the expected profit from future collusion is

$$S'_{collusion} = \delta \frac{V - \frac{E(c_{(1,N)} + c_{(2,N)})}{2}}{(1-\delta)N} < S_{collusion}, \tag{20}$$

as the allocation is less efficient. Neither the firm with the lowest or the second lowest cost have an incentive to deviate, as it would not affect their probability of winning. The firm with the third highest cost could gain $\frac{V - c_{(3,N)}}{2}$ by doing so. Finally, the expected discounted future profit with competition is

$$S'_{competition} = \delta \frac{E(c_{(3,n)} - \frac{(c_{(1,N)} + c_{(2,N)})}{2})}{(1-\delta)N}, \tag{21}$$

with high higher than $S_{competition}$ as long as the expected profit of the winning firm is higher in the third price lottery than in the second price auction. This holds for instance with uniform distributions (as the gap between order statistics is constant) or (using simulations) for standard log-concave density functions such as the Normal distribution. Collusion can be sustained as long as $\frac{V - c_{(3,N)}}{2} < S'_{collusion} - S'_{competition}$.

When comparing the second price auction and the third price lottery, we see that the immediate gains from a deviation are smaller in the latter. This is because a least efficient firm would deviate, and would not get the contract with certainty. There are two effects on the right-hand side of the conditions for collusion to be sustained, and both go in the direction of making collusion easier to sustain with second price auctions. In the third price lottery, punishment is smaller, because the firm's expected surplus is higher in a competitive environment, and smaller in a collusive one. This makes collusion more difficult to sustain, similar to the intuition in Chassang and Ortner (2019).

C Simulations

C.1 Contracted discount, default rate, consumer and producer surplus

We start with six bidders, $l = 0, h = V = T = 1, \theta, D = 2/5$ and τ sufficiently high for firms of type H to never default.

Figure 2 displays the average winning bid for each of the possible procedures. As expected, the procedure extracting the highest contracted discount, slightly above 0.6, is the second-price auction. By definition, the lottery does not lead to any discount, and the third-price lottery fares somewhere in between the two, slightly under 0.5. A key point here is the number of bidders and the dispersion of the cost function. Adding more bidders would make the two above lines closer to each other, while removing some would put them further away.

Figure 3 then shows the main advantage of the two procedures with a lottery component when compared to the second price auction is the fact that more firms actually deliver the project at the contracted price. In the standard lottery, the share of winners with no reputation concern and without cost of defaulting (type L) is equal to their share in the population, 50%. Of them, only those with the highest cost fail to deliver in the bad state, so that 80% of the winning firms always deliver. In the third price lottery, two-third of the selected firms have the undesirable type L, but around 50% of the firms always deliver nonetheless. In the second price auction, three-quarters of the selected firms are of type L, and only around 30% of the selected firms always deliver.

The first difference between a lottery procedure and the auction is thus that the latter is more likely to select a bidder without cost of defaulting (type L). Indeed, we have seen that this kind of firm bids more aggressively, and are therefore more likely to win a competitive tender. The second is that, by guaranteeing such a high price to the firm, procedures involving a lottery ensure that even those with a high cost often do not default. The exception is when the firm selected by the lottery is of type L and has such a high cost that delivering at the reserve price is too costly for her when the state is bad.

Figure 4 shows the expected consumer surplus from each of the procedures. For the chosen parameter values, the third-price lottery is superior in that dimension. The reason is that the difference between the second and the third highest bid is not very high, while the difference in the probability of default is important. The choice of parameter values is however crucial. For instance, with a lower cost of default for society $T = 0.4$ (instead of $T = 1$), the second-price auction performs better. The same holds if we reduce the number of bidders (see below). Consumer surplus is always negative with the Lottery, as the project is awarded without any discount, but society bears the expected cost of default.

Figure 5 shows that producer surplus, is the highest for the standard lottery. One reason is mechanical: the winning firm does not offer any discount, and thus charges the highest possible price. Another comes from the benefit from actually delivering the project in every state of the world. In comparison, there is a big loss of surplus in the second-price auction coming from the fact that the winning firm is likely to default and receive no surplus at all. Lotteries are however not perfect in that dimension, as there is an allocative efficiency problem: The winning firm is expected to have a highest cost than in all other procedures.

For comparison we copy below (figures 6, 7, 8 and 9 the results of simulations with 4 bidders, uniform cost distribution $l = 0, h = V = T = 1, \theta, D = 2/5$ and τ sufficiently high for firms of type H to never default.

Figure 2: Contracted discount with 6 bidders

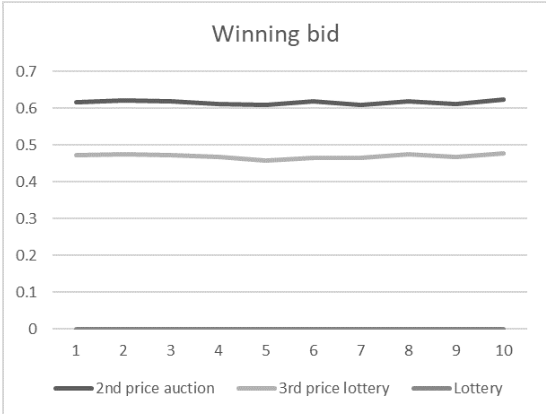

Figure 3: Default/renegotiation rate with 6 bidders

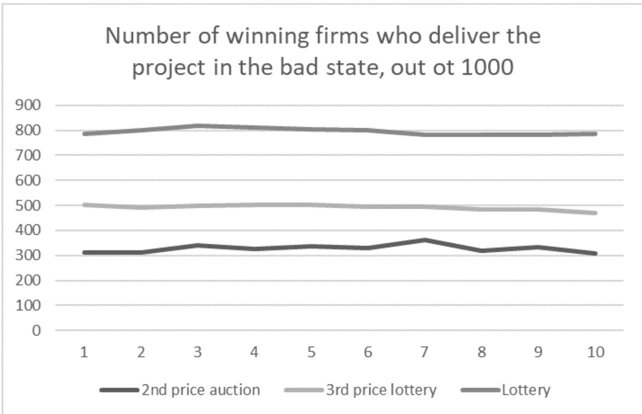

Figure 4: Welfare with consumer standard with 6 bidders

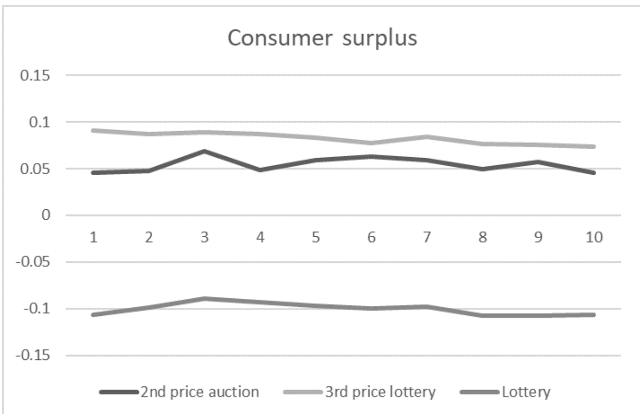

Figure 5: Profits with 6 bidders

Figure 6: Contracted discount with 4 bidders

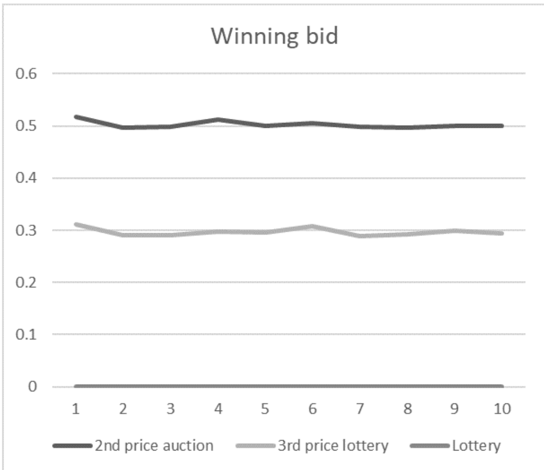

Figure 7: Default/renegotiation rate with 4 bidders

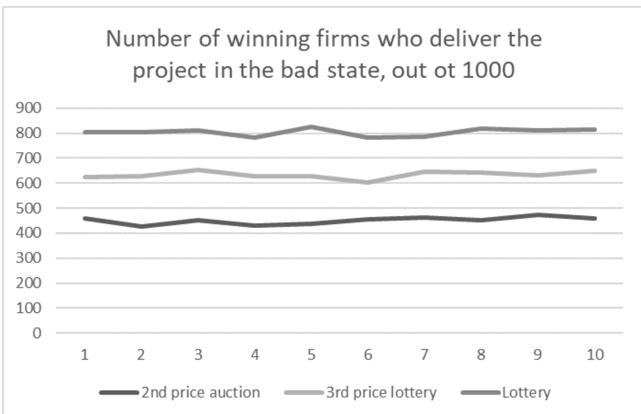

Figure 8: Welfare with consumer standard with 4 bidders

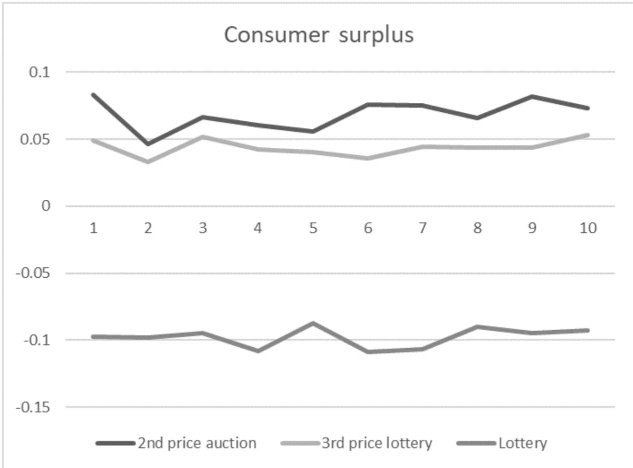

Figure 9: Profits with 4 bidders

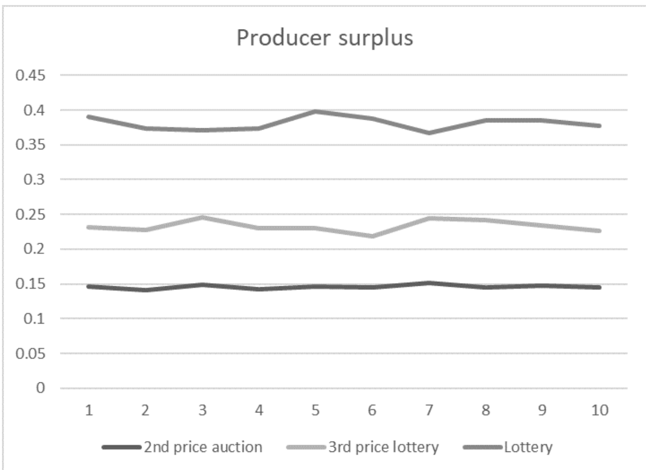

Table 3: Consumer surplus (auction ; 3rd price lottery ; standard lottery)

D/T	0.5	0.75	1
0.1	.53 ; .47 ; - .01	.50 ; .45 ; -.02	.47 ; .43 ; -.02
0.2	.40 ; .35 ; -.03	.34 ; .32 ; -.03	.26 ; .28 ; -.05
0.3	.30 ; .27 ; -.04	.21 ; .21 ; -.05	.15 ; .17 ; -.08
0.4	.22 ; .20 ; -.05	.13 ; .14 ; -.07	.06 ; .09 ; -.10

Table 4: Producer surplus (auction ; 3rd price lottery ; standard lottery)

D/T	0.5	0.75	1
0.1	.12 ; .19 ; .47	.10 ; .17 ; .45	.12 ; .18 ; .47
0.2	.11 ; .17 ; .46	.11 ; .17 ; .44	.11 ; .18 ; .45
0.3	.10 ; .15 ; .42	.11 ; .16 ; .41	.11 ; .16 ; .41
0.4	.10 ; .15 ; .38	.10 ; .16 ; .37	.10 ; .15 ; .39

C.2 Tables for a broader range of parameter values

We report in Table 3 the consumer surplus in the form of (second-price auction, third-price lottery, standard lottery) for different values of T and D, average of 10000 simulations. Table 4 reports the producer surplus for the same parameter values. Table 5 is the aggregate welfare with a weight $\alpha = 1/2$ on producer surplus, and Table 6 with a weight on producer surplus equal to the weight on consumer surplus, $\alpha = 1$.

Table 5: Welfare wtih $\alpha = 1/2$ (auction ; 3rd price lottery ; standard lottery)

D/T	0.5	0.75	1
0.1	.59 ; .57 ; .23	.55 ; .53 ; .28	.53 ; .52 ; .21
0.2	.45 ; .43 ; .23	.39 ; .40 ; .19	.32 ; .36 ; .17
0.3	.35 ; .35 ; .21	.27 ; .29 ; .16	.20 ; .25 ; .12
0.4	.27 ; .28 ; .19	.18 ; .22 ; .9	.11 ; .17 ; .10

Table 6: Welfare wtih $\alpha = 1$ (auction ; 3rd price lottery ; standard lottery)

D/T	0.5	0.75	1
0.1	.65 ; .66 ; 46	.60 ; .62 ; .43	.59 ; .61 ; .45
0.2	.51 ; .52 ; .43	.45 ; .49 ; .41	.37 ; .46 ; .40
0.3	.40 ; .42 ; .38	.33 ; .37 ; .36	.26 ; .33 ; .33
0.4	.32 ; .35 ; .33	.23 ; .30 ; .30	.16 ; .24 ; .29