THE UNIVERSITY *of York*

Discussion Papers in Economics

No. 12/26

Fuzzy Price-Quality Ratio Procurement under Incomplete Information

Thomas Giebe and Paul Schweinzer

Department of Economics and Related Studies
University of York
Heslington
York, YO10 5DD

Fuzzy Price-Quality Ratio Procurement under Incomplete Information*

Thomas Giebe† Paul Schweinzer‡

Abstract

We analyse a procurement auction in which sellers are distinguished on the basis of the ratios of quality per unit of money they offer. Sellers are privately informed on the quality of the technology or good they offer. We assume that the procurer cannot perfectly identify the best offer. Thus, with positive (and decreasing) probability, the second, third, etc. best ratio offered is selected as the winner of the auction. We model the decision process as based on a general noisy ranking of offers. We show that, although the problem seems to be analytically intractable in general, there exists a simple symmetric pure-strategy equilibrium, provided that the procurer's ranking technology employs the right degree of noisiness. (JEL *C7, D7, H57.* Keywords: *Auctions, Contests, Procurement.*)

1 Introduction

In many procurement settings, bidders are compared, and winners selected, on the basis of the offered quality per unit of money, or the quality-price ratio.[1] In these applications, 'quality' can be a multi-dimensional property of the product in question, typically summarized by some score. While the bidders might be fully aware of the quality they offer, the procurer is often not able to perfectly assess the offered qualities, for instance, when the procurement decision is based on designs or prototypes. Thus, with positive probability, sub-optimal offers are selected as the winner and the best offer does not get the award with certainty.

Moreover, in many scenarios, the procured objects may have aspects of credence or experience goods such that the actual quality is not fully revealed (or may not be verifiable by a third party), making it impossible to make payments conditional on actual quality. Examples seem to abound in the government procurement of, for instance, long-term defense capabilities which are only developed

*Thanks to Dan Kovenock, Ella Segev, and Cédric Wasser for helpful discussions. †Microeconomics, Technische Universität Berlin, Str. des 17. Juni 135, 10623 Berlin, Germany, thomas.giebe@tu-berlin.de. ‡Department of Economics, University of York, Heslington, York YO10 5DD, United Kingdom, paul.schweinzer@york.ac.uk.

[1] In its *Guide to Greener Purchasing*, the Organisation for Economic Co-operation and Development (2000, p.12) writes that the objective of procurement rules in member countries is "to achieve a transparent and verifiable best price/quality ratio for any given product or service." Quality-price ratios (or, synonymously throughout the paper, price-quality ratios) are thus used explicitly for assessing bids for procurement purposes by many governments. An example is the Scottish Government's Construction Procurement Manual (http://www.scotland.gov.uk/Publications/2005/11/28100404/AnnexA).

on the basis of an award. In such a setting, bidders, offering a certain quality, face the strategic task to bid such that their expected payoff is maximised, taking into account that the offers might be misjudged by the procurer. This problem of a buyer's imprecise evaluation of the sellers' offers is at the core of our paper.

We explore this problem by assuming that the procurer cannot perfectly rank the sellers' offers. Instead, the procurer's decision procedure is modeled as a general noisy or fuzzy ranking technology that determines a winner from among all offers in a sealed-bid procurement auction. Each seller's offer consists of an object of given quality and a financial bid. The latter is the payment the seller demands in case she wins the auction and has to deliver the object in return. At the time of bidding, the object's quality is assumed to be fixed.[2] Thus, the seller's strategic variable is the 'value for money' implied by his financial bid for the given object.

The procurer wants to select the best quality per unit of money among all offers. The ranking technology determines a winner on the basis of quality-price ratios, i.e., the ratio of the offered quality and the financial bid. The offer with the best actual quality-price ratio is most likely to win, but does not win with certainty. The probability of winning depends on the particular ranking technology used by the buyer.

Our main motivation in this paper is to find a tractable way of solving a fuzzy auction assignment problem in the context of a procurement setting. In order to analyse this problem, we need to assign well-defined and smooth winning probabilities on the basis of the submitted bids. For precisely this purpose, the complete information contests literature developed the well-founded and axiomatised generalised Tullock success function (a generalisation of) which we adopt for our model. From a theoretical point of view, the contribution of the present paper is to integrate this fuzzy assignment technology into an incomplete information auction framework.

Although the problem turns out to be difficult to handle in general, we can identify a simple symmetric pure-strategy equilibrium of our stylised procurement game, in which all sellers offer the same quality-price ratio, i.e., sellers with larger quality demand larger payments. This equilibrium requires a particular target precision of the ranking technology employed by the buyer; it does not exist if the ranking is too precise or too imprecise. This notwithstanding and as the basis for a behavioural rule, we are able to confirm that, for small deviations from this target precision, a seller benefits from bidding close to the prescribed equilibrium bids.

We demonstrate that a class of winning probability assignments, which contains the generalised Tullock success function, is an example of a feasible technology in our incomplete information setting.

[2] As an example of an actual procurement which seems to feature elements of what we model, consider the US Joint Strike Fighter (JSF) acquisition which resulted in the development of the Lockheed Martin F-35 airplane. After fierce competition between Boing and Lockheed, the JSF system development and demonstration contract was awarded on 26-Oct-01 to Lockheed Martin on a 'winner takes all' basis. In relation to our key model assumptions, evidently both Boing and Lockheed Martin were capable of delivering a fighter aircraft technology prior to selection; the companies' costs of acquiring this technology was sunk at the competition stage. Since the advanced features required in the JSF specification were not yet developed at the award stage, a perfect discrimination between projected qualities seems to have been impossible. Finally, even if serious quality problems should surface ex-post, recouping a non-negligible share of the estimated US$323 billion development and procurement cost of the F-35 is hardly conceivable as Lockheed Martin is also the exclusive supplier of other key US weapons programmes such as, for instance, the F-22 Raptor aircraft. (All sources: `wikipedia`, 7-Sep-11.)

For the Tullock case, the equilibrium requires a certain precision parameter above one and below two, implying on the one hand that the well-known 'lottery contest' cannot implement the equilibrium, and on the other hand that we should not expect severe equilibrium existence problems. If the buyer can design the procurement auction such that the target precision parameter is implemented, then this equilibrium always exists.

Literature

This paper combines ideas from fuzzy, or 'imperfectly discriminating' contests under complete information with a procurement problem that makes use of a price/quality scoring rule under incomplete information. In other words, we introduce the probabilistic assignment of winners into an auction setting. Therefore, we combine different strands of literature as follows.

The assumption of private information is standard and has been extensively analyzed in auction theory, and the theory of perfectly discriminating contests (all-pay auctions). In the realm of imperfectly discriminating contests, however, the literature has largely avoided incomplete information.[3]

Due to technical difficulties, general solutions remain elusive and, at present, very little is known about the case of incomplete information.[4] The literature has found closed-form solutions only for special cases, mostly standard two-player Tullock contests where one or both players are privately informed about their (discrete) valuation of the prize or their (constant) marginal cost. Examples, among others, are Hurley and Shogren (1998b), Malueg and Yates (2004), and Schoonbeek and Winkel (2006). In Katsenos (2010), players can signal their marginal cost prior to the contest. Münster (2009b) looks at repeated contests. Wärneryd (2003) and Wärneryd (2009) assume a common value of winning.

Existence proofs as well as comparative statics (e.g., with respect to rent dissipation and aggregate effort) have been provided by Fey (2008), Prada-Sarmiento (2010), Ryvkin (2010), Wasser (2010), Wasser (2012). Ko (2012) works backwards from the equilibrium distribution of efforts in order to shed light on the solution. Numerical results were obtained by, e.g., Hurley and Shogren (1998a). Ewerhart (2010) provides the only available analytical solution for a given (specially designed) continuous distribution.

All of the above cited works are related to our model only in the sense that they are concerned with incomplete information contests. The obvious difference is that our private information only affects a player's price/quality score and, thus, the probability of winning, rather than the cost of effort or the valuation of winning. For our setting, we provide a closed-form symmetric equilibrium for any number of players and a contest success function that is more general than the Tullock form. Due to the simplicity of the equilibrium strategy, we can allow for general joint distributions of private information.

The widely used Tullock contest specification has been justified axiomatically (e.g. Skaperdas (1996), Münster (2009a), Arbatskaya and Mialon (2010)), with micro foundations (e.g. Corchón and Dahm (2010), Fu and Lu (2012), Jia (2008)) and, recently, using a mechanism design ap-

[3] See, for example, the surveys Corchón (2007), Garfinkel and Skaperdas (2006), and Konrad (2008).
[4] See the extensive discussion of the state of the literature in Ryvkin (2010) and Wasser (2010).

proach Polishchuk and Tonis (2011)). Our paper is technically related to Arbatskaya and Mialon (2010). They study multi-armed contests, in which players choose several efforts simultaneously. Our price/quality ratios can be seen as two kinds of effort that a player brings into the contest. However, the clear difference is that quality in our case cannot be changed at the contest stage and, thus, is a type, rather than a strategic choice. Moreover, quality is part of our contest designer's preferences.

Our paper is clearly related to the literature on (standard) scoring auctions. Che (1993), Che and Gale (2003), and Asker and Cantillon (2008) study quasilinear scoring rules, the latter for multidimensional types. We are, however, concerned with price/quality ratios, i.e., a nonlinear scoring rule. Standard first- and second-score auctions for this case are analysed by Hanazono, Nakabayashi, and Tsuruoka (2012). In contrast to our paper, all of the above assume that the highest score wins with certainty, and, with the exception of Che and Gale (2003), that production takes place after the auction, i.e., the solution is driven by player's privately known cost functions which play no role in our case, where production cost is sunk at the bidding stage.

Our basic assumption that the buyer cannot perfectly determine the quality of an offered good or service seems to be natural and has been made previously, for instance, by Dranove and Satterthwaite (1992). In a theoretical and empirical study of procurement, Decarolis (2010) gives another justification to our approach that the highest or best bidder does not necessarily win. Under widely used rules, high bids are eliminated if they differ too much from some weighted average of bids in order to reduce the risk of defaults and renegotiation. This is equivalent to saying that the probability of winning for the best bidder is less than one.

2 Model

Consider a two-dimensional procurement setting where $n \geq 2$ sellers offer objects of privately known and fixed quality. Qualities (types) $\theta = (\theta_1, \ldots, \theta_n)$ are distributed according to the joint cumulative distribution function $F_{(0,\infty)^n}$ with positive joint density f. Denote the type space by Θ. As usual, θ_{-i} and Θ_{-i} indicate vectors with the i^{th} component removed. The joint density of i's rivals' qualities is $f_{-i}(\theta_{-i})$. Sellers simultaneously quote (verifiable) ask prices $b_i > 0$, which we call bids, in a sealed-bid auction.[5]

At the time of bidding, the number of participating bidders is commonly known. Moreover, we assume that the objects to be procured are already in existence and worthless for the sellers.

[5] We exclude zero bids for lack of economic sense. See Decarolis (2010) for the common procurement practice to exclude extreme bids that are 'too good to be true'. Notice that we study a first-price setting here. In subsection 5.2 we briefly look at a second-score auction. There, a zero bid is feasible, unless the buyer chooses to specify a minimum acceptable bid, because the winner's payment is independent of his bid.

Hence, production costs are sunk at that stage.[6] We do not explicitly model the production process of the objects on offer, but solve our game assuming that qualities are distributed according to the cumulative distribution function F. As will become clear, the equilibrium we derive is entirely independent of the assumptions on F. Thus, there may be an unmodeled production or innovation stage before the procurement stage, in which sellers decide strategically how much to invest into producing, resp. innovating, an object in some noisy production or innovation process. Making such choices before deciding whether to participate in the bidding game and how much to bid there, however, does not change any of our results. The cumulative distribution function of qualities, F, captures the uncertainty of sellers and the procurer as to what qualities the (other) sellers' objects might have at the procurement stage.

As a particular 'scoring' rule, we assume that the buyer is only interested in each player i's ratio of quality over the bid $c_i = \theta_i/b_i$. The procurer prefers higher ratios to lower ratios, but cannot perfectly assess an offered object's quality. Therefore, his decision making is based on an imperfect ranking technology, the details of which are specified below. The financial component of a seller's bid is perfectly observable to the buyer; the quality component stems from the individual seller's specifications in the tender documents. Combined, the buyer can infer a fuzzy ranking of quality-price ratios on the basis of which he makes the award.

In practical applications, the buyer might be able to verify that the offered items exceed a certain threshold quality, and, moreover, the buyer might have a 'target' quality-price ratio in mind, based on an existing, currently used technology that is to be replaced by the offered technology.[7]

Finally, we would like to point out that selecting the best quality-price ratio is consistent with standard profit maximisation, if, e.g., we assume the procurer's profit to be $\theta - b$. Then a larger quality θ, resp., a lower ask price b, imply both a larger ratio θ/b as well as a larger profit.

In the following, we consider a procurement game, where n sellers, owning one object of fixed quality each, simultaneously submit a sealed financial bid. The procurer employs a noisy (imprecise) ranking technology that determines a winner as a function of sellers' actual quality-price ratios. The winning bidder receives his financial bid in return for his object. Before we introduce and analyse the general ranking technology in section 4, we present an example for such a technology in order to illustrate the main result. We complement our analysis of fuzzy rankings with that of an infinitely precise ranking in section 5 and conclude with a short discussion in section 6. All proofs are in the appendix.

[6] This is only a conceptual requirement—it is perfectly possible to think about equivalently applicable situations such as the procurement of some production technology where no physical object exists but the capability to produce the technology is fixed during the procurement process. Che and Gale (2003) explicitly model production (innovation) which, similar to our setting, happens before the procurement stage. The authors mention that all that matters is that quality and financial bids are chosen sequentially, whereby bids are chosen when rivals bidders' qualities are still unknown. In this sense, Che and Gale (2003) provide a micro-foundation for the distribution of qualities assumed in our model.

[7] In the context of our airplane procurement example of footnote 2, one part of the minimum requirements at the system development and demonstration stage was that the involved prototypes could actually fly, implying a minimum verifiable quality threshold. The reserve quality-prices ratio in that example stems from those of the General Dynamics F-16 and Fairchild Republic A-10 aircrafts the new F-35 is, among others, intended to replace.

3 Example

Suppose that seller i, who owns an object of quality θ_i, faces the problem of choosing a bid b_i in order to participate in a sealed-bid procurement auction. The winning bidder receives his own bid. We further assume that the procurer employs a ranking technology that determines seller i's probability of winning the auction, π_i, as a function of each seller's actual quality-price ratio as

$$\pi_i(\theta_i, b_i) = \frac{(\theta_i/b_i)^r}{(\theta_i/b_i)^r + \sum_{j \neq i} (\theta_j/b_j)^r}, \quad r > 0 \tag{1}$$

In the complete information contest literature, this functional form is known as the generalised Tullock success function.[8] The parameter r is sometimes called the ranking's precision parameter. A higher value of r implies a higher probability for the best ratio to win but no finite r results in the highest ratio winning with certainty. Thus, a privately informed seller i maximises expected payoff

$$u_i(\theta_i, b_i) = b_i \int_{\Theta_{-i}} \pi_i(\theta_i, b_i) f_{-i}(\theta_{-i}) d\theta_{-i}. \tag{2}$$

This general problem seems to be analytically intractable and has, to the best of the authors' knowledge, not been solved. However, we conjecture that in symmetric, pure-strategy equilibrium, a player of type θ_j uses her financial bid in order to achieve a *constant* ratio $\theta_j/\beta(\theta_j) = c$, $c > 0$, implying a candidate equilibrium bid of $\beta(\theta_j) = \theta_j/c$. Then, the distribution of qualities, F, does not matter because player i only cares about his rivals' quality-price ratios, not about the components of those ratios. In a two-player setting, i's best-response problem (2) simplifies point-wise to

$$u_i(\theta_i, b_i) = b_i \frac{(\theta_i/b_i)^r}{(\theta_i/b_i)^r + c^r}. \tag{3}$$

Computing the derivative with respect to b_i, and evaluating at the candidate $\theta_i/b_i = c$, player i's first-order condition simplifies to $r = 2$, implying that this simple equilibrium with constant quality-price ratios $c = \theta_j/b_j$ for all $j = 1, \ldots, n$ exists only if the ranking technology has the precision $r = 2$. Under this provision, any commonly chosen price-quality ratio $c > 0$ constitutes an equilibrium of our procurement game.

Having characterized the set of symmetric equilibria we briefly discuss our result. First, participation in procurement contests is usually restricted to sellers who meet certain criteria, e.g., there might be a minimum verifiable threshold quality (as mentioned in the previous section). Then participation is restricted to the players that can provide sufficient quality. Thus, in the following, the number of sellers, n, is understood to be the number of players who satisfy the buyer's participation criteria. This number of 'short listed' sellers is assumed to be commonly known, as is usually the

[8] There are two popular interpretations of the Tullock success probabilities: one is the lottery ticket idea where a player's winning probability is a function of the number of tickets bought over the total number of tickets (where the equilibrium slope of this function is controlled by the exponent r). The second interpretation—which may be equally interesting in the present framework—is to see the winning probabilities as contract shares, i.e., the buyer might simultaneously award a contract to several suppliers. For details and an application to the production of a divisible good (of differing qualities) across several bidders, see Gong, Li, and McAfee (2011).

case in government procurement.[9] As we show in the general analysis, the equilibrium requires a ranking precision that is a function of the commonly known number of participating sellers. Second, the buyer might announce a reserve or target quality-price ratio, possibly based on the ratio provided by an existing technology that is going to be replaced. This announcement might act as a coordination device, selecting an equilibrium.[10] Finally, and in the tradition of a large literature on contest success functions, we might want to view the ranking technology itself as a black box, coming up with the required probabilities on the basis of information which is neither contractible nor, in principle, verifiable.[11]

In a numerical example with ranking precision $r = 2$ and equilibrium price-quality ratio $c = 2$, we show deviation utilities for a quality type of $\theta_i = 3$ and competitor quality $\theta_j = 10$ in the below figure.

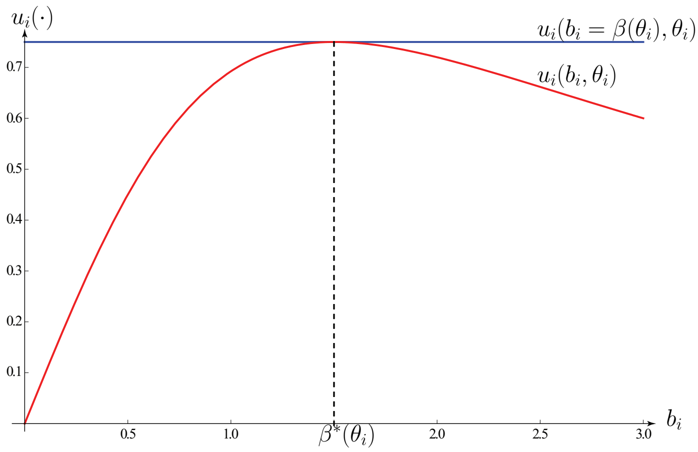

Figure 1: The top, horizontal line is player i's equilibrium utility from using the equilibrium bid $\beta(\theta_i)$. The curve below shows i's unilaterally deviation utility from using b_i. (The objective has no other extrema.)

We would like to reiterate that in this example we can only identify an equilibrium for the particular design parameter of $r = 2$. The general characterisation of equilibria for all values of r—perhaps interpretable as the level of detail applied to the scrutiny of the tender documents—is still elusive. The parameter of r, however, is a *design* parameter of the used ranking and can be chosen by the designer. The progress we can report, therefore, seems to constitute a first step with respect the auctioning of goods with fuzzy or stochastic assignment rules under incomplete information. Equipped with the two-players result, we now present our general analysis and results.

[9] Regardless of whether there is such a verifiable quality level, short listing is a common procurement procedure. It reduces cost duplication at the bid preparation stage and makes participation more profitable. Moreover, it simplifies the sellers' decision problem.

[10] Decarolis (2010) provides an analysis of the common procurement practice to reject seller bids that are 'too good to be true' in the sense that their price-quality relation is unrealistic or too far away from the average offer. This practice might also help sellers to converge on the average offered ratio in the way required for our equilibria.

[11] For a recent discussion see, for instance, Corchón and Dahm (2010) and the references therein.

4 Analysis

Assume that the procurer can make use of a noisy and partial but verifiable ranking of quality-price ratios[12]

$$\Gamma(\tilde{\mathbf{c}}) = [\pi_1(\tilde{c}_1), \ldots, \pi_n(\tilde{c}_n)] \tag{4}$$

where $c_{ij} = c_i/c_j$ is the ratio of sellers i and j's quality-price ratios, and $\tilde{c}_i = (c_{i1}, c_{i2}, \ldots, c_{in})$, with the i^{th} element equal to 1. Thus (4) ranks the players on the basis of ratios of quality-price ratio pairs such that $\pi_i(\tilde{c}_i)$ is player i's probability of being ranked first given the vector of quality-price ratios $\mathbf{c} = (c_1, \ldots, c_n)$. Together with a smoothness assumption, we assume the following on $\pi_i(\cdot)$:[13]

A1 Symmetry: every opponent of player i affects the winning probability of i in a similar way. Thus, if players l and m exchange their quality-price ratios, this does not affect the winning probability of player $i \notin \{l, m\}$. This assumption implies that, in symmetric equilibrium where $c_1 = \cdots = c_n$, the slope of π_i with respect to any ratio c_{ij} is the same for all i, j, $j \neq i$ and each ratio is equal to 1. We denote this slope by $\pi'(\mathbf{1})$.

A2 Responsiveness: i's probability of being ranked first should react positively to both an improvement in i's quality-price ratio and a deterioration of a rival's quality-price ratio.

The following result characterizes a class of symmetric equilibria for this game.

Proposition 1. *Consider a procurement game with $n \geq 2$ sellers, where the winning probability is given by $\pi_i(\tilde{c}_i)$ and has the property $\pi'(\mathbf{1}) = \frac{1}{n(n-1)}$. This game has a continuum of symmetric pure-strategy equilibria where all players bid the same constant quality-price ratio, c.*

Our general class of quality rankings (resp. winning probabilities) includes the generalised Tullock success function (1) that we have used in the example section in order to illustrate our result. To see this, consider

$$\pi_i(\tilde{c}_i) = \left(\sum_{j=1}^{n} c_{ij}^{-r}\right)^{-1} = \frac{c_i^r}{\sum_{j=1}^{n} c_j^r}, \quad r > 0. \tag{5}$$

For this particular ranking technology, and $j \neq i$,

$$\frac{\partial \pi_i(\tilde{c}_i)}{\partial c_{ij}} = -\left(\sum_{j=1}^{n} c_{ij}^{-1}\right)^{-2} (-r) c_{ij}^{-r-1}. \tag{6}$$

Evaluated in symmetric equilibrium, where $c_{ij} = c/c = 1$, this simplifies to r/n^2. From this we can determine the optimal 'monitoring precision' r as follows. By proposition 1, the (necessary) equilibrium condition is $\frac{\partial \pi_i(\tilde{c}_i)}{\partial c_{ij}} = \pi'(\mathbf{1}) = \frac{1}{n(n-1)}$. Thus,

$$\frac{r}{n^2} \overset{!}{=} \frac{1}{n(n-1)} \iff r^* = \frac{n}{n-1}, \tag{7}$$

[12] The verifiability in the procurement context is often a legal requirement satisfied by means of scores defined in the tender procedure. The technical specification which follows is adapted from Gershkov, Li, and Schweinzer (2009), changed to reflect the present quality-price ratio based setup.

[13] We only state an intuitive version of the required assumptions in the main body of the paper. The precise requirements are relegated to the appendix.

corresponding to our result in the example section. There, we explicitly compute the result, rather than deriving it from our main result. We summarize our findings in the following corollary.

Corollary 1. *Consider a procurement game with $n \geq 2$ sellers, where the winning probability is given by the generalized Tullock assignment probability*

$$\pi_i(c_i, c_{-i}) = \frac{c_i^r}{\sum_{j=1}^n c_j^r}, \quad r > 0. \tag{8}$$

If the ranking precision is equal to $r = n/(n-1)$, then the game has a continuum of symmetric pure-strategy equilibria where all players bid the same constant quality-price ratio, c.

In proposition 1 and its corollary, we only look at the necessary first-order condition(s) for symmetric equilibria. The following result demonstrates existence for the class of Tullock technologies used in corollary 1.

Proposition 2. *Consider the equilibrium characterised in corollary 1. This equilibrium exists for all feasible values of model parameters.*

Let us now briefly examine the robustness of the proposed equilibrium by checking a slight deviation from the equilibrium ranking precision r^* by the buyer. Assume that all but one sellers stick to bidding the same constant ratio c. Consider a Tullock technology with precision parameter $r_\varepsilon = (1+\varepsilon)r^*$ where (as above) $r^* = n/(n-1)$ and the deviation $\varepsilon \neq 0$ is small. Then the first-order condition of player i's best response problem, given in (19), is satisfied if[14]

$$\left(\frac{\theta_i}{b_i}\right)^{r_\varepsilon} = (n-1)(r_\varepsilon - 1)c^{r_\varepsilon}. \tag{9}$$

Now, denote some player i's strategy in terms of deviations from the other players' constant ratio: $\theta_i/b_i = \delta c$ for $\delta > 0$. Hence, $\delta = 1$ implies no deviation and following the ratios of the other sellers. Thus, c^{r_ε} cancels out in the above and (9) simplifies to

$$\delta^{\frac{(1+\varepsilon)n}{n-1}} = \varepsilon n + 1 \quad \iff \quad \delta^*(\varepsilon) = (\varepsilon n + 1)^{\frac{n-1}{(1+\varepsilon)n}} \quad \Rightarrow \quad \lim_{\varepsilon \to 0} \delta^*(\varepsilon) = 1. \tag{10}$$

Thus, the smaller the buyer's 'error' ε, the closer is a deviating seller's best response to those of the other 'constant ratio' sellers'. This result provides a behavioral continuity argument implying that, as long as the buyer's ranking precision is reasonably close to the 'optimal' precision, it is desirable for a seller to bid 'close to' the constant-ratio equilibrium under the optimal precision. Hence, although we cannot determine sellers' equilibrium behaviour under precision parameters different from (7), we know that utility gains or losses in a small neighbourhood about the equilibrium bids are bounded for small deviations from the target precision.

[14] Note that only the big parenthesis in the numerator is of importance, given that all players' ratios are positive.

5 Precise ranking technologies

In this section, we assume that the buyer can observe the sellers' qualities perfectly. Thus, the procurer still decides on the basis of quality-price ratios, but he can now verify qualities, implying that the sellers' ratio offers can be ranked precisely. Notice that, among the sellers, rivals' qualities and financial bids are not observable, as before.

The present section serves two purposes. First, it contrasts and supplements our main results for imprecise, fuzzy ranking technologies by assuming an infinitely precise ranking, as in standard auction theory. Second, it highlights the differences between our sunk-cost setup and the standard setup in which scoring auctions are analysed in the literature. There, procurement offers are essentially promises, where production takes place after procurement. Consequently, the results are driven by the sellers' private information on their (future) production cost. In our setting, by contrast, the tradeoff between a seller's expected revenue and production cost is missing, implying that sellers just maximize their expected revenue. The latter is typically not concave, which might explain the nature of the equilibria described below.

As mentioned in our discussion of the existing literature, standard first- and second-score auctions with nonlinear scoring rules and private ex-post production cost are discussed in detail by Hanazono, Nakabayashi, and Tsuruoka (2012). Following the standard auctions literature, we now refer to a quality-price ratio as a 'score', and to the respective auctions as 'scoring auctions'. For simplicity, for this section, we consider only the unit interval $[0, 1]$ as support for qualities.

5.1 First-score auction

In a first-score auction, bidders submit a financial bid b_i for their quality θ_i from which the buyer computes the ratio, or score, θ_i/b_i.[15] The highest score wins and the winner is paid his financial bid b_i in return for his object.

Proposition 3. *The first-score auction has a class of symmetric pure-strategy equilibria, characterized by the financial bid $\beta(\theta) = \theta/(kF^{n-1}(\theta))$ where $k > 0$ is a constant. This implies a monotonically increasing quality-price ratio equal to $kF^{n-1}(\theta)$.*

The proof is constructive and shows that our solution class includes the full family of bidding functions which feature a monotonically increasing ratio as a function of quality and are consistent with the first-order approach. In these equilibria, bidders with higher quality have larger expected payoffs although they bid larger ratios, i.e., weaker bidders demand *relatively* more money for their qualities. Moreover, in these equilibria, a seller's payoff, as a function of the financial bid, is constant and, thus, sellers are indifferent between financial bids.

Another class of equilibria arises when we introduce a minimum bid.[16] If the c.d.f. of qualities is such that its reverse hazard rate is larger than a certain multiple of the uniform distribution's

[15] Recall, that the quality of seller i's object is now observable between seller i and the buyer, but unknown to i's rivals.

[16] A similar result obtains if we assume that, in case of zero financial bids (infinite ratios), the largest quality wins.

reverse hazard rate $(1/\theta)$, then a seller's expected payoff is decreasing in the financial bid, given that all other sellers also bid the minimum bid. An example for such distributions is $F(\theta) = \theta^s$ for $s > 1/(n-1)$.

Proposition 4. *Consider the first-score auction with a minimum bid. If the cumulative distribution function of qualities satisfies*

$$\frac{f(\theta)}{F(\theta)} > \frac{1}{n-1}\frac{1}{\theta}, \tag{11}$$

then the first-score auction has a symmetric pure-strategy equilibrium where the financial bids are equal to the minimum bid.

5.2 Second-score auction

In this subsection, we again assume that qualities are perfectly observable to the buyer but are private information among sellers at the procurement stage. We now look at 'second-score' auctions defined as follows: Suppose that, as before, every player submits an object of quality θ_i and a financial bid b_i from which the score $c_i = \theta_i/b_i$ is computed. The highest scoring bidder wins and receives a financial payment such that the winner's ratio is reduced to that of the second best bidder, i.e., if bidder i wins and bidder j is the second best bidder, then i receives a payment of p such that $\theta_i/p = \theta_j/b_j$. In order to rule out division by zero in a ratio, we assume a minimum financial bid.[17,18]

Proposition 5. *In the second-score auction with minimum bid, bidders have the (weakly) dominant strategy to bid the minimum bid.*

The above derived (precise ranking) equilibria for the first- and second-score auctions are efficient in the sense that, in each case, the best quality object as well as the best available quality-price ratio is procured. This efficiency result does not seem surprising, given that a seller (with an item of given quality) can match any ratio-offer made by a seller with lower quality. Thus, in a symmetric equilibrium, it seems intuitive that the sellers with the largest qualities should win. For both pricing rules, we identify equilibria where sellers compete with the maximum permissable quality-price ratios (implied by the minimum bid). For the first-score auction, we identify a class of equilibria where each seller is made indifferent between financial bids.

6 Concluding remarks

This paper is the first to analyse the 'fuzzy' or probabilistic assignment of winning probabilities in a procurement auction. Our study is motivated by procurement settings where the buyer evaluates

[17] Again, an alternative assumption is that, in case of zero bids, the largest quality wins.

[18] We mention that a second-score auction requires verifiability of the first- and second-best qualities because the winner's payment depends on both. This is not necessary in the case of the first-score auction, where the promised payment (the winner's financial bid) is verifiable, anyway. Our purpose here is to contrast theoretical results obtained in the different auction formats, rather than recommend a mechanism.

offers by their price-quality ratios, while sellers' offers are composed of a financial bid and a design or prototype rather than a finished product. We take it as given, that, in such situations, it might be impossible to identify the best offer with certainty. Thus, we do not recommend the use of probabilistic rankings, we take them to be descriptive of reality. Nevertheless, we assume that the degree of the ranking's precision can be influenced by the buyer who may specify the contents of tender documents, capabilities of prototypes, etc.

In order to capture this uncertain assignment, we employ a general ranking technology that includes the widely used generalized Tullock success function as a special case. Although the fuzzy assignment problems turns out to be intractable in general, we show that a particular design case lends itself to a simple solution—the ranking precision can generally be chosen such that bidding a constant price-quality ratio is an equilibrium, involving a very simple strategy. Moreover, we argue that there is a 'myopic' robustness property to this equilibrium: if the ranking precision is only 'near' the target level, a behaviorally optimal prescription is to still bid (close to) the derived simple equilibrium strategy.

We summarise that an appealing and unique equilibrium can be reached in this complicated strategic setting if the buyer can guide the sellers by announcing a targeted price-quality ratio. Then each seller faces a simple strategic task, to bid such that this target ratio is met. For this equilibrium to exist, the seller must design the procurement process to implement a certain ranking precision of offers. Thus, another conclusion from this model may be that, given that an infinitely precise ranking is impossible, the buyer should not necessarily try to make the ranking as precise as possible—our simple equilibrium requires a ranking that is neither too fuzzy nor too precise. We show that, even if the ranking precision is only 'almost' correct in the context of our equilibrium, coordinating on a target ratio is a sensible behavioral prescription for the buyers, due to the strategic complexity of the problem.

Appendix

A. Assumptions on n-player ratio-based ranking technology

The complete statement of the required assumptions on the ranking technology $\pi_i(\cdot)$ introduced in section 4 is:

A1 Symmetry: For any two players $l \neq m$ and for any two vectors of quality-price ratios, $(c_1, \ldots, c_n)$ and $(c'_1, \ldots, c'_n)$ with $c_k = c'_k$ for $k \notin \{l, m\}$ and $c_l = c'_m$ and $c_m = c'_l$, we have $\pi_l(\tilde{c}_l) = \pi_m(\tilde{c}'_m)$. Moreover, for any player i, let the elements of a quality-price ratio vector $\tilde{c}'_i$ be arbitrary permutations of those in $\tilde{c}_i$ except for the element at the ith position. For these we require $\pi_i(\tilde{c}_i) = \pi_i(\tilde{c}'_i)$.

A2 Responsiveness: For any $l \in \{1, \ldots, n\}$ and $l \neq i$, $\dfrac{\partial \pi_i(\tilde{c}_i)}{\partial c_{il}} > 0$.

A3 $\pi(\cdot)$ is twice continuously differentiable.

Notice that the technical third assumption is not stated explicitly in the main text. An interpretation of the other assumptions is given in the body of the paper (section 4).

Proof of proposition 1. Consider seller i's utility maximisation problem, given that all other sellers $j \neq i$ choose their bids b_j such that they all offer the same constant quality-price ratio c, i.e., $b_j = \beta(\theta_j) = \theta_j/c$. Seller i needs to choose a bid b_i, given the quality θ_i, in order to maximise

$$u_i(\tilde{c}_i) = b_i \int_{\Theta_{-i}} \pi_i(\tilde{c}_i) f_{-i}(\theta_{-i}) d\theta_{-i} \tag{12}$$

where

$$\tilde{c}_i = (c_{i1}, c_{i2}, \ldots, c_{ii}, \ldots, c_{in}) = \left(\frac{c_i}{c}, \frac{c_i}{c}, \ldots, 1, \ldots, \frac{c_i}{c}\right) \tag{13}$$

and $c_i = \theta_i/b_i$. Thus,

$$u_i(\tilde{c}_i) = b_i \pi_i(\tilde{c}_i) = b_i \pi_i \left(\frac{\theta_i}{b_i c}, \frac{\theta_i}{b_i c}, \ldots, 1, \ldots, \frac{\theta_i}{b_i c}\right). \tag{14}$$

Note that the slope of c_{ii} with respect to b_i is identically zero because $c_{ii} \equiv 1$. Thus, we only need to consider $\frac{\partial \pi_i(\tilde{c}_i)}{\partial b_i}$ for $j \neq i$ in the following. The first-order condition is given by

$$\frac{\partial u_i(\tilde{c}_i)}{\partial b_i} = 0 \iff \pi_i(\tilde{c}_i) + b_i \sum_{j \neq i} \left(\frac{\partial \pi_i(\tilde{c}_i)}{\partial c_{ij}} \frac{\partial c_{ij}}{\partial b_i}\right) = 0. \tag{15}$$

In our symmetric pure-strategy equilibrium candidate, all players choose the constant ratio c, i.e., $c_i = c$ and $\tilde{c}_i = (\frac{c}{c}, \frac{c}{c}, \ldots, 1, \ldots, \frac{c}{c}) = (1, 1, \ldots, 1) = \mathbf{1}$, implying $\pi_i(\mathbf{1}) = 1/n$. By **A1**, the slopes $\frac{\partial \pi_i(\tilde{c}_i)}{\partial c_{ij}}$ for $j \neq i$ are equal. Recall that the uniform ratio c implies $c_{ij} = 1$ in which case that slope was denoted by $\pi'(\mathbf{1})$. Finally,

$$\frac{\partial c_{ij}}{\partial b_i} = -\frac{\theta_i}{b_i^2 c}. \tag{16}$$

Evaluated at the equilibrium candidate (where $\theta_i/b_i = c$), this is $\frac{\partial c_{ij}}{\partial b_i} = -\frac{1}{b_i}$. Therefore, the first-order condition, evaluated at the equilibrium candidate, can be written as

$$\pi_i(\mathbf{1}) + b_i(n-1)\pi'(\mathbf{1}) \left(-\frac{1}{b_i}\right) = 0 \iff \pi'(\mathbf{1}) = \frac{1}{n(n-1)} \tag{17}$$

which establishes our claim. □

Proof of proposition 2. Inserting the equilibrium candidate $b_j = \theta_j/c$, (2) simplifies to

$$u_i(\theta_i, b_i) = b_i \int_0^\infty \frac{(\theta_i/b_i)^r}{(\theta_i/b_i)^r + (n-1)c^r} f(\theta_j) d\theta_j = b_i \frac{(\theta_i/b_i)^r}{(\theta_i/b_i)^r + (n-1)c^r}. \tag{18}$$

The first-order condition with respect to b_i is

$$\frac{\left(\frac{\theta_i}{b_i}\right)^r\left(\left(\frac{\theta_i}{b_i}\right)^r-(n-1)(r-1)c^r\right)}{\left(\left(\frac{\theta_i}{b_i}\right)^r+(n-1)c^r\right)^2}=0. \tag{19}$$

Inserting the symmetric candidate $b_i=\theta_i/c$, simplifying, and solving for r we get

$$r=\frac{n}{n-1}. \tag{20}$$

In order to demonstrate existence, we evaluate player i's incentives to deviate from the equilibrium candidate quality-price ratio c. In order to simplify the exposition, we denote i's deviations from quality-price ratio c by kc with $k>0$, i.e., we consider bids $b_i=\theta_i/(kc)>0$ which implies that i offers a quality-price ratio of kc, with $k=1$ in equilibrium. Thus, the payoff from 'deviation k' is computed by inserting $b_i=\theta_i/(kc)$ in (18),

$$u_i^{dev}(k)=\frac{\theta_i}{kc}\cdot\frac{(kc)^r}{(kc)^r+(n-1)c^r}=\frac{\theta_i}{c}\cdot\frac{k^{r-1}}{k^r+n-1},\quad r=\frac{n}{n-1}. \tag{21}$$

In the equilibrium candidate, $k=1$,

$$u_i^{dev}(1)=\frac{\theta_i}{cn}. \tag{22}$$

We claim that $k=1$ is the most profitable deviation, confirming the equilibrium candidate. Thus, we want to show that

$$u_i^{dev}(1)>u_i^{dev}(k),\quad \forall k\neq 1,\quad r=\frac{n}{n-1}. \tag{23}$$

This is equivalent to

$$\frac{1}{n}>\frac{k^{\frac{1}{n-1}}}{k^{\frac{n}{n-1}}+n-1}\iff k^{\frac{n}{n-1}}+n-1>nk^{\frac{1}{n-1}}\iff n-1>(n-k)k^{\frac{1}{n-1}}. \tag{24}$$

The latter inequality is obviously satisfied for $k\geq n$ (then the right-hand side is nonpositive). Thus, it remains to consider $k<n$. Note that the left-hand side of the inequality is constant, while the right-hand side depends on k. In the following, we maximize the right-hand side. The first-order condition for a maximum is

$$\frac{\partial}{\partial k}\left((n-k)k^{\frac{1}{n-1}}\right)=0\iff\ldots\iff k^{\frac{1}{n-1}}(k-1)=0. \tag{25}$$

Since $k>0$, the only solution is $k=1$. Note that the second derivative is negative:

$$\frac{\partial^2}{(\partial k)^2}\left((n-k)k^{\frac{1}{n-1}}\right)=\frac{n}{(n-1)^2}k^{\frac{3-2n}{n-1}}(2-n-k)<0. \tag{26}$$

Thus, the right-hand side of our condition is uniquely maximised at $k=1$. The value of the right-hand side at $k=1$ is $n-1$ which implies that the condition holds for all $k\neq 1$. Finally, since (22) is positive, $b_i=0$ can be ruled out as a best response. This completes the proof. □

Proof of proposition 3. By assumption, qualities are fixed and observable by the buyer. Therefore, a seller's choice variable is the financial bid, $\beta(\theta)$, and each bid implies a unique quality-price ratio. W.l.o.g., we express strategies in terms of quality-price ratios where a ratio is denoted by $\tilde{\beta}(\theta) = \frac{\theta}{\beta(\theta)}$. We conjecture that there is a symmetric pure-strategy equilibrium, where every seller bids a monotonically increasing (and, thus, invertible) quality-price ratio, i.e., $\tilde{\beta}'(\theta) > 0$. Denote the c.d.f. of the largest quality among $n-1$ sellers by $G(\theta) = F^{n-1}(\theta)$ and its density by $g(\theta)$. Suppose seller i submits the financial bid b_i. We derive the first-order condition of seller i's expected profit maximization given that the other $j \neq i$ sellers bid according to the ratio $\tilde{\beta}(\theta_j)$.

$$\pi_i(\theta_i) = \Pr\left\{\frac{\theta_i}{b_i} > \max_{j\neq i} \tilde{\beta}(\theta_j)\right\} b_i = \Pr\left\{\tilde{\beta}^{-1}\left(\frac{\theta_i}{b_i}\right) > \max_{j\neq i} \theta_j\right\} b_i \tag{27}$$

$$= G\left(\tilde{\beta}^{-1}\left(\frac{\theta_i}{b_i}\right)\right) b_i. \tag{28}$$

$$\pi_i'(\theta_i) = \left(-\frac{g\left(\tilde{\beta}^{-1}\left(\frac{\theta_i}{b_i}\right)\right)}{\tilde{\beta}'\left(\tilde{\beta}^{-1}\left(\frac{\theta_i}{b_i}\right)\right)} \frac{\theta_i}{b_i^2}\right) b_i + G\left(\tilde{\beta}^{-1}\left(\frac{\theta_i}{b_i}\right)\right) = 0. \tag{29}$$

In symmetric equilibrium, $\tilde{\beta}(\theta_i) = \frac{\theta_i}{b_i}$. Thus, the first-order condition simplifies to the differential equation

$$-\frac{g(\theta_i)}{\tilde{\beta}'(\theta_i)} \tilde{\beta}(\theta_i) + G(\theta_i) = 0 \iff \frac{G(\theta_i)}{g(\theta_i)} = \frac{\tilde{\beta}(\theta_i)}{\tilde{\beta}'(\theta_i)} \tag{30}$$

The solution of this differential equation is

$$\tilde{\beta}(\theta) = kG(\theta), \tag{31}$$

where $k > 0$ is a constant. Thus, the ratio is indeed monotonically increasing in quality. The associated financial bid $\beta(\theta)$ is determined using $\tilde{\beta}(\theta) = \theta/\beta(\theta)$, i.e., $\beta(\theta) = \theta/(kG(\theta))$. In order to see that the above candidate is an equilibrium, we insert it into i's expected profit from above, (27).

$$\pi_i(\theta_i) = \Pr\left\{\frac{\theta_i}{b_i} > \max_{j\neq i} kG(\theta_j)\right\} b_i = \Pr\left\{\frac{\theta_i}{b_i} > kG\left(\max_{j\neq i} \theta_j\right)\right\} b_i \tag{32}$$

$$= \Pr\left\{G^{-1}\left(\frac{\theta_i}{kb_i}\right) > \max_{j\neq i} \theta_j\right\} b_i = G\left(G^{-1}\left(\frac{\theta_i}{kb_i}\right)\right) b_i = \frac{\theta_i}{k}. \tag{33}$$

Thus, i's expected profit is constant, implying that the ratio $\tilde{\beta}(\theta_i)$ is a best response to the other sellers' ratios $\tilde{\beta}(\theta_j)$, $j \neq i$. □

Proof of proposition 4. Suppose seller i's rivals bid the minimum financial bid, denoted by ε. Denote the c.d.f. of the largest quality among $n-1$ sellers by $G(\theta) = F^{n-1}(\theta)$ and its density by $g(\theta)$.

Then i's expected payoff associated with the financial bid b_i is

$$\pi_i = \Pr\left\{\frac{\theta_i}{b_i} > \frac{\theta_j}{\varepsilon}\right\} b_i = G\left(\frac{\theta_i \varepsilon}{b_i}\right) b_i \tag{34}$$

The first derivative w.r.t. b_i is

$$\pi_i' = G\left(\frac{\theta_i \varepsilon}{b_i}\right) - g\left(\frac{\theta_i \varepsilon}{b_i}\right)\frac{\theta_i \varepsilon}{b_i}. \tag{35}$$

Observe that the argument $\left(\frac{\theta_i \varepsilon}{b_i}\right) \in (0,1)$. Thus, if $G(\theta) < g(\theta)\theta$, seller i's expected payoff is decreasing in the financial bid, implying that $b_i = \varepsilon$ is a best response.

It is easy to verify that the condition $G(\theta) < g(\theta)\theta$ is equivalent to (11). □

Proof of proposition 5. In the second-score auction, each bidder i maximises expected income since there is no cost. Bidder i's own bid only affects the probability of winning but not the payoff in the event of winning. Bidder i's payoff in the event of winning, $p_i = \theta_i b_j / \theta_j$ is entirely determined by i's (fixed) quality, and the best rival j's quality-price ratio. Thus, bidder i needs to maximise the probability of winning. This is achieved by bidding the minimum bid. □

References

Arbatskaya, M., and H. Mialon (2010): "Multi-activity contests," *Economic Theory*, 43(1), 23–43.

Asker, J., and E. Cantillon (2008): "Properties of Scoring Auctions," *Rand Journal of Economics*, 39, 69–85.

Che, Y.-K. (1993): "Design competition through multidimensional auctions," *Rand Journal of Economics*, 24(4), 668–80.

Che, Y.-K. C., and I. Gale (2003): "Optimal Design of Research Contests," *American Economic Review*, 93, 643–51.

Corchón, L. C. (2007): "The theory of contests: a survey," *Review of Economic Design*, 11, 69–100.

Corchón, L. C., and M. Dahm (2010): "Foundations for contest success functions," *Economic Theory*, 43, 81–98.

Decarolis, F. (2010): "When the Highest Bidder Loses the Auction: Theory and Evidence from Public Procurement," *University of Chicago*, Working paper.

Dranove, D., and M. A. Satterthwaite (1992): "Monopolistic Competition when Price and Quality are Imperfectly Observable," *Rand Journal of Economics*, 23(4), 518–34.

Ewerhart, C. (2010): "Rent-seeking contests with independent private values," *University of Zurich*, Working Paper(#490).

Fey, M. (2008): "Rent-seeking contests with incomplete information," *Public Choice*, 135, 225–36.

Fu, Q., and J. Lu (2012): "Micro foundations of multi-prize lottery contests: a perspective of noisy performance ranking," *Social Choice and Welfare*, 38, 497–517.

Garfinkel, M. R., and S. Skaperdas (2006): "Economics of Conflict: An Overview," in *Handbook of Defense Economics*, ed. by T. Sandler, and K. Hartley, vol. 2, chap. 3. Elsevier.

Gershkov, A., J. Li, and P. Schweinzer (2009): "Efficient Tournaments within Teams," *Rand Journal of Economics*, 40(1), 103–19.

Gong, J., J. Li, and R. P. McAfee (2011): "Split-award Auctions with Investment," *Journal of Public Economics*, 96, 188–97.

Hanazono, M., J. Nakabayashi, and M. Tsuruoka (2012): "A theory of multi-dimensional auctions with non-quasi-linear scoring rules," *Working Paper*, March 10, 2012.

Hurley, T. M., and J. F. Shogren (1998a): "Asymmetric information contests," *European Journal of Political Economy*, 14, 645–65.

——— (1998b): "Effort levels in a Cournot Nash contest with asymmetric information," *Journal of Public Economics*, 69, 195–210.

Jia, H. (2008): "A stochastic derivation of the ratio form of contest success functions," *Public Choice*, 135, 125–30.

Katsenos, G. (2010): "Long-Term Conflict: How to Signal a Winner?," *University of Hannover, Working Paper*, April 20.

Ko, G. (2012): "Two-player Rent-seeking Contests with Private Values," *School of Humanities and Social Sciences, Nanyang Technological University, Working Paper*, June 1.

Konrad, K. (2008): *Strategy and Dynamics in Contests*. Oxford University Press, Oxford.

Malueg, D. A., and A. J. Yates (2004): "Rent seeking with private values," *Public Choice*, 119, 161–78.

Münster, J. (2009a): "Group constest success functions," *Economic Theory*, 41, 345–57.

——— (2009b): "Repeated Contests with Asymmetric Information," 11, 89–118.

Organisation for Economic Co-operation and Development (2000): *Greener Public Purchasing: Issues and Practical Solutions*. OECD, Paris.

Polishchuk, L., and A. Tonis (2011): "Endogenous contest success functions: a mechanism design approach," *Economic Theory*, forthcoming.

Prada-Sarmiento, J. D. (2010): "Uncertainty in conflicts," *CEDE, Univesidad de las Andes, Working Paper*, October 29.

Ryvkin, D. (2010): "Contests with private costs: Beyond two players," *European Journal of Political Economy*, 26, 558–67.

Schoonbeek, L., and B. M. Winkel (2006): "Activity and inactivity in a rent-seeking contest with private information," *Public Choice*, 127, 123–132.

Skaperdas, S. (1996): "Contest Success Functions," *Economic Theory*, 7(2), 283–90.

Wärneryd, K. (2003): "Information in Conflicts," *Journal of Economic Theory*, 110, 121–36.

——— (2009): "Multi-player contests with asymmetric information," *Economic Theory*, forthcoming.

Wasser, C. (2010): "Incomplete Information in Rent-seeking Contests," *Economic Theory*, forthcoming.

——— (2012): "A Note on Bayesian Nash Equilibria in Imperfectly Discriminating Contests," *Humboldt University of Berlin, Discussion Paper*, August 7.