Dipartimento di Scienze Economiche
Università degli Studi di Firenze

Working Paper Series

Design of Multidimensional Franchise Auctions by an Ignorant Principal

Pier Angelo Mori

Working Paper N. 13/2008
September 2008

Dipartimento di Scienze Economiche, Università degli Studi di Firenze
Via delle Pandette 9, 50127 Firenze, Italia
www.dse.unifi.it

Design of Multidimensional Franchise Auctions by an Ignorant Principal

Pier Angelo Mori*
Dipartimento di Scienze Economiche
Università di Firenze
via delle Pandette 9, I–50127 Firenze
pierangelo.mori@unifi.it

This version: August 2008

Abstract

The paper's main aim is to identify under which conditions the criterion of prior-independent optimality is applicable in the design of multidimensional franchise auctions. We first establish an impossibility result for second-score auctions by showing that in single-crossing environments necessary and sufficient condition for score functions to be optimal in this sense is that bidders have equal variable cost functions. Then we show that the result is not confined to the second-score format but holds for any scoring auction under stochastic independence. Therefore, a regulator who has no information at all about firms' costs cannot in such circumstances avail himself of prior-independent optimality as choice criterion. Conversely, if variable cost functions are equal across potential contractors, as is likely in certain public services markets, and the regulator knows it, it is possible for him to implement a prior-independent optimum by scoring bids according to the social welfare function under various auction formats, including first- and second-score auctions. This simple prescription however no longer applies if the regulator is ignorant about market demand too. In this case a fully rational choice of the score function is precluded, though it may be possible to make a reasonable one: a brief discussion of this point closes the paper.

JEL: D44, H57, L51.

*I am grateful to Fioravante Patrone for a few substantial suggestions and to Nicola Doni and Domenico Menicucci for their comments.

1 Introduction

The paper studies the choice by an ignorant principal of the score function through which multidimensional bids are evaluated in scoring auctions. The context of reference is franchise bidding in the field of public services where competition is on the tariff charged to customers and the subsidy granted by the regulator to the contractor.

The score-function choice was originally studied by Che (1993) and subsequent developments like Branco (1997); David et al. (2006); Naegelen (2002) for the case of mono-dimensional cost uncertainty. In these papers bidders' costs are assumed to depend on a single uncertain parameter, which means that the principal is uncertain either about marginal or fixed costs but not both, unless they are perfectly correlated. This assumption is of course hardly realistic, since there is no reason why the principal should be perfectly informed about fixed costs while being uncertain about marginal ones or viceversa and the only motivation is to remove the difficulties of mechanism design with multidimensional types (cf. the literature on the related problem of non-linear monopoly pricing, Armstrong and Rochet, 1999; Laffont et al., 1987; McAfee and McMillan, 1988). Here we relax this restriction by assuming that the regulator is uncertain about both costs and allow for imperfect correlation between them. A number of contributions are now available which analyze various aspects of scoring auctions under multidimensional uncertainty (Asker and Cantillon, 2008; Bichler and Kalagnanam, 2005; Bushnell and Oren, 1994; Müller et al., 2007), including optimal design (Asker and Cantillon, 2005), within the traditional Bayesian framework. In the present paper we depart from this approach and focus on the design of scoring auctions by a principal who has no prior on firms' costs.

The biggest obstacle for a principal who ignores the probability distribution of agents' types is that he is generally unable to a priori identify the optimal mechanism. One way out of this difficulty is to look for weaker choice criteria than expected-value optimality. For instance in the literature on selling mechanisms for electronic platforms there has recently been an upsurge of interest in the design of procedures for use by uninformed sellers. Choice criteria have been developed such as "competitiveness" (Goldberg and Hartline, 2003; Goldberg et al., 2006; Hartline and McGrew, 2005) which allow to design satisfactory, though not optimal auctions in the complete absence of seller information about buyers. The rationale for this strand of research is that, though some information can often be collected by the seller, its direct costs and inaccuracy grow steeply with market size and in mass markets standardized procedures which make no use of a priori information turn out preferable to Bayesian ones. A different research direction is based on the idea that mechanisms may not only convey information on the actual realization of private values but also on the distributions from which they are extracted. The task of designing prior-free mechanisms capable to elicit agents' type distributions is a general one and a literature is developing which tackles the problem in a variety of context (cf. Baliga and Vohra, 2003; Bergemann and Morris, 2005). Segal (2003) focuses on learning in auctions. His main result is that, by designing auction mechanisms so as to work as consistent estimators of the unknown cost distribution, they behave asymptotically as optimal mechanisms in the Bayesian sense, even without the auctioneer's knowledge of the prior. The problem with this approach is that, while suited to sales of large numbers of possibly small-value units, it is not appropriate for franchise-awarding contexts where transactions are typically of high value and involve few bidders. In a similar vein and more closely to our topic, Beil and Wein (2003) study a multistage multidimensional auction mechanism which allows a completely ignorant auctioneer to gain information from myopic suppliers about their costs by progressively changing the score function from round to round. This approach, however, has the important limitation of being applicable only to informal auctions, where the scoring rule is not set and announced beforehand—which is generally not the case when the auction object is a franchise and the auction is called by a public body.

Here we take up the issue from a different angle. We go back to the fundamentals of auction design under ignorance by investigating what can be attained within the class of scoring auctions with the least possible sophistication in design and the least possible information. Though an ignorant principal will generally be unable to single out an optimal mechanism, it need not always be so. Indeed, there may exist mechanisms that dominate every other irrespective of the prior—i.e. prior-independent optimal—and in such a case even an ignorant principal is able to identify and implement them. Our main aim here is to find out if and when such mechanisms exist.

Two sets of results are developed. We start by asking whether in the class of second-score auctions there exist prior-independent optimal ones. More precisely, we ask if there are score functions which, when applied to a second-score auction, obtain a better outcome than any other, irrespective of the prior. When variable cost functions display the single-crossing property, it is indeed possible to characterize the environments which admit prior-independent optimal score functions: such optima exist if and only if all variable cost functions are equal across bidders and they always include the social welfare function itself (Proposition 1). The "only if" part of the characterization is in fact an impossibility result for the universal domain of priors. A completely ignorant regulator is one who has no information at all about environment parameters, i.e. who knows that the prior belongs to the universal prior domain. The result tells that there exists no optimal score function with respect to all priors in this domain, and hence the regulator will be unable to make use of the prior-independent optimality criterion. By contrast, if variable cost functions are equal across firms, there does exist a prior-independent optimal score function for second-score-auctions. Since this condition in fact amounts to a restriction of the prior field (though not knowing the prior, the regulator knows that it belongs to a well-defined restricted domain), what we actually find is prior-independent optima with respect to the priors belonging to a restricted domain. This result seems especially relevant for the market of public-service franchises: for many services such as bus transportation, waste collection, etc., costs are indeed likely to display this property and the regulator may be aware of this even if he has no prior on them.

A second set of results extends the basic insight in several directions and allows to cast light on the existence of prior-independent optima in a wider context. One might suspect that, when prior-independent optimal second-score auctions do not exist, there may nonetheless exist optimal auctions of different kinds: as aa matter of fact, we are able to prove that in a few relevant cases this is certainly not true. If variable and fixed cost parameters are i.i.d. across potential contractors, it is possible to compare welfare levels supported by different auction mechanisms and the result is that all scoring auctions with the same score function are welfare-equivalent. Then if a prior-independent optimal score function exists, it must also be optimal with respect to the subset of priors with i.i.d. costs, which means that a prior-independent optimal second-score one should exist even when variable cost functions differ across firms: since our characterization result for second-score auctions rules out this, the conclusion is that no prior-independent optimal auction of any kind ever exist in the universal domain of priors (Proposition 2). On the positive side, when variable cost functions are equal across firms and certain additional conditions hold (namely, fixed costs independently and identically distributed across bidders) we can prove that first-score and second-score auctions which employ the social welfare function to score bids are optimal in expected value both with respect to all priors in this restricted domain and all kinds of auction mechanisms (Proposition 3).

These prescriptions all have in common the use of the social welfare function for scoring bids. Of course, this requires that the regulator knows it, which is however not always the case: if the welfare function is specified in the usual way as a weighted sum of consumer and producer surpluses, a regulator who ignores market demand ignores it too. Thus even under the equality of variable cost functions across bidders the possibility of a fully rational choice by an ignorant regulator vanishes altogether and a new problem arises. When a generic score function is applied to a first- or second-score auction under the equality of variable-cost parameters across bidders, our results tell that what is actually maximized is an objective coincident with that score function. In other words, the regulator in fact behaves as if he pursued this objective, which will in general be different from his natural one. To stress the difference we call this second sort of objectives pseudo-objectives. Then, instead of assessing different score functions against an a priori given natural objective, he has to find a pseudo-objective that is reasonable in some sense. How this is to be done is a question that falls outside the paper's scope but we want to stress one point. Though ignorant in a Bayesian sense, regulators may and do often have a priori information that can help restrict the choice field in a significant way. The restriction process will not generally single out a unique score function but may however be enough to drop from consideration a good number of them and even some that are commonly used in regulatory practice.

The paper is organized as follows. After laying out the model in section 2, we study prior-independent optimal second-score auctions in section 3. Section 4 presents an impossibility result for generic scoring auctions. In section 5 a few results are developed about optimal auctions under the restriction of equal variable cost functions across firms. In section 6 we tackle the problem of ignorance about market demand and we show how our results on prior-independent optimality allow in certain cases to devise pragmatic solutions to the auction design problems under this further restriction.

2 The model

A regulator is to select a contractor for the provision of a public service. Firm i's cost of service supply y is $c(y, m_i) + F_i$, where F_i is the fixed cost, $0 < F_i < \infty$, and m_i is a variable-cost parameter which varies over some interval ($c(0, m_i) = 0$ for all m_i). These parameters are private information to the firm. We denote service tariff

p and the (possibly negative) subsidy granted to the contractor s. Service demand at every p is $y(p)$ and the contractor's profit is

$$\pi(p, s, m_i, F_i) = py(p) + s - c\big(y(p), m_i\big) - F_i.$$

Demand and cost functions are of usual shape (respectively decreasing and increasing in y) and continuous. The regulator's objective is to maximize social welfare which we assume to be specified in the standard fashion after Baron and Myerson (1982)

$$W(p, s, m, F) = CS(p) - s + \alpha\pi(p, s, m, F) \tag{1}$$

where $CS(p)$ is the consumer surplus and α is a distributive parameter such that $0 \leqslant \alpha < 1$. Moreover we assume that the regulator has no prior on firms' costs (later on we shall discuss the implications of ignorance about demand functions as well).

Scoring auctions are perhaps the most frequent procedure for awarding franchises and quite common in other contexts too. The focus of our analysis is the design of such auctions and to this purpose we assume that the regulator's choice is a priori restricted to them to the exclusion of other formats (like e.g. the full-surplus extraction auctions of Crémer and McLean, 1988, and McAfee and Reny, 1992). A *scoring auction* is a particular sort of one-shot multidimensional auction where bids are contract proposals, in our case denoted (p, s), each bid gets a score $V(p, s)$ according to some score function $V\colon \mathbb{R}_+ \times \mathbb{R} \to \mathbb{R}$, the highest-score bidder wins and the implemented contract is chosen by the winner under the constraint that it matches some score computed from realized individual scores through a predetermined algorithm. Throughout we assume that losers pay nothing. This class of auction rules includes two popular awarding mechanisms like first- and second-score auctions which have been extensively studied in the literature (Asker and Cantillon, 2008; Branco, 1997; Che, 1993) and which will be a major focus of our analysis too. Under the first-score rule the implemented contract is the winner's bid. Second-score auctions are extensions of second-price ones where the winner is free to choose any contract with a score at least as great as the second-highest one.

Score functions are chosen by the regulator and the choice is restricted to functions that are continuous, decreasing in p and strictly decreasing in s. Auction mechanisms also usually specify a reserve value for each bid variable but here, in view of the regulator's severe informational constraints, are assumed as given and hence fall outside the scope of our analysis. As a consequence, in this model the choice of an auction mechanism in fact reduces to the choice of a score function and a rule of contract determination based on realized individual scores. Moreover, to simplify calculations we also assume that, unless indicated otherwise, reserve values are infinite, as in Branco (1997) and Che (1993), and that all firm types always participate in the auction as e.g. in Branco (1997); Che (1993); Riordan and Sappington (1987) (hence we eschew the problems of exclusion/non-participation that typically arise in multidimensional adverse-selection environments, cf. Armstrong, 1996; Jehiel et al., 1999; Rochet and Choné, 1998; Rochet and Stole, 2002).

In sum, our model coincides with that of Che (1993) except in two main respects—that here we have two uncertain cost parameters instead of one and, above all, the regulator has no prior on cost functions. Without such information full rationality is generally out of reach for an ignorant and unsophisticated regulator but there are cases where even such a decision maker can make fully rational choices: this is the object of our enquiry in the next sections.

3 Second-score auctions

The task of a Bayesian regulator is straightforward: to select the score function which implements the contract yielding the highest expected welfare with respect to his prior distribution of m_i and F_i. A regulator who has no prior information is instead unable to work out the expected welfare attainable through alternative mechanisms. There is however a case in which this limitation is irrelevant and even an unsophisticated regulator can compute social optima—when there exists a *prior-independent optimal* score function, i.e. one which maximizes expected welfare for all possible priors. A special case of it, with which we shall extensively deal in the subsequent analysis, is when a score function is *pointwise optimal*, i.e. maximizes not only expected welfare for all priors but also ex-post welfare state by state. The study of such optima is one of the main themes of this section. As we shall see, restrictive though the prior-independence criterion is, the optimal design problem does admit solutions under it in a few economically significant environments.

We proceed in two steps. A social welfare function is independent of the state of the world when it is such that $W(p, s, m, F) = \hat{W}(p, s)$ for every p, s, m, F. We first focus on welfare functions of this kind with $\hat{W}$ strictly decreasing in s and characterize prior-independent optimal auctions in the class of second-score ones under this restriction. Then we will extend the results to different auctions and more general welfare functions.

Second-score auctions have dominant-strategy equilibria like second-price ones: the following is a preliminary result that is obtained as a straightforward extension to our context of the standard characterization of second-price equilibria (the proof is an easy exercise and is therefore omitted).

Lemma 1 *Second-score auctions with a score function V continuous and monotonic decreasing in s always have a dominant-strategy equilibrium where equilibrium bids are those which maximize the bidder's score subject to the zero-profit constraint.*

A convenient way to approach prior-independent optimality is to start from pointwise optimality which is easier to characterize. Score functions that are optimal in the former sense are also optimal in the latter and, as we shall see, in many situations the converse holds true too. In the cases where this one-to-one correspondence holds the characterization of prior-independent optimal score functions reduces to a characterization of pointwise optimal ones. In this subsection the main focus is on the latter and the first result is a sufficient condition for the existence of a pointwise-optimal score function.

Lemma 2 *If variable cost functions are equal across firms for all production levels the regulator's welfare function is a pointwise-optimal score function for second-score auctions.*

Proof. Consider any realization of the parameters under the restriction $m_i = m$ for all i, i.e. $m, F_1, F_2, \ldots, F_n$, and without loss of generality assume that $F_i \leqslant$

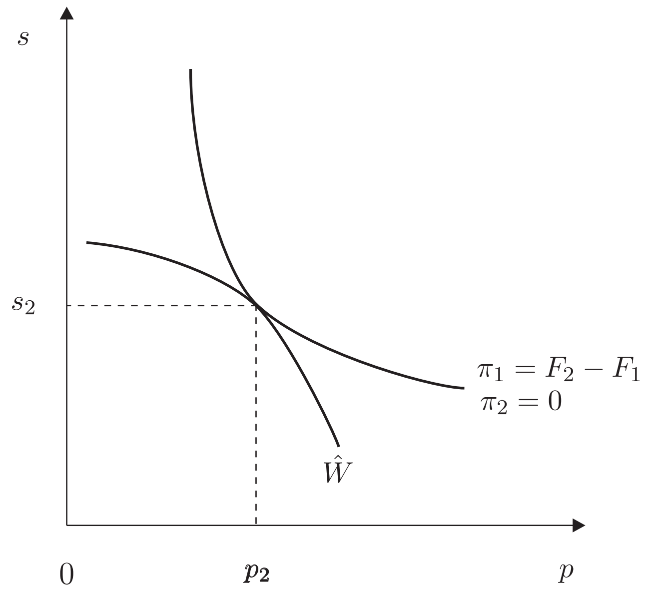

Figure 1: Equality of variable cost functions across firms is sufficient for pointwise optimality

$F_{i+1}, i = 1, \ldots, n-1$. By Lemma 1, under welfare function $\hat{W}$ as score function, the second-highest bid belongs to bidder 2's zero-profit locus and to the highest iso-welfare curve, as in Fig. 1. This is also the auction outcome since the winner, tough free to choose a contract different from (p_2, s_2), has no incentive to do so (this is just the most advantageous contract among those totalling a score equal to the second-highest one): whenever bidders differ only in fixed costs, and hence all firms' iso-profit curves have the same shape, the implemented contract and the second-highest bid always coincide. Now take any score function that at least in one state induces the second-highest bidder to choose a point on his iso-profit curve $\pi_2 = 0$ that does not maximize the regulator's welfare (if this were not true, the score function would be equivalent to $\hat{W}$). Then also the contract implemented under the new score function will entail a (weakly) lower welfare than that implemented under $\hat{W}$ (i.e. (p_2, s_2) of Fig. 1). In conclusion, by adopting a score function different from $\hat{W}$ the regulator obtains a (weakly) lower welfare in every $m, F_1, F_2, \ldots, F_n$. In other words, $\hat{W}$ dominates every V as score function in every state. □

Remark 1 When parameter m is common knowledge we are in the particular case studied by Che (1993) and the existence result can be obtained in a different way too. Che shows (Proposition 4) that in his special set-up the distortion of the optimal score function relative to the social welfare depends multiplicatively on the cross-derivative c'_{ym}. Since the latter annihilates if m is common knowledge and fixed costs are uncertain, the optimal distortion annihilates too and the optimality of the social welfare function as score function immediately follows.

The economic intuition underlying this result is quite simple. When m_i is the same for every i, bidders differ at most in fixed costs (i.e. have the same isoprofit

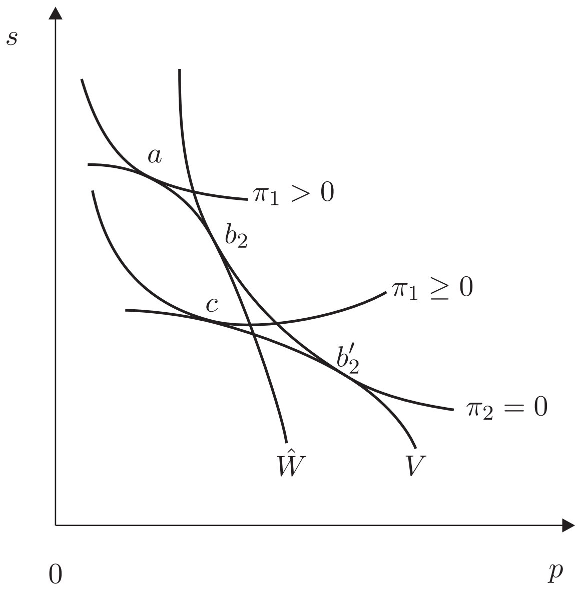

Figure 2: A $\hat{W}$ that is not pointwise optimal

curves with a different scale, $\pi(p, s, m, F_i) = \pi(p, s, m, F_j) - F_i + F_j, \forall i, \forall j)$. Then, since by definition the second-highest bid is the welfare-maximizing one subject to the zero-profit constraint, in every state this also coincides with the highest-profit bid for the winner subject to the second score. Therefore a displacement of the score function from $\hat{W}$ shifts the second bid and the implemented contract to a (weakly) lower-welfare position along the curve $\pi_2 = 0$.

To see what happens when the implemented contract does not coincide with the second highest bid let us have a look at Fig. 2. Here the welfare function $\hat{W}$ is used to score bids and a and b_2 are respectively the implemented contract and the second-score bid in the state of the world to which the depicted curves π_i refer (to simplify notation, here and in the following we assume without loss of generality that the i-th ranked bidder coincides with i). In the given situation there also exists another score function V (throughout V will denote generic score functions) under which the second-score bid is b_2' and the implemented contract is a couple c such that $\hat{W}(c) > \hat{W}(a)$. This means that in the state of the world represented in the figure $\hat{W}$ supports an outcome inferior to that supported by V and hence is not pointwise optimal. Note that for the result it is not required that V be welfare-improving over $\hat{W}$ in all states of the world but it is enough that it is so in at least one state. It is also to be noted that the mere existence of a curve like V in some state is not enough to conclude that $\hat{W}$ is not optimal. Suppose that in the figure we had $\pi_1(a) = 0$ in place of $\pi_1(a) > 0$, i.e. that the iso-profit curve through a were the zero-profit one. Then there would actually be two equilibria in that state—one with a and the other with b_2 as implemented contract —both of which are full-information optima, i.e. contracts that could be indifferently chosen by the regulator if he had full knowledge of the state of the world: when this is the case, the divergence between first and second-score bids in one state does not cause $\hat{W}$ to be dominated. Except in this special case, however, $\hat{W}$ cannot be pointwise optimal if there exist curves like V of Fig. 2 and this occurs whenever the implemented contract and the second-score bid do not coincide. Lemma 4 of the appendix provides a formal and general proof which applies to any score function, not only $\hat{W}$: in conclusion, the pointwise optimality of any score function requires in every state either the coincidence between implemented and second-score bid or that the implemented contract be a full-information optimum.

Intuitively, the divergence between implemented contract and second bid in one state reflects different trade-offs between tariff and subsidy across bidders (i.e. iso-profit curves are not vertical translations of one another). This causes bidders to reply differently to displacements of the score function and it is just these differences that allow to attain welfare improvements in that state. In particular, in the example of Fig. 2 the winner has a lower marginal substitution rate of s for p, i.e. he demands less than the second-ranked firm as compensating tariff increase for a given subsidy reduction. Then, by switching from $\hat{W}$ to a flatter score function like V (i.e. by putting less weight on tariff) the second bidder is induced to substitute subsidy for tariff and the new iso-score touches a lower iso-profit curve of the winner. If variations are not too large, the shift of the implemented contract from a to c brings about an improvement in welfare as is represented in the figure (note that here bidder 2 over-reacts, as it were, by "asking" for a larger tariff increase and thus reduces the welfare level attached to the second bid, but this is not relevant to the regulator who is only interested in the welfare attached to the implemented contract).

This result allows to characterize environments—essentially cost functions—which admit pointwise optimal score functions for second-score auctions. As we have seen, the alternative requirements for a score function to be pointwise optimal are that in every state it supports a full-information optimum or that the implemented contract coincides with the second-score bid. On the other hand, if fixed costs F_i vary continuously, the latter requirement is always to be met and hence there will never exist a pointwise optimal score function unless the coincidence condition holds in *every* state, as is easy to verify. Suppose that in some state we have $a \neq b_2$ and the winner's profit is null, $\pi_1(a) = 0$, i.e. outcome a is a full-information optimum in that state (without loss of generality we keep assuming that 1 is the winner and 2 the second-highest bidder). If F_1 is the winner's fixed cost, a slight reduction of it will cause no change in the outcome if parameters $F_2, \ldots, F_n, m_1, \ldots, m_n$, do not change (the winner will still be 1, a the implemented contract, and so on) but the winner's profit will now be positive. This means that there exists at least one state of the world where the outcome of a second-score auction is not a full-information optimum, which proves that the coincidence of implemented contract and second bid in every state is necessary here for pointwise optimality (a formal statement of this argument is found in the proof of Proposition 1 in the appendix). Then what we must look for is environments where the coincidence condition holds in every state of the world.

We have already found a sufficient structural condition which guarantees it—the equality of variable cost functions across firms, $m_i = m, \forall i$ (Lemma 2)—but unfortunately this condition is not necessary as Fig. 3 makes clear. In this case all isoprofit curves $\pi_1 \geqslant 0$ lie in the epigraph of $\pi_2 = 0$ and $\hat{W}$ is optimal as score function in the state to which the depicted curves refer (if we take any score function V, the iso-score curve through the second bid b_2 has no intersection with the interior of the epigraph of $\pi_2 = 0$ and hence at no point on it profits can be larger than at a).

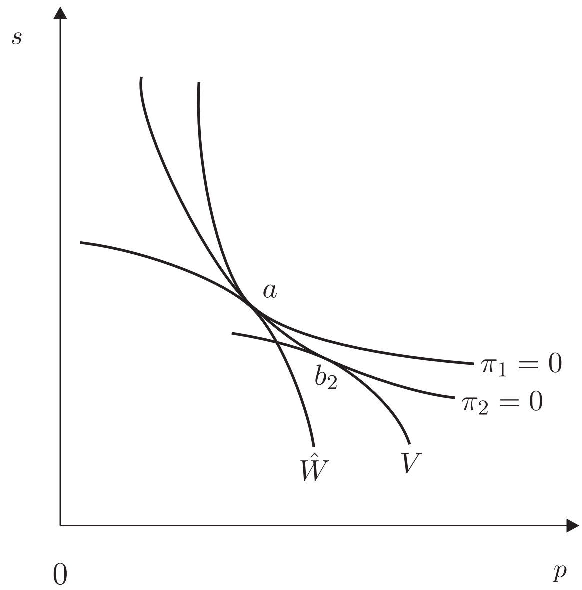

Figure 3: Equality of variable cost functions across firms is not necessary for pointwise optimality

If this holds true in every state of the world, there clearly exists no V that strictly dominates $\hat{W}$ pointwise and therefore this is pointwise optimal despite variable-cost functions are not equal across firms.

A natural question at this point is whether by coupling the equality condition with appropriate restrictions on the environment one can come to a complete characterization of pointwise optimal score functions. The answer is positive but it is possible to identify more than one class of such functions in this way. Here we present a characterization result that is especially interesting for it holds under a property which is common to many models of mechanism design – the *single-crossing* property for cost functions, i.e.

$$\frac{\partial^2 c}{\partial y \partial m_i} > 0, \ \forall y, \forall m_i.$$

Under it different types' iso-profit curves cross once and only once and the case of Fig. 3 never occurs: with differentiable functions contract a has the characteristics of the figure only if the implemented tariff is such that $c'_m\big(y(p), m_1\big) = c'_m\big(y(p), m_2\big)$ and $c\big(y(p), m_1\big) - F_1 = c\big(y(p), m_2\big) - F_2$, but the single crossing property is incompatible with the first condition for $m_1 \neq m_2$. In the class of environments displaying this property there holds the following characterization of prior-independent optimal score functions (proof in the appendix).

Proposition 1 *In environments where cost functions display the single-crossing property there exist pointwise optimal score functions for second-score auctions if and only if variable cost functions do not vary across firms for all production levels, and one of them is $\hat{W}$.*

Proposition 1 has immediate implications for prior-independent optimality. In second-score auctions any pointwise-optimal score function is also prior-independent optimal in the universal domain of priors and viceversa. The first implication is obvious and requires no comment. The reverse can easily be established ex absurdo by assuming that V is prior-independent optimal but not pointwise optimal. For this to be true there must be at least one state where some $\tilde{V}$ supports a strictly better contract with respect to social welfare $\hat{W}$ than that supported by V (recall that in dominant-strategy equilibria both outcomes and ex-post utilities are independent of the prior in every state). Then, by taking a prior with a sufficiently large probability mass concentrated on this state, we can make the expected social welfare attained by $\tilde{V}$ larger than that attained by V with respect to this prior, thus contradicting the assumption that V is prior-independent optimal. In other words, in the universal domain of priors there is a one-to-one correspondence between pointwise and prior-independent optimality in second-score auctions. This and Proposition 1 imply that in the class of single-crossing environments there is a boundary beyond which the search for a prior-independent optimal score function fails: such functions simply do not exist for second-score auctions except when variable-cost functions are equal across firms. In other words, there exists no score function prior-independent optimal with respect to all possible priors. Therefore for a completely ignorant regulator who is unable to delimit the prior domain in any way, and hence has to work with the universal one, the prior-independent optimality criterion will be wholly useless. However, even a non-Bayesian regulator may sometimes have available a priori information that, though insufficient to identify the true prior, allows to restrict the domain which contains it in such a way that the prior-independent optimization problem has a solution constrained to the restricted domain. The "if" part of Proposition 1 gives a precise answer to this question: in the family of environments characterized by $m_i = m_j$ for every i, j, the social welfare function is prior-independent optimal for second-score auctions.

Note that, even though we have so far assumed that welfare functions are independent of the winner's type, i.e. are of the form $\hat{W}(p, s)$, the result holds with the more general functions W defined by Eq. (1) too. When variable cost functions differ across firms and the single-crossing property holds, implemented contract and second bid differ in every state and the argument of Lemma 4 of the appendix (Fig. 2) can be replicated for functions W too: the negative conclusion on the existence of a prior-independent optimal V then remains valid even when the social welfare function is of form W. In the next section this impossibility result is extended to generic auction mechanisms.

4 Impossibility of optimal generic auctions in the universal prior domain

Proposition 1 allows to pinpoint precisely in which prior domains the prior-independent optimization problem has solutions for *second-score* auctions. A further question is if anything substantial changes when we remove the restriction on the auction type, and particularly if better outcomes are attainable from the regulator's standpoint through other auction mechanisms. The comparison in terms of expected welfare between different auction structures is generally hard to make unless score functions are separable—an assumption on which most of the existing literature on regulatory design and scoring auctions is founded. Here we follow suit and from now on we restrict ourselves to social welfare functions that are linearly separable in subsidy, i.e. $\hat{W}(p,s) \equiv w(p) - s$, and also restrict the choice to separable score functions, $V(p,s) \equiv v(p) - s$.

The interesting fact is that under the single-crossing property allowing for auction rules other than the second-score one brings about no welfare improvement, since either there exists no solution to the prior-independent optimization problem or the solution is the social welfare function itself in essentially the same cases as before.

To check this we first verify that an impossibility result for the universal prior domain holds true here too. The basic instrument for comparing different auction mechanisms is the revenue-equivalence argument. Asker and Cantillon (2008) develop a version of it for scoring auctions and multidimensional types that fits precisely into the present framework. Consider the set of stochastically independent distributions of types $(\theta_1, \theta_2, \ldots, \theta_n)$, where $\theta_i = (m_i, F_i)$. Firm i's maximum apparent social surplus

$$k_i(m_i, F_i) \equiv \max_p \left[v(p) + py(p) - c(y(p), m_i) - F_i\right]$$

is called the *pseudotype* of that firm by Asker and Cantillon (the expression in square brackets is just ordinary surplus with the willingness to pay replaced by the tariff score $v(p)$). Under the conditions assumed here (namely, separability and risk neutrality) scoring auctions are equivalent to auctions where bidders make statements about their own pseudotype and the highest one wins (Asker and Cantillon, 2008, Theorem 1). Since these auctions are actually equivalent to standard price auctions where reservation values are replaced with pseudotypes, they enjoy the same basic properties and in particular they are awarded to the highest pseudotype in every state of the world (equilibrium strategies are monotonic increasing in the pseudotype). Then, if any two scoring auctions with the same separable score function are such that the winner is the same in every vector of pseudotypes and the lowest pseudotype gets the same expected utility in both of them, the auctioneer's expected utility is invariant across the two auctions (Asker and Cantillon, 2008, Theorem 2).

In light of this revenue-equivalence result it is immediate to verify that even for generic scoring auctions there exists no prior-independent optimal score function in the universal prior domain. Suppose ex absurdo that there exists a score function V^* that is prior-independent optimal in some scoring auction A, i.e. there exist no other score function nor auction rules in the class of scoring auctions that allow to attain a higher level of regulator expected welfare. Take now the subset of stochastically independent priors: if the given score function is to be prior-independent optimal, the dominance relation must hold for these too. Asker and Cantillon's revenue-equivalence result implies that with stochastically independent priors the auctioneer's expected welfare is the same both in auction A and in a second-score one under the same score function: therefore if V^* is prior-independent optimal under A it must be so in the second-score auction too. But the latter is impossible, since by Proposition 1 such score functions never exist for second-score auctions unless $m_i = m_j$, and hence do not exist in the set of stochastically independent priors either. Then a fortiori there do not exist in the universal prior domain either, thus contradicting the assumption that V^* and auction mechanism A are prior-independent optimal in it. Note that by an argument similar to that of the previous subsection, this negative conclusion actually holds for general welfare functions whose values depend on the winner's type too, $w(p, m, F) - s$. All this is summarized in the following proposition.

Proposition 2 (Impossibility in the universal prior domain) *In environments where cost functions display the single-crossing property there exist no prior-independent optimal score functions over the universal prior domain for any kind of scoring auction.*

In conclusion, if we look for a prior-independent optimal score function in the universal domain of priors, or in other words if the decision-maker has no a priori information at all about firms' costs, there is no solution under whatever scoring auction. But again, if the auctioneer has some ex ante information that allows to restrict the prior field appropriately, there do exist auctions that are prior-independent optimal in the restricted domain and over *any* scoring auction, not only second-score ones, as we show in the next section.

5 Optimal generic auctions in a restricted domain

From now on we assume that the regulator, though ignoring the exact prior shape, has the following a priori information: 1) $m_i = m_j$ for every i, j, 2) fixed costs F_i are equally and independently distributed according to a generic distribution function $F_i \mapsto G(F_i)$ with density $F_i \mapsto g(F_i)$, which is also common knowledge to firms, 3) m and F_i are stochastically independent. Note that to analyse this case we cannot apply Asker and Cantillon's equivalence result, since individual types θ_i are not stochastically independent, and then we have to prove an ad hoc result. Since all firms have the same m and the only uncertain parameters for them are their competitors' fixed costs, we can calculate the expected value of a given score function with respect to $G(\cdot)$ in the usual way. That is, we first work out the winner's expected profit by applying the revelation principle. Then we verify through standard differential methods that any two auction mechanisms that induce the same ex-post allocation are equivalent in terms of expected score (for the proof see the appendix).

Lemma 3 ("Revenue equivalence") *If $m_i = m$ for all i, m and F_i are stochastically independent and moreover F_i are i.i.d., all auctions with the same score function additively separable in s that award the contract to a firm with the lowest F_i and entail the same tariff outcome $m \mapsto p^*(m)$ realize the same expected value of the score function.*

The lemma establishes the equivalence with respect to expected score of all auction rules that: 1) are efficient (i.e. the lowest type F_i wins and the highest type has a null probability of winning), 2) induce the same tariff function $m \mapsto p^*(m)$ as outcome. Now note that both first- and second-score have the two properties. Since parameter m_i is equal for all firms, the winner in both first- and second-score auctions is the bidder with the lowest F_i (the probability of winning is $(1-G(F_i))^{n-1}$).

Moreover, given any separable score function of the type $V(p,s) \equiv v(p) - s$, in first- and second-score auctions the tariff is

$$p^*(m) = \text{argmax}_p \left[v(p) + py(p) - c(y(p), m) - F_i - \bar{V} \right]$$

for every m, irrespective of the score $\bar{V}$ to be met by the winner (this is actually true for any nth score auction). Then by Lemma 3 first- and second-score auctions are equivalent in terms of expected score. Moreover, there is no other auction that can do better in expected terms than first- and second-score ones (or, for that matters, than any efficient auction with the same tariff outcomes): welfare functions additively separable in s dominate any other scoring criterion in first-score auctions for any distribution function $G(\cdot)$ too. To see this, suppose ex absurdo that there exists a score function that obtains a higher expected welfare under the first-score rule. By the equivalence result just established dominance should also hold under the second-score rule but this would contradict Lemma 2, which asserts the reverse dominance order state by state and hence in expected value as well.

The reason is intuitively clear if one considers that the tariff implemented under them is ex-post efficient, whereby gross social surplus is the largest possible, and the selected bidder is that with the lowest F_i, which implies that social surplus net of fixed costs is maximized too. For an auction rule to improve on first- and second-score ones in terms of the regulator's welfare, it would be necessary that it pushed down the contractor's expected profits. But, whatever the prior, such an auction mechanism does not exist, since in the present conditions to ensure incentive compatibility the contractor must always be granted the same expected profit, irrespective of the auction rules (score function). Therefore, for the regulator there is nothing better than an auction that obtains the highest social welfare in every state, as both first- and second-score auctions actually do when bids are scored in terms of social welfare. In other words, first- and second-score auctions are optimal in expected welfare with respect to any distribution function $G(\cdot)$. All this is summarized in the following proposition.

Proposition 3 *With $m_i = m$ for all i, fixed costs i.i.d. across firms and social welfare additively separable in s, first- and second-score auctions which use the social welfare function to score bids are optimal in expected welfare with respect to any distribution function $G(\cdot)$.*

Under the proposition's conditions second-score and first-score auctions with social welfare as score function are optimal irrespective of the prior and can thus be used interchangeably by a non-Bayesian regulator who just knows that F_i's are independent and that variable-cost functions do not vary significantly across firms. Note that these auctions (or for that matters any auction where the most efficient bidder wins and an ex-post efficient allocation is implemented) are optimal not only over other scoring auctions but also over *any mechanism*, i.e. there is no mechanism, however structured, that can improve on them.

That prior-independent optimal score functions exist in particular environments is not a novelty. Che's (1993) results for instance imply that they exist under the assumption that all firms have the same variable cost parameter and this is known with certainty by everybody (cf. Remark 1 above). Che's model, however, is unable to say anything about the possibility of prior-independent optimality when the assumption is relaxed and specifically whether the equality of variable cost parameters across firms keeps playing a role when the principal is uncertain both about marginal and fixed costs. Here we have made a step forward by identifying a general class of environments—single-crossing ones—where the equality property is necessary and sufficient for prior-independent optimality.

So far we have worked with welfare functions that do not depend on the state of the world. Now let us go back to Eq. (1). The previous results directly apply to it for $\alpha = 0$ and ensure that when $m_i = m$ for all i, the regulator should simply use the net consumer surplus

$$\hat{W}(p, s) = CS(p) - s \tag{2}$$

to score bids. But what if social welfare depends also on firms' profits besides consumer surplus, i.e. $\alpha > 0$? As a matter of fact the answer does not change: net consumer surplus remains optimal for scoring bids. To see this, suppose for a moment that the regulator knows the state of the world $((m, F_1), \ldots, (m, F_n))$ with certainty. Since nothing in Lemma 2 depends on whether welfare is affected or not by the state of the world, the lemma holds for general welfare functions $(p, s) \mapsto W(p, s, m, F_i)$ too. In other words, by the argument of Lemma 2, we can rule out the existence of a V supporting better outcomes than those obtained by employing W as score function. Of course, the problem is that the regulator actually ignores the state of the world and W cannot be used in practice. However, net consumer surplus Eq. (2) is equivalent to Eq. (1), since the second-score auction outcome is the same under both of them in every state of the world. Therefore, net consumer surplus is optimal among all score functions even with $\alpha > 0$. The intuition is straightforward. In second-score auctions the winning firm's profit is always equal to $(F_2 - F_1)$ irrespective of the equilibrium tariff, i.e. a change in the score function can only affect the implemented tariff and the consumer surplus. Then, if the latter has a positive impact on welfare, however small, a score function that induces the highest possible consumer surplus is a solution to the problem. From Lemma 2 we know that to obtain the highest consumer surplus the trick is to use it to score bids: therefore, even when social welfare is specified as Eq. (1), it is in fact enough to adopt Eq. (2) as score function (for a formal proof see the appendix).

In the previous analysis a critical role is played by the equality of variable cost functions across firms. It is then natural at this point to ask how likely such a situation is in regulatory practice. A brief discussion of the point will suffice here. Take for example service contracts for the operation of bus lines. Such contracts usually impose on the contractor the duty to drive certain routes with given frequency, i.e. to total a fixed amount of kilometres per time unit. In these circumstances the marginal service cost—that here coincides with the cost of an extra passenger/kilometre—is unlikely to differ significantly across potential contractors, since it essentially reflects variations in fuel consumption and tyre wear which can hardly differ among them. By contrast, fixed costs depend on the efficiency of the firm's overall organization which can considerably vary from firm to firm. Therefore cost differentials among potential contractors, if any, are in this case to be ascribed to fixed costs only. Similar remarks hold for other services of akin nature like waste disposal and also for some that are entirely different from transportation. Let us consider by way of example a contract for building and exploiting a bridge (but the same argument applies to motorways, tunnels, etc.). Construction entails large fixed (sunk) costs which may vary sensibly across potential contractors according to their efficiency. Operating costs are mostly fixed (non-sunk) too: the bridge requires the employment of so many surveillance staff, the light of so many lamp columns, etc., irrespective of how many vehicles cross it, unless it is kept closed (in which case these costs are zero). Maintenance costs are instead partially variable with use (like e.g. the cost of asphalt paving) but the variable part is again unlikely to differ significantly across potential contractors, who usually contract out maintenance and have access to the same subcontractors. These remarks suggest that the equality of variable cost functions across firms is not an unlikely event in public services markets and therefore prior-independent optimality is a potentially relevant choice criterion for auction practice in this field.

6 An extension

The main upshot of the previous analysis is that a completely ignorant regulator will be unable to act according to the prior-independent optimality criterion, given the lack of an optimal score function in the universal prior domain. However, if there is available a priori information that allows to restrict the prior domain appropriately, there may indeed exist constrained solutions according to this criterion too and even a non-Bayesian regulator may be able to make a fully rational choice. For an important class of environments we have characterized the structural conditions and the a priori information required for a non-Bayesian regulator to implement a prior-independent optimal auction: under the single-crossing property for cost functions, the only case where (constrained) prior-independent optimal score functions exist is when variable cost parameters do not vary across firms. Therefore, when the latter do not meet this requirement or the regulator is completely ignorant about them, there will be needed weaker optimality criteria than prior-independence. This topic is left to future research but before closing we want to discuss a further problem that may arise in practice.

Our model's basic prescription is that, when certain conditions are met—essentially $m_i = m$ for all i—, bids must be scored by the social welfare function. Of course, such a procedure is viable only if the welfare function is known to the regulator but this is not always the case. As is usual in regulation theory, we have so far assumed that production costs are uncertain to him while market demand is common knowledge. By contrast it is not implausible that a regulator who is ignorant about firms' costs is ignorant about market demand too. The problem is that in such a case the social welfare function Eq. (1) is unknown too and, even if it is prior-independent optimal as score function, it cannot be implemented.

There is more to it. The results of the previous sections apply not only to social welfare functions $(p, s) \mapsto W(p, s)$ but actually to any function $(p, s) \mapsto V(p, s)$ that were taken as the regulator's objective. This means that any score function is prior-independent optimal with respect to itself in the restricted prior domain identified by $m_i = m$ for all i. In order to distinguish generic objectives $V(p, s)$ from the social welfare function we call them *pseudo-objectives.* From a practical standpoint, the regulator knows that in such a situation, whatever score function he chooses, he will behave as if he optimized that function as pseudo-objective, i.e. he will get the best outcome with respect to it. As a matter of fact, he may even occasionally attain a social optimum but he can never know whether this is the case and, if not, how big the welfare losses associated to different scoring rules are: what the regulator knows is just that in choosing a certain score function, in fact he sets it as his maximand.

This clearly poses a new problem to the regulator—to make a reasoned choice of the pseudo-objective. That is, instead of assessing different score functions against a given a priori natural objective, he is now to evaluate the merits of different pseudo-objectives in order to find a reasonable one to implement. Hence, the regulator's question is no longer "by what instruments can I pursue my given objective?" but "which objective am I to pursue?". It is to be noted that the latter is extraneous to standard mechanism theory and the design problem here essentially reduces to making value judgements about alternative objectives. To gain a better grasp of the problem let us see an example.

Example. Service contracts sometimes set a rigid capacity constraint on supply. Consider a contract whose object is the operation of a parking lot: the number of parking hours "produced" per day is equal to the number of parking places times the daily opening hours and not a single hour more can be made available by the given facility (it is irrelevant here if the parking facility already exists or must be built under the contract: a BOT contract which sets a given size for a parking lot to be built presents the same problems as contracts for the operation of an existing one). The maximum capacity $\bar{y}$ of course coincides with the fixed service supply and all costs are here fixed ($c(\bar{y}, m)$ is an uncertain constant). In these circumstances the first-best tariff is $P(\bar{y})$ if $y(0) \geqslant \bar{y}$, where $y \mapsto P(y)$ is the inverse of the demand function, and 0 if $y(0) < \bar{y}$. Then, if market demand is completely unknown, the first-best tariff is not computable by the regulator. The production capacity, however, can be used to define scores in many ways, like e.g.

$$K(p, s) = p\bar{y} + s \tag{3}$$

which represents the maximum cost—subsidy plus tariff times maximum capacity—that citizens commit themselves to bear for a maximum capacity equal to $\bar{y}$ when the contract awarded through an auction is (p, s). Since Eq. (3) is a cost, score function/pseudo-objective $K(\,\cdot\,)$ is obviously to be minimized, not maximized, and hence the regulator's objective will be to *minimize the maximum cost of service*. Minimizing pseudo-objective $K(\,\cdot\,)$ is the same as maximizing $-K(\,\cdot\,)$ and the latter is monotonic decreasing in both variables, continuous and additively separable. Then, if $m_i = m$ for all i, Lemma 2 applies to this function and, in order to minimize the expected value of maximum service cost, the regulator has just to run a second-score auction or equivalently a first-score auction, if fixed costs are i.i.d. too (Proposition 3), in both of which the winner is the bidder whose bid solves $\min_i(p_i\bar{y} + s_i)$.

Can such a score function be judged reasonable? In the absence of any information on service use, it is not implausible that citizens may indeed consider reasonable to minimize the maximum cost of service supply but even this simple objective admits several specifications and Eq. (3) is just one. For instance, if the whole subsidy is paid upfront, one can make a more precise estimate of the maximum cost of service for society by replacing the right-hand member of Eq. (3) with $[p\bar{y} + (1 + r)s]$, where r is an appropriate interest rate that measures the cost of capital advance (this specification corresponds to payment of the subsidy in this period and payment of the tariff in the next one; of course any other timing structure can be allowed for by modifying it suitably). A further alternative is to allow for the shadow cost of public funds, λ (subsidies are usually financed through distortionary taxes) by transforming the right-hand member of Eq. (3) into $[p\bar{y} + (1+\lambda)s]$, and others can be found in similar ways. What we have here is the situation that will typically be faced by the regulator: several alternative pseudo-objectives—all of them plausible—among which a choice is to be made. This problem is too complex to be addressed here but one point is clear enough. If the conditions identified by the previous analysis occur, choosing a score function amounts to choosing the pseudo-objective pursued and both regulator and citizens fully know what society will attain even under severe ignorance. Then an obvious minimal requirement in these circumstances is that the decision-maker explain to citizens why maximizing the chosen pseudo-objective is more desirable than maximizing others. Trivial though it may appear, the requirement has a relevant implication for practice: in order to be communicated, the choice must in fact be restricted to objectives that have a plain economic meaning—like minimizing the maximum cost of service—easily understandable to citizens. The example also shows a further fact that is worth stressing: sometimes satisfactory auction design can be accomplished just by good sense without resorting to complicated calculations, contrary to the impression auction theory often conveys to laymen.

7 Conclusions

The paper provides a few basic insights into the parsimonious design of multidimensional auctions. We have started from the extreme situation where the principal knows nothing at all about its potential counterparts' costs and tried to apply the most demanding choice criterion of all—prior-independent optimality. We have shown is that complete ignorance is incompatible with it but also that, on the other hand, prior-independent optimization does admit constrained solutions when a minimum amount of a priori information about cost structure is available to the regulator. In particular we have characterized the structural conditions and the least amount of information for this to be possible—in short that cost variations among firms be imputable to differences in fixed costs.

Our analysis traces the boundaries beyond which the criterion of prior-independent optimality becomes useless and weaker choice criteria or learning mechanisms are needed. One direction for future research is to assay if the criterion can be slackened in some way to be applicable to more general environments. Approximate optimality is a possible answer: if differences in variable cost functions across firms are sufficiently small, we expect that the social welfare function, though not exactly prior-independent optimal, might be approximately so and an interesting task is to determine which conditions, structural and informational, guarantee approximate implementation by an ignorant regulator. A further question raised by our analysis is the choice of pseudo-objectives. As we have seen, when ignorance concerns market demand too, the regulator is in some cases called to make a choice of the pseudo-objective he pursues: such a choice can hardly be rational but ought to be reasonable in some sense and appropriate criteria are needed to define reasonableness in this context. A last problem to which we draw attention is the limitation of the present analysis to scoring auctions. These are some of the most frequently used formats for multidimensional auctions but certainly not the only ones. A possible cause of their diffusion is that they are simple to understand and use in practice but the fact is that theory has very little to say in this regard. What are the exact reasons that make one prefer these mechanisms over others? And more basically: are we sure that there are no other mechanisms as simple as scoring auctions but better than them in some circumstances? Of course, such questions can find no answer here since the scoring-auction format has been assumed as an a priori constraint of the analysis and more general models will be needed to make advances in this direction.

Appendix

Lemma 4 *In second-score auctions if in some state $(m_i, F_i)_{i \in I}$ the second-ranked bid under score function V is not an optimal choice for the winner and the implemented contract yields strictly positive profits, V is not a pointwise-optimal score function.*

Proof. Let a denote the contract implemented in state $(m_i, F_i)_{i \in I}$ by a second-score auction under score function V, i.e.

$$a \in \mathrm{argmax}_{V(p,s) \geqslant V^*} \pi_{(1)}(p, s),$$

where subscript (i) denotes the i-th ranked bidder and $V^* \equiv V(b_{(2)})$, with the second-ranked bid $b_{(2)}$ defined as

$$b_{(2)} \in \mathrm{argmax}_{\pi_{(2)}=0} V(p, s).$$

Note that to avoid cumbersome notation we keep implicit the state of the world on which a and $b_{(2)}$ depend. Notation can be further simplified by assuming without loss of generality that the first- and second-ranked bidders are respectively 1 and 2, whereby index i can be used in place of index (i).

Recall that by assumption $\pi_1(a) > 0$. Moreover, by assumption we also have $b_2 \notin \mathrm{argmax}_{V(p,s) \geqslant V^*} \pi_1(p, s)$, i.e. $a \neq b_2$, which implies $p_a \neq p_{b_2}$, for V is strictly decreasing in s (the equilibrium condition $V(p_a, s_a) = V(p_{b_2}, s_{b_2})$ can be satisfied for $s_a \neq s_{b_2}$ only if $p_a \neq p_{b_2}$). In these circumstances there always exists a point $c \in [0, \infty[\times \mathbb{R}$ such that

$$\begin{aligned} &\pi_1(c) \geqslant 0 \\ &\hat{W}(c) > \hat{W}(b_2) \end{aligned}$$

and $\pi_1(c) > \pi_1(b_2)$: since by assumption $\pi_1(a) > \pi_1(b_2)$, c can be chosen so close to a that the second bid lies below bidder 1's isoprofit curve of level $\pi_1(c)$ (for example we can set $c = a + (0, -\delta) = (p_a, s_a) + (0, -\delta)$, $\delta > 0$ and, given $\pi_1(p_a, s_a - \delta) = \pi_1(a) - \delta$, δ can be chosen so that $\pi_1(a) > \delta \geqslant \pi_1(a)/2$). Strict monotonicity of $\hat{W}$ in s ensures $\hat{W}(c) > \hat{W}(b_2)$.

Let $\psi_i(p) \equiv py(p) - c(y(p), m_i) - F_i$ denote the operational profits (commercial revenue minus cost) and consider the following functions of p

$$\begin{aligned} &p \mapsto -\psi_1(p) + d_1 \\ &p \mapsto -\psi_2(p) \end{aligned}$$

where d_1 is chosen so that $-\psi_1(p_a) + d_1 = \pi_1(a) - \delta$. Note that, given the continuity of the two functions, the minimum function $p \mapsto \min\{-\psi_1(p) + d_1, -\psi_2(p)\} \equiv h(p)$ is continuous too. Then define $\mu : [0, \infty[\to \mathbb{R}$ such that

$$\begin{aligned} &\mu(p_c) = 0 \\ &\mu(p_{b_2}) = 0 \\ &\mu(p) < 0, p \neq p_{b_2}, p_c \end{aligned}$$

(functions with such characteristics, even belonging to class $C^\infty([0, \infty[)$, can be defined without difficulty) and a function $\tilde{V}$ such that

$$\tilde{V}(p, s) \equiv h(p) + \mu(p) - s.$$

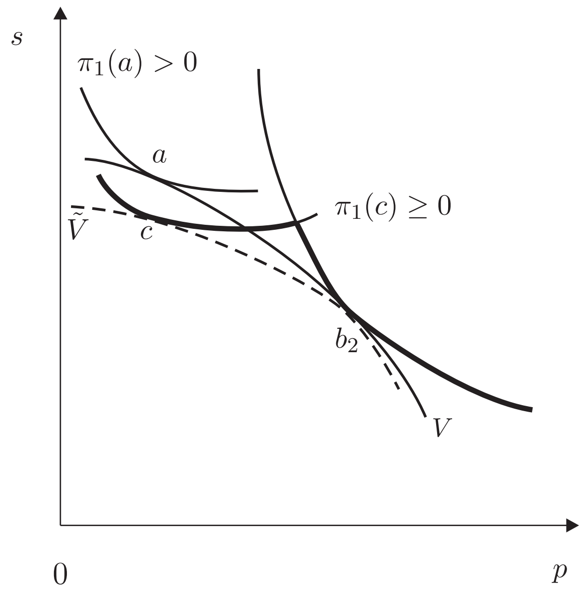

Figure 4: Construction of a score function $\tilde{V}$ that dominates a generic V in some state.

It is immediate to see that $\tilde{V}$ is continuous and has value 0 if and only if $h(p) + \mu(p) = s$. Moreover, $\tilde{V}(p, h(p)) \equiv \mu(p) \leqslant 0$ (that is, it is always < 0 except at p_c and p_{b_2}). Therefore points c and b_2 are maximizers of $\tilde{V}$ for $s \geqslant h(p)$, since

$$\tilde{V}(p, s) \equiv h(p) + \mu(p) - s \leqslant h(p) + \mu(p) - h(p) = \mu(p) \leqslant 0,$$

and this proves the existence of an iso-score curve like $\tilde{V}$ in Fig. 4. By translating the curve vertically we get a score function $\tilde{V} : [0, \infty[\times \mathbb{R} \to \mathbb{R}$ that dominates V in the state of the world under consideration and therefore we can conclude that in the given conditions V is not pointwise optimal.

$\square$

Proof of Proposition 1.
(Sufficiency) The sufficiency part is an immediate implication of Lemma 2.
(Necessity) Given any m_1 and m_2 such that $m_1 \neq m_2$, there always exists a state of the world where bidder 1 is the winner and bidder 2 is ranked second: to find it simply set 2's fixed cost high enough for him to be second and all others' costs higher. Then if parameters m_i are allowed to differ across firms, there is always a state of the world where the highest and the second-highest bidders have different variable-cost parameters.

It is immediate to verify that second-score auctions with a given score function V cannot implement a contract a with $\pi_1(a) = 0$ in every state of the world in which $m_1 \neq m_2$. In second-score auctions this case occurs only if both the first and the second bid are equilibrium ones in some state (i.e. there is a "tie"), i.e.

$$V(p_1^*, s_1^*) = V(p_2^*, s_2^*) \tag{4}$$

with

$$(p_i^*, s_i^*) \equiv \text{argmax}_{\pi(p,s,m_i,F_i)=0} V(p,s), i = 1,2.$$

If in some state Eq. (4) holds, there exists an equilibrium where 1 wins and implements contract (p_1^*, s_1^*): since $V(p_1^*, s_1^*) = V(p_2^*, s_2^*)$ is the highest score subject to the zero-profit constraint, it is impossible to attain higher profits under the constraint $V(p_1, s_1) = V(p_2^*, s_2^*)$. Similarly, there exists a second equilibrium where 2 wins and implements (p_2^*, s_2^*). Of course for fixed costs $F_1 - \varepsilon \geqslant 0, \varepsilon > 0$ we have $\pi(p_1^*, s_1^* - \varepsilon, m_1, F_1 - \varepsilon) = 0$ and $V(p_1^*, s_1^* - \varepsilon) > V(p_2^*, s_2^*)$, i.e. in state $((m_1, F_1 - \varepsilon), (m_2, F_2), \ldots)$ there is no tie. We have thus proved that for any V there are states of the world with $m_1 \neq m_2$ in which contract a implemented through a second-score auction does not meet $\pi_1(a) = 0$, i.e. there are states where the winner realizes strictly positive profits, $\pi_1(a) > 0$. The same argument holds when the winner is any i and the second is any j (i.e. there are states in which any couple of bidders is ranked in this way). Since the single-crossing property implies that in no state of the world the second bid is an optimal choice for the winner ($b \notin \text{argmax}_{V(p,s) \geqslant V^*} \pi_1(p,s)$), Lemma 4 applies to every state $(m_i, F_i)_{i \in I}$ such that $m_i \neq m_j$ for some i, j. Then we can conclude that no V is pointwise optimal if m_i are not equal across firms in all states of the world. □

Proof of Lemma 3. The proof is based on the differential method of equilibrium analysis introduced by Milgrom (1989, 1985) (here all functions are assumed differentiable). Since m is common knowledge to firms, each firm i is uncertain about $(F_1, \ldots, F_{i-1}, F_{i+1}, \ldots, F_n)$. Variable m is a common parameter of firms' choices and therefore we can represent the problem as if we had as many auctions as the values of m. Given the optimal winning probability $\phi^*(m, F_i)$, the optimal expected profit for a generic firm of type F_i is

$$\pi^*(m, F_i) = \phi^*(m, F_i) \left[p^*(m, F_i) y(p^*(m, F_i)) + s^*(m, F_i) - c(y(p^*(m, F_i)), m) - F_i\right]$$

By the Envelope Theorem we have

$$\pi_F^{*\prime}(m, F_i) = -\phi^*(m, F_i)$$

from which by integration there follows

$$\pi^*(m, F_i) = \int_{F_i}^{\bar{F}_i} \phi^*(m, z) dz$$

($\pi^*(m, \bar{F}_i) = 0$: the probability of winning for the worst type $\bar{F}_i$ is null and so is the optimal expected profit in $\bar{F}_i$). Therefore a generic bidder's expected profit conditional on m is

$$E_F(\pi^*(m, F_i)) = \int_{\underline{F}_i}^{\bar{F}_i} \int_{F_i}^{\bar{F}_i} \phi^*(m, z) g(F_i) dF_i dz$$

The winner's expected profit is therefore $n E_F(\pi^*(m, F_i))$. In all auctions awarded to the lowest F_i (given m) the winning probability $\phi^*(m, F_i) = [1 - G(F_i)]^{n-1}$ is the same and so is the expected profit conditional on m. The conventional surplus, defined as

$$\left[v(p) + py(p) - c(y, m) - F_i\right],$$

is in expected value the same in all auctions that select the most efficient firm and entail the same tariff outcome (that is if $p^*(m, F_i)$ is invariant to the auction rules), i.e.

$$E_{F_{(1)}}\Big\{v(p^*(m, F_{(1)})) - s^*(m, F_{(1)}) \\ + \big[p^*(m, F_{(1)})y(p^*(m, F_{(1)})) + s^*(m, F_{(1)}) - c(y(p^*(m, F_{(1)})), m) - F_{(1)}\big]\Big\}$$

where $F_{(1)}$ is the first-order statistics of F. The sum in square brackets is the winner's expected profit, which is equal across all auctions under consideration. Then, by difference, both the expected score conditional on m

$$E_{F_{(1)}}\left[v(p^*(m, F_{(1)})) - s^*(m, F_{(1)})\right]$$

and its expected value with respect to m are equal too. □

Equivalence of welfare functions derived from social surplus

First we verify that net consumer surplus is equivalent to Eq. (1) as score function, in the sense that, even if the regulator knew each bidder's parameters and were able to compute the social welfare Eq. (1) for each bid, the outcome of second-score auctions awarded in this way would be the same as if he used net consumer surplus as score function. So for the argument's sake suppose that the regulator knows the state of the world $(m_i, F_i)_{i\in I}$ at the moment of evaluating bids. To prove equivalence we must ascertain that in switching from a score function to the other there is no change in: a) the winner, b) the implemented contract.

It is immediate to verify that the winner does not change with the score function. Denote operational profits (i.e. net of subsidy) $\psi(p, m, F) \equiv py(p) - c(y(p), m) - F$ and let (p_i^*, s_i^*) be a solution to

$$\max_{p,s}\left[CS(p) - s + \alpha\psi(p, m_i, F_i) + \alpha s\right] \tag{5}$$

subject to

$$s + \psi(p, m_i, F_i) = 0.$$

Since by Eq. (5)

$$CS(p_i^*) - s_i^* + \alpha\psi(p_i^*, m_i, F_i) + \alpha s_i^* \geqslant CS(p_j^*) - s_j^* + \alpha\psi(p_j^*, m_j, F_j) + \alpha s_j^*$$

implies

$$CS(p_i^*) - s_i^* \geqslant CS(p_j^*) - s_j^*$$

then the ordering of bidders is the same both under Eqs. (1) and (2) as score functions.

Let us now turn to the second problem. Assume that under both score function the winner is 1 (identified by (m_1, F_1)) and the second-highest is 2 (m_2, F_2). The second-highest bid is a solution to

$$\max_{p,s}\left[CS(p) - s + \alpha\psi(p, m_2, F_2) + \alpha s\right] \tag{6}$$

subject to

$$s + \psi(p, m_2, F_2) = 0$$

Firm 2's tariff is clearly also a solution to

$$\max_p \left[CS(p) + \psi(p, m_2, F_2)\right]$$

i.e. it maximizes social surplus irrespective of α. Denote W_2^* the social welfare associated to the second-highest bid. The tariff implemented in equilibrium is a solution of

$$\max_{p,s} \pi(p, s, m_1, F_1) \tag{7}$$

subject to $CS(p) - s + \alpha(\psi(p, m_1, F_1) + s) = W_2^*$, i.e. it is a solution of

$$\max_p \left[\psi(p, m_1, F_1) + \frac{1}{1-\alpha}\left(CS(p) + \alpha\psi(p, m_1, F_1) - W_2^*\right)\right]$$

which we can also write as

$$\max_p \left\{\frac{1}{1-\alpha}\left[CS(p) + \psi(p, m_1, F_1)\right] - \frac{W_2^*}{1-\alpha}\right\}.$$

Therefore the outcome tariff $p^*(m_1)$ does not vary with parameter α, nor does the outcome subsidy $s^*(m_1, F_1)$. In other words, it is the same that would obtain by applying the consumer surplus as score function, i.e. by solving Eqs. (6) and (7) for $\alpha = 0$.

The last task is to check that when we use $(p, s) \mapsto CS(p) - s$ to score bids there exists no other $(p, s) \mapsto V(p, s)$ that allows to attain a better outcome with respect to the social welfare $CS(p) - s + \alpha\pi(p, s, m, F)$ in some state. If the regulator knew the state of the world, he could use $W(p, s, m, F) \equiv CS(p) - s + \alpha\pi(p, s, m, F)$ as score function. Since both welfare and profit functions are vertical translations when $m_i = m$ for all i, the implemented contract and the second bid coincide and the argument of Lemma 2 can be applied state by state to this function too. This means that there exists no score function V under which the auction outcome obtains a strictly higher welfare $CS(p) - s + \alpha\pi(p, s, m, F)$. But, since net consumer surplus $CS(p) - s$ allows to attain the same outcomes as W in every state, there exists no V that strictly dominate the former in terms of welfare $CS(p) - s + \alpha\pi(p, s, m, F)$. In conclusion, by resorting to Eq. (2) there is no loss due to the informational constraint that makes Eq. (1) actually unusable for scoring bids.

References

Armstrong M. (1996). Multiproduct nonlinear pricing. *Econometrica*, **64**(1):51–75. doi:10.2307/2171924.

Armstrong M. and Rochet J.C. (1999). Multi-dimensional screening: A user's guide. *European Economic Review*, **43**(4-6):959–979. doi:10.1016/S0014-2921(98)00108-1.

Asker J. and Cantillon E. (2005). Optimal procurement when both price and quality matter. Centre for Economic Policy Research, London, Discussion Paper no. 5276.

Asker J. and Cantillon E. (2008). Properties of scoring auctions. *Rand Journal Of Economics*, **39**(1):69–85. doi:10.1111/j.0741-6261.2008.00004.x.

Baliga S. and Vohra R. (2003). Market research and market design. *Advances in Theoretical Economics*, **3**(1):1059–1059. doi:10.2202/1534-5963.1059.

Baron D.P. and Myerson R.B. (1982). Regulating a monopolist with unknown costs. *Econometrica*, **50**(4):911–930. doi:10.2307/1912769.

Beil D.R. and Wein L.M. (2003). An inverse-optimization-based auction mechanism to support a multiattribute rfq process. *Management Science*, **49**(11):1529–1545. doi:10.1287/mnsc.49.11.1529.20588.

Bergemann D. and Morris S. (2005). Robust mechanism design. *Econometrica*, **73**(6):1771–1813. doi:10.1111/j.1468-0262.2005.00638.x.

Bichler M. and Kalagnanam J. (2005). Configurable offers and winner determination in multi-attribute auctions. *European Journal of Operational Research*, **160**(2):380–394. doi:10.1016/j.ejor.2003.07.014.

Branco F. (1997). The design of multidimensional auctions. *Rand Journal of Economics*, **28**(1):63–81. doi:10.2307/2555940.

Bushnell J.B. and Oren S.S. (1994). Bidder cost revelation in electric-power auctions. *Journal of Regulatory Economics*, **6**(1):5–26. doi:10.1007/BF01065387.

Bushnell J.B. and Oren S.S. (1995). Internal auctions for the efficient sourcing of intermediate products. *Journal of Operations Management*, **12**(3-4):311–320. doi:10.1016/0272-6963(94)00013-5.

Che Y.K. (1993). Design competition through multidimensional auctions. *Rand Journal of Economics*, **24**(4):668–680. doi:10.2307/2555752.

Crémer J. and McLean R.P. (1988). Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica*, **56**(6):1247–1257. doi:10.2307/1913096.

David E., Azoulay-Schwartz R., and Kraus S. (2006). Bidding in sealed-bid and english multi-attribute auctions. *Decision Support Systems*, **42**(2):527–556.

Goldberg A. and Hartline J. (2003). Competitiveness via consensus. In *Proceedings of the 14^{th} Symposium on Discrete Algebra*. ACM/SIAM.

Goldberg A.V., Hartline J.D., Karlin A.R., Saks M., and Wright A. (2006). Competitive auctions. *Games and Economic Behavior*, **55**(2):242–269. doi:10.1016/j.geb.2006.02.003.

Hartline J. and McGrew R. (2005). From optimal limited to unlimited supply auctions. In *Proceedings of the 6^{th} ACM conference on Electronic commerce*, pages 175–182. ACM Press New York, NY, USA.

Jehiel P., Moldovanu B., and Stacchetti E. (1999). Multidimensional mechanism design for auctions with externalities. *Journal of Economic Theory*, **85**(2):258–293. doi:10.1006/jeth.1998.2501.

Laffont J.J., Maskin E., and Rochet J.C. (1987). Optimal nonlinear pricing with two-dimensional characteristics. In *Information, Incentives, and Economic Mechanisms*, edited by Groves T., Radner R., and Reiter S., pages 256–266. University of Minnesota Press, Minneapolis.

McAfee R.P. and McMillan J. (1988). Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory*, **46**(2):335–354. doi:10.1016/0022-0531(88)90135-4.

McAfee R.P. and Reny P.J. (1992). Correlated information and mechanism design. *Econometrica*, **60**(2):395–421. doi:10.2307/2951601.

Milgrom P. (1989). Auctions and bidding - a primer. *Journal of Economic Perspectives*, **3**(3):3–22.

Milgrom P.R. (1985). The economics of competitive bidding: A selective survey. In *Social Goals and Social Organization: Essays in Memory of Elisha Pazner*, edited by Hurwicz L., Schmeidler D., and Sonnenschein H., pages 261–289. Cambridge University Press, Cambridge, UK.

Müller R., Perea A., and Wolf S. (2007). Combinatorial scoring auctions. Meteor, University of Maastricht.

Naegelen F. (2002). Implementing optimal procurement auctions with exogenous quality. *Review of Economic Design*, **7**(2):135–153. doi:10.1007/s100580200067.

Riordan M.H. and Sappington D.E.M. (1987). Awarding monopoly franchises. *American Economic Review*, **77**(3):375–387.

Rochet J.C. and Choné P. (1998). Ironing, sweeping, and multidimensional screening. *Econometrica*, **66**(4):783–826. doi:10.2307/2999574.

Rochet J.C. and Stole L.A. (2002). Nonlinear pricing with random participation. *Review of Economic Studies*, **69**(1):277–311. doi:10.1111/1467-937X.00206.

Segal I. (2003). Optimal pricing mechanisms with unknown demand. *American Economic Review*, **93**(3):509–529. doi:10.1257/000282803322156963.