



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Predicting the performance of conservation tenders when information on bidders' costs is limited

Authors: [Steven Schilizzi](#) *
& Uwe Latacz-Lohmann

(*) Corresponding author: Steven Schilizzi, SARE, UWA, 35 Stirling Hwy, Crawley WA 6009
Email: Steven.Schilizzi@uwa.edu.au

ABSTRACT

Buying environmental services from private landholders using tendering mechanisms are usually subject to a budget constraint. Auction theory has mostly focused on target-constrained auctions and is not well developed for this type of auction. This paper examines the predictive capacity of a simple model developed for budget-constrained tenders, already used to design new conservation programs, by submitting it to controlled lab experiments. We study the capacity of the model to predict both experimental bids and the performance of the auction institution, based on the kind of limited information typically available to a conservation agency. We conclude there exists an optimal level of information on bidders' costs, neither too large nor too small, making the tender worth considering as a policy option as well as allowing an ex-ante assessment of its economic performance.

Key words: Auctions, procurement, tenders, conservation, learning, economic experiments

JEL Classification: C91, C92, D44, Q24, Q28

I. INTRODUCTION

Buying environmental services from private landholders using tendering mechanisms usually involves budget-constrained, procurement-type auctions. This poses a problem to the extent that auction theory has been well developed, since Vickrey's 1962 paper (less well-known than his much-cited 1961 paper), for target-constrained (TC) auctions, but much less so for budget-constrained (BC) auctions (Müller and Weikard, 2002). In a target-constrained auction or tender¹, the number of contracts or hectares of land to come under contract is decided upon and is known; the risk is with what it might end up costing. In a budget-constrained auction or tender, the programme's budget is decided upon and is known; the risk is with the number of contracts or hectares that might not come under contract, that is, with the degree of environmental effectiveness of the policy. It seems that target-constrained tenders are used where government cannot fall short of its objectives, as is typically the case with military procurement programmes. In the field of environmental policy, governments' use of the budget-constrained tenders probably reflects their general political priorities. As a result, in the field of environmental policy, there is a gap between theory and practice for BC tendering mechanisms. A better theory would allow agencies to improve tender design and perhaps decide whether such a mechanism is worth going ahead with or not, given existing alternatives.

This study sets out to investigate, using the techniques of experimental economics, the predictive capacity of a new model developed for BC tenders. This model was first proposed by Latacz-Lohmann and Van der Hamsvoort (henceforth, LH) in 1997 and further refined in 1998, where policy implementation was investigated. To the best of our knowledge, this is to date the only extension of auction theory which captures the particular features of conservation tenders. The model has been criticised, however, for not conforming to the standard assumptions of auction theory regarding optimal bid formulation. As explained by Müller and Weikard (2002), making the same assumptions in the BC auction model as in the better known TC model creates one more degree of freedom (since the final number of winners is not known), leading to a complex situation with multiple Nash equilibria and no dominant solution for choosing an optimal bid. LH (1997) solve this problem by introducing an exogenous parameter, the bidders' expectation of the highest acceptable bid, knowing the budget constraint and the number of bidders. Bidders then use this best guess of theirs to form their optimal bids. The result is a very simple model, much simpler than the more standard TC model. This simplicity comes however at a cost, in that no theory or model for the formation of bidders' expectations is offered.

The purpose of this study is to investigate the validity and credibility of this new BC auction model. The focus here is thus on the performance of the BC *model* as opposed to the performance of the BC auction *institution*. The latter issue was investigated in Schilizzi and

Latacz-Lohmann (2007), where, in addition, they compared the performance of the BC and TC auction institutions relative to an equivalent fixed-price scheme. To the extent that government agencies have, to date, almost exclusively used the BC format in environmental policy, it seems important to test any model that might be used to inform the design of this type of policy instrument. In Australia, for example, the Victoria BushTender conservation program was directly inspired by the LH 1997 model. (Stoneham et al. 2003).

To carry out this ‘proof of concept’ study, we implemented the institutional and informational conditions of the BC model using controlled economic experiments. In addition, since agri-environmental contracts are often issued over a sequence of years, as the United States Conservation Reserve Program typifies, we extend the problem to repeated auctions. With repetition, as shown in Hailu and Schilizzi (2004), bidders learn to bid so as to extract increasing information rents at the expense of budgetary cost-effectiveness, eventually defeating the purpose of the tender, as bidders learn to bid the government’s implicit reserve price.

By testing the validity and credibility of the BC model we mean two things. First, we test whether optimal bids as computed by the model predict sufficiently well the bids as expressed by experimental subjects. Of course, this test holds only to the extent that the experiments correctly implement the model’s assumptions. Secondly, we test whether the model is capable of predicting the performance of the auction institution using *ex-ante* predicted bids rather than *ex-post* actual bids. In other words, using both model predictions and experimental results, would the model recommend the use of this institution when in fact it should not do so, and vice-versa? Importantly, how does the model’s as well as the institution’s performance depend on the cost-related bidder information available to the procurer? In general, the procurer only has a rough knowledge of bidders’ average participation costs. A more accurate knowledge would of course make the use of a tendering system redundant, since an auction also functions as a price or cost revelation mechanism.

The remainder of the paper is organised as follows. Section two summarises the role of controlled laboratory experiments in relation to existing theory for allocating conservation contracts. Section three outlines the BC auction model and highlights the key differences with the more standard TC model. Section four describes the economic experiments, and section five provides and discusses the results. Section six concludes as to whether the LH (1997) model for conservation tenders is a credible tool or not for auction design and environmental policy.

II. THEORY AND EXPERIMENTS FOR CONSERVATION TENDERS

The use of laboratory experiments to study tendering outcomes originated in the fundamental complexity of the auction institution. Following Vickrey's seminal work in 1961, it was soon recognised that a large number of parameters influenced auction performance and that outcomes were very sensitive to the values of these parameters. These included for example the distribution of information (private-value versus common-value auctions), auction format (sealed-bid versus open call), and payment format (first-price versus second-price), to name but a few (Klemperer, 1999, 2002, 2004; Milgrom, 1989). Theoretical investigations, which are constrained by analytical tractability, could only investigate the effect of one or a small number of parameters at a time, assuming all others constant. Major reviews of this literature include Cassady's book (1967) and survey papers by Engelbrecht-Wiggans (1980), McAfee and McMillan (1987), Milgrom (1985, 1989), Wilson (1992), and Klemperer (1999). As a result, the theoretical literature on auctions remained divorced from the practical needs of auction implementation. This is well reviewed by Rothkopf and Harstad (1994) and Klemperer (2002).

Economic experiments were called upon to bridge the gap between theory and practical implementation. Kagel's review, in Kagel and Roth's (1995) *Handbook of Experimental Economics*, remains a key reference for the contributions of the experimental effort up to that date. The *Handbook of Experimental Economics Results* (Plott and Smith, 2008), will, when all the volumes have been published, provide a very comprehensive update.

The situation is exacerbated in the case of conservation tenders, as these are usually procurement (reverse), repeated, multi-unit auctions. They are procurement auctions in that the auctioneer (the government agency) buys rather than sells environmental services. They are multi-unit auctions in that landholders sell units of different quality (environmental services per unit area vary across the landscape), they can sell several units each, and there is more than one winner. Conservation tenders are also repeated over time, as witnessed by the US Conservation Reserve Program (CRP) which has been run as a multiple sign-up scheme (Riechelderfer and Boggess, 1988; Johansson, 2006). Auction theory is less well developed for procurement than for direct (selling) auctions, for multiple-unit than for single-unit auctions, and for repeated than for one-shot auctions. The main reason, on which we shall not dwell here, is the level of complexity involved by the characteristics of conservation tenders.

Accordingly, conservation tenders have begun to be studied experimentally. This refers, strictly speaking, to controlled laboratory experiments, but can also be understood in a broader sense to mean the sequential combination of laboratory experiments and small-scale field trials. This was done in Australia in connection with the BushTender trials in the state of Victoria, where certain design problems, in particular the amount and choice of the information to be

communicated to landholders before the bidding session, was investigated experimentally (Cason et al., 2003). In the State of Georgia, USA, auctions for buying back water abstraction licences from irrigators in times of drought were not implemented before a number of controlled laboratory experiments had been carried out (Laury, 2002; Cummings *et al.*, 2004). Cason and Gangadharan (2005) report the results of an economic experiment to investigate the outcome properties of uniform versus discriminatory-price auctions for reducing non-point source pollution. They find that although overbidding was more pronounced in the discriminatory-price auction, the discriminatory format had superior overall market performance.

The present paper contributes to the experimental effort in the field of conservation tenders. In contrast to previous studies, which have investigated the outcome properties of alternative auction design options, the focus of this paper is on testing new theory – the first bidding model that attempts to capture the particular features of conservation tenders.

III. THE BUDGET-CONSTRAINED BIDDING MODEL

Key features of the budget-constrained model

The sealed-bid discriminatory price budget-constrained (BC) model examined in this paper was first proposed by Latacz-Lohmann and Van der Hamsvoort (LH) in 1997. They considered landholders to hold private information about their opportunity costs of participating in the government's conservation programme. These costs arise when management prescriptions divert farmers' land management practices away from their current plan, assumed to be the most profitable one. The government's problem, in order to attract farmers into the scheme, is to compensate them for the lost profits without knowing the magnitude of their opportunity costs. Auctions have the property of revealing at least part of this information. In order for the landholder to participate in the scheme, the payment he or she receives must be at least equal to his or her opportunity cost of participation.

LH (1997) first assume that landholders' bidding strategies are predicated on the belief that the conservation agency (the procurer) will decide on a maximum acceptable bid, or payment level, β . This is a common practice when the agency is subjected to a constrained budget. In actual fact, this maximum bid β is determined *ex post*, after all bids have been received, as the last (highest) bid accepted within the available budget. In other words, no individual bids above β will be accepted. β represents an implicit reserve price per unit of environmental service, unknown to bidders (and also unknown to the procurer until all bids have been received). This external parameter β represents a deviation from standard TC auction theory, where optimal bids are determined endogenously as a function of the number of bidders, the distribution of bidders'

opportunity costs (assumed common knowledge), and the target to be achieved. In the BC auction, this target – the number of winners or hectares contracted – is unknown. A landholder will tender a bid b if the expected utility in case of participation exceeds his or her reservation utility.

The second assumption in LH's model is that bidders, not knowing the value of the bid cap β , will form expectations about it, which can be characterised by the density function $f(b)$ and by the distribution function $F(b)$. The probability that a bid is accepted can then be expressed as

$$p(b \leq \beta) = \int_b^{\bar{\beta}} f(b)db = 1 - F(b) \quad (1)$$

where p is probability and $\bar{\beta}$ represents the upper limit of the bidder's expectations about the bid cap, or the maximum expected bid cap. The essence of the bidding problem is to balance out net payoffs and probability of acceptance. This means determining the optimal bid which maximises the expected utility over and above the reservation utility.

Further assumptions are that there are no transaction costs in bid preparation and implementation, that payment is only a function of the bid (discriminatory price auction), and that bidders are risk-neutral². A risk-neutral bidder simply maximises expected payoff. The optimal bid, b^* , derived by LH (1997) is given by equation (2), where c represents the opportunity costs of participation:

$$b^* = c + \frac{1 - F(b)}{f(b)} \quad (2)$$

LH (1997) further assume that bidders' individual expectations about the bid cap β are uniformly distributed in the range $[\underline{\beta}, \bar{\beta}]$, where the lower and upper bounds represent the bidder's minimum and maximum expected bid cap. A bidder's expectations are that any bid equal to or below $\underline{\beta}$ has a probability of 1 of being accepted, and any bid equal to or above $\bar{\beta}$ has a probability of zero of getting accepted. Then the expression for the optimal bid becomes (LH, 1997):

$$b^* = \max \left[\frac{1}{2}(c + \bar{\beta}), \underline{\beta} \right] \quad \text{s.t.} \quad b^* > c \quad (3')$$

This is true for each of the i bidders, so that expression (3') also reads as:

$$b_i^* = \max \left[\frac{1}{2}(c_i + \bar{\beta}_i), \underline{\beta}_i \right] \quad \text{s.t.} \quad b_i^* > c_i \quad (3'')$$

Expressions (3) show that the optimal bidding strategy of a risk-neutral bidder increases linearly with both the bidder's opportunity costs c_i and his or her expectations about the bid cap, characterized by $\underline{\beta}_i$ and $\bar{\beta}_i$. Bids thus convey information about opportunity costs, which are

private information unknown to the procurer; they thereby reduce the information asymmetry, but not completely: indeed, the auction’s cost revelation property is blurred by the fact that bids also reflect bidders’ beliefs about the bid cap. This creates room for bidders to bid above their true opportunity costs and thereby to secure for themselves an information rent. With repetition, when bidders have the opportunity to learn from the results of past bidding strategies, this blurring is expected to be further exacerbated: from Hailu and Schilizzi’s (2004) results, one would expect bids to depend increasingly on bid cap expectations (the β ’s) and less on opportunity costs (the c_i).

Key differences with the better known target-constrained model

As discussed by Müller and Weikard (2002), target-constrained (TC) auctions differ from budget-constrained (BC) auctions by allowing endogenous expectations to form and optimal bids to be formulated without the need for exogenous bid caps. As a result, the nature of the two models is different. While the BC model is a best-response model, the TC model is a Nash equilibrium model. This is because, unlike a limited budget, knowing the target also tells bidders how many winners there will be or contracts will be allocated, and thereby yields fewer degrees of freedom than the BC auction. This explains why TC auctions were the first to be modelled.

Following his 1961 work, Vickrey (1962) formulated a Nash equilibrium bidding model in the case of single-unit sealed-bid discriminative price auctions (when agents bid only for one unit) and demonstrated that the Revenue Equivalence Theorem holds for risk-neutral bidders with individual values for the auctioned objects drawn from a uniform distribution. Harris and Raviv (1981) generalised the Vickrey model for bidders’ valuations drawn from general distribution functions and when all bidders have identical concave utility functions. All subsequent extensions (Milgrom and Weber, 1982; Cox et al, 1984) focused on “selling” auctions, of little relevance to conservation programmes. In the literature, optimal bid formulas have been explicitly given for direct or selling auctions (e.g. Cox et al. 1984) but not for procurement tenders. Hailu et al. (2005) customised the Harris and Raviv (1981) approach to model the Nash equilibrium risk-neutral bid functions in a procurement multi-unit sealed-bid tender, relevant for government conservation schemes. After showing that, even in the single-unit case, the optimal bid formulas for discriminatory-price ‘selling’ and procurement auctions are not symmetrical³, they derived a general formula for the optimal bid in a multiple-unit procurement auction:

$$B(v) = \frac{\int_v^{v_{\text{sup}}} uG'(u)du}{\int_v^{v_{\text{sup}}} G'(u)du} \quad (4)$$

where n is the number of bidders, m the number of units sought (i.e. the target), v the bidder’s own value, and u the integrand between v and v_{sup} , the value or cost of the highest cost bidder. $G(v)$ is

derived from the expression of the probability distribution of the m^{th} order statistic on a $(n-1)$ sample in that distribution. The formula is made explicit for bidder values (or opportunity costs) uniformly distributed on $[0, 1]$, yielding:

$$b(v) = \frac{\int_v^1 u^m (1-u)^{n-m-1} du}{\int_v^1 u^{m-1} (1-u)^{n-m-1} du} \quad (5)$$

where here u is the integrand for values varying between v and $v_{sup} = 1$.

The assumptions underlying this model are that bidders are risk-neutral, each bidder can sell more than one unit, winners get paid their own bid (discriminative price auction), and the uniform probability distribution of values on support $[0,1]$ is common knowledge to all. A key difference with the BC model appears in the upper bound of the integrals: whereas bidders in a TC auction are modelled as forming their bids in reference to other bidders' values (v_{sup}), bidders in the BC auction form their bids in reference to their expected bid cap (β). In addition, both models make bids depend (equivalently) on bidders' values or costs, but the BC model's bid dependence is only on the bidder's own value or cost, whereas in the TC model the bid depends on an assumed distribution over all values or costs.

In terms of outcomes, both models predict that the optimal bidding strategy is one of overbidding ($b > c$ or $b > v$). Moreover, the level of overbidding is high for low-value bidders and low for high-value bidders. Overbidding decreases as the private value or cost increases, with the bids from high-value or high-cost bidders asymptotically approaching their respective values: $\lim_{v \rightarrow \infty} (b) = v$ or $\lim_{c \rightarrow \infty} (b) = c$.

Thus both BC and TC models predict overbidding as an optimal strategy. This is in accordance with the theory of discriminative (first-price) sealed-bid auctions⁴. However, in a real policy setting, as opposed to controlled laboratory experiments, such over-bidding is likely to be somewhat dampened, due to other-than-profit motivations, such as environmental stewardship, cross-compliance constraints, or fear of future regulatory action, all of which would push their bids down somewhat, for at least some of the bidders. In this setting, however, we focus on the prime profit motive in order to link theory and experiment. This can be considered as a 'worse case scenario' in terms of the institution's economic performance.

IV. EXPERIMENTAL SETUP

The purpose of the experiments described below was to assess the predictive capacity of the non-standard BC model, in order to decide whether it is a credible tool or not for informing budget-constrained auction design for allocating conservation contracts. By 'predictive capacity',

we mean two things: first, the difference between the observed experimental bids and those calculated based on equation 3; and second, whether the performance of the tendering institution using predicted bids is close to that using actual ex-post bids, as expressed by experimental subjects. This should shed some light on whether experimental results can be used for guiding the use of BC tendering mechanisms.

Preliminary bidder surveys

Prior to holding the experiment, we surveyed our experimental subjects along two dimensions: their attitude towards environmental conservation, and towards risk. The first question was asked so as to be able, after the experiment, to relate the amount of bid shading to environmental attitudes, since the auction was set in a land conservation context. One would assume that in a real policy setting, the more environmentally concerned bidders would shade their bids less than the less concerned. Whether such a reduction in bid shading would also be observed in laboratory experiments would depend on the extent to which the context is effective in influencing participants' decisions.

Bidders' risk attitudes were measured using a certainty-equivalent method, whereby they were asked to state the minimum price they would accept from selling a lottery ticket that had been given to them. This measure was also hypothesised to explain possible differences in bid shading, whereby more risk-averse bidders would shade their bids less than the less risk-averse. As it turned out, environmental attitudes, as measured in this survey, did not appear to be related in any way to bid shading, whereas, as will be detailed later, risk attitudes, as measured, did have some impact in the expected direction. The implication of this is that contextual effects such as environmental concerns did not affect experimental outcomes – a positive feature in terms of experimental control.

General experimental setup

Experiments were first carried out at the University of XXX, K, then at the University of YYY in P.¹ The P experiment replicated the K experiment, in order to check for the robustness of results.

The K experiment was carried out with first-year students in agricultural economics. The total number of students was about 44 (the number varied slightly across sessions). The auction setup referred to reductions in nitrogen fertiliser on a wheat crop, in order to meet EU regulations

¹ K and P are used in lieu of actual institution and location names to preserve anonymity in the reviewing process: they will be replaced by the original names in the final version of this paper.

regarding limits to nitrate concentration in groundwater (50 mg/litre). This is a serious concern in rural areas of northern K, and one which students in K would be aware of and sensitive to. Participants were offered would-be contracts for committing themselves to reduce applications of nitrogen fertiliser from their currently *most profitable* level down to a predefined constrained level, equal to 80 kg per hectare. Each participant was given a different production function for nitrogen fertiliser in wheat production and thus faced a different opportunity cost resulting from the adoption of the nitrogen reduction programme. Participation costs were spread uniformly between €5 (the lowest-cost farmer) and €64 (the highest-cost farmer). Bidders knew their own opportunity costs but not those of rival bidders (see appendix). Participants were told that not all of them would be able to win contracts and that they were therefore competing against each other. To keep things very simple, each participant could put up just one land unit of wheat, the same area for all participants. They were told that if they won a contract, they would be paid the difference between their bid and their opportunity cost.

Three identical rounds were held. The purpose of this was to investigate the performance of the model with repetition. That is, was the BC model able to maintain the quality of its predictions as bidders get to “play the game” several times? Three rounds are not many, but if already in the third round a clear tendency was observable, it would be an indication that repetition does affect predictive power of the model. In rounds two and three, exactly the same setup was used, except that bidders knew of their own result in the previous round(s), and successful bidders had been paid their net gains at the end of each round. More rounds were not run due to time and resource limitations.

Specifics of experimental setup

Since auctions are very sensitive to information structure, it was important to control for this aspect. In the first round, bidders were informed of the available budget for the current session. The cost range (€5 to €64) was not given, but bidders were told that costs were uniformly distributed. Each bidder knew his or her own opportunity cost and was given a rough estimate of where he or she stood compared to rival bidders in terms of opportunity costs. This was done by informing bidders in which cost quartile they belonged: lower quarter, second quarter, third quarter, upper quarter (see appendix). It was assumed that bidders could look around and estimate the number of competitors in their group: between 40 and 44 depending on sessions in the K experiment, and 27 in the P experiment.

The budget constraint announced (€3900) was clearly distinguished from the actual payments made at the end of the session⁵. Actual bidder payments would be proportional to their gains calculated as own bid minus participation cost. Bidders were asked two pieces of numerical

information, their estimate of the maximum expected bid cap ($\bar{\beta}_i$), and their bid (b_i). We did not ask for the lower bound $\underline{\beta}_i$ as it could have confused the participants and it was unlikely to be binding in the formulation of their bids. In the following two rounds, bidders also knew whether they had previously been successful or not, and if so, what their net gains were. No information regarding other bidders was given, as e.g. who the winners or what their gains were.

The P replicate

The P experiment was in all points identical to the K experiment, save for the following logistical details. Participants were mostly second-year students, with a few third and fourth years as well as a handful of postgraduates – all in the area of agriculture or natural resource management. They totalled about 27 in number, with a variation of one or two between sessions. To reflect the smaller number of participants in the P experiment, the budget constraint was lowered proportionately, compared to the K experiment (\$2300).

A slight difference in the P experiment was the twist given to the story. Rather than nitrogen leaching into the groundwater, the government agency was buying back from horticulturalists in the P catchment (around P) a composite good made of nitrogen and phosphorus, and the problem was eutrophication in the P river following excess runoff of these two nutrients – a socially and politically sensitive issue in P.

V. RESULTS AND DISCUSSION

Organizing the results

The results are organized as follows. We first examine the BC model's predictive performance by confronting the theoretically optimal bids, computed using equation (3''), with experimentally observed bids. Secondly, we examine whether the model's performance is affected by repetition, by confronting computed and experimental bids in subsequent rounds. We try to understand the model's performance by examining whether bid and *beta* formation obey the rationale assumed by the model. We are then ready to extend the analysis to how well the model predicts the economic performance of the institution. We first do this by comparing its predictions using the experimentally observed bids to that using the computed bids, computing the bids using the individually expressed β 's (third subsection). We then shift the focus from the experimenter's point of view to that of the policy-maker's by constructing a counterfactual in terms of the information available on bidders' costs (fourth subsection). We do this by assuming only average costs used in the experiments, a single-point overall average cost and a four-point quartile cost distribution. Recall that the quartile cost information was also that provided to the bidders themselves.

Using the BC model to examine whether it can predict *ex-ante* the economic performance of a tendering mechanism warrants further explanation. The question is this: can model predictions based on experimental results provide a reliable basis for predicting the performance of a budget-constrained tender? That is, instead of evaluating the tender *ex-post*, after the bids have been received, we would like to know whether an *ex-ante* evaluation, based on predicted bids rather than actual bids, is possible. The rationale for this is provided in Table 1.

Table 1 about here

Table 1: Use of the BC model for predicting the performance of a BC tender

Knowledge by procurer	Costs (and N)	Bid caps β	Computed bids b^*	Auction performance
Pre-experimental knowledge	c_Q (and N_Q) or c_A (and N_A)	---	---	---
Experimental knowledge	c_i (experimental)	$\beta_i = f_i(c_i)$	$b_i^* = f_i(c_i, \beta_i)$	Experimental auction, based on b_i^*
Post-experim., pre-auction knowledge	c_Q or c_A	$\beta_Q = f_i(c_Q)$ or $\beta_A = f_i(c_A)$	$b_Q^* = f_i(c_Q, \beta_Q)$ or $b_A^* = f_i(c_A, \beta_A)$	Based on b_Q^* or b_A^*
Post-auction knowledge	c_Q or c_A	---	---	Based on actual bids from the field

N = amount of Nitrogen abated

The key issue is the information available to the procurer agency. In terms of the BC model, the only information a government agency can reasonably be assumed to have is some rough estimate of the participation costs incurred by program participants. Of course, this information must be sufficiently rough and inaccurate for a cost-revealing mechanism to be worth implementing; otherwise, the conservation contracts could be priced as a function of participation costs. In the analysis that follows, we examine two limited-information scenarios:

- A minimal information scenario, where the procurer only has a single point estimate of the overall average cost of implementing the environmental program: this can be understood as representing the participation costs of a typical, average farm in the policy region.
- A medium information scenario, where the procurer has been able to obtain a four point estimate of program implementation costs, and which represents the quartile distribution. Recall that bidders in the experiments were given this level of information on the distribution of costs, by telling them only which cost quartile they belonged to.

The procuring agency is assumed to have one or the other level of information. We call the single-point cost estimate the ‘average cost scenario’ or ACS and the four-point cost estimate the ‘quartile cost scenario’ or QCS. The costs in the ACS case will yield a single estimate of β , the average

expected bid cap, and a single bid estimate, the overall average bid, which will be used to determine the auction's performance.

In the absence of experimental implementation, the model could not be directly used for policy guidance other than in broad qualitative terms. The extra information provided by its experimental implementation is the dependence of the β 's on individual costs: indeed, the model provides no β -theory on the formation of bid cap expectations, which underlie bid formulation. Transposing this relationship from the experimental setting to the field is of course fraught with danger and subject to the 'external validity' criticism (Schramm, 2005; Guala and Mittone, 2005); the danger is however minimized if the extrapolation, as we shall clarify shortly, is kept under control. The $\beta = f(c)$ relationship, empirically estimated using the experimental data, is then applied to whatever knowledge of costs, 'average' or 'quartile', is available, and the optimal bid formula in equation (3) is then applied to estimate the bids. These are then used for computing the performance criteria of the institution, such as budgetary cost-effectiveness and the rate of bidder information rents.

Underlying this approach are careful assumptions on the distribution of information. In general, in experiments for test-bedding tendering mechanisms, there are four players: the experimenter, the policy maker, the experimental bidders, and the bidders in the field. Deciding who knows what is crucial to the validity of the analysis. In the present case, the distinction between experimental and field bidders is not essential, whereas that between the experimenter and the policy maker is. We simulate the latter's information level by assuming unknown what the former knows; namely, the experimental β_i 's and the experimental bids b_i . The policy maker, instead, can only know the computed average b_A^* or quartiles b_Q^* , depending on the information scenario. The experimenter can then evaluate the performance of the institution by comparing its estimate using the experimental bids and its estimate using the computed bids. The focus is on the direction and order of magnitude of the 'error' made by using the computed bids. It is the analysis of that 'error' which provides the information necessary for policy guidance.

The raw experimental data are provided in Appendix II.

How well does the BC model predict the one-shot tender?

The two frames in Figure 1 plot predicted optimal bids against experimentally observed bids for the BC auction in replicates K and P. Optimal bids were computed for each bidder using equation (3''). The 45 degree line represents perfect prediction, if all data points were situated on it. The closer the linear coefficient is to 1.00, the better the fit, provided the coefficient's t statistic is good enough. Two things can be observed. Firstly, prediction is less than perfect, as one would have expected. Secondly, the model underestimates the experimental bids in K slightly but systematically, the linear fit being everywhere above the 45 degree line.

Figure 1 about here




Figure 1: One-shot model performance: First-round theoretically computed versus experimentally observed bids in K and P for a BC auction. The 45 degree lines of perfect fit are shown. The *** indicate significance at the 1% confidence level.

One feature of the model may explain this slight overbidding: bidders are assumed in equation (3) to be risk-neutral. The bidders in the K experiment were found to be somewhat risk-prone, with an average certainty equivalent ratio of 107%, slightly greater than the 100% reflecting risk neutrality. This may partly explain the underestimation of the model, since risk-prone bidders can be shown to optimally over-bid relative to risk-neutral bidders.

The P data confirm this. P participants were risk-averse, with an average certainty equivalent ratio of 88%, and the model no longer underestimates the bids, but rather, as expected, overestimates them. However, in both K and P experiments, the linear fit has a smaller slope than 1, the 45 degree line, with the difference more marked in P. The BC model slightly predicts low bids higher than it does high bids. This is because the model computes optimal bids which, as mentioned in section 3, reflect higher bid shading for low cost bidders than for high cost bidders. Although these results are statistically significant at the 1% level in both replicates (Figure 1), the K results appear to be more reliable, probably because of the larger number of participants (44 instead of 27 in P). Specifically, as reflected by a higher R^2 (0.83 instead of 0.61), the K bids are less dispersed than in P along the 45 degree line.

How robust is the BC model to repetition and potential bidder learning?

In order to obtain some insight into the predictive capacity of the BC model with repetition, we first pooled the data for the three bidding rounds and conducted the following regressions for each treatment: Actual (experimental) bids = f (optimal bids, round dummies, risk attitudes, environmental attitudes), where the last two variables were informed by the pre-experimental surveys; the dummies were defined for each round to allow the pooling of data over multiple rounds. The results, reported in Table 2:

- show that a learning effect over the three repetitions is manifest provided the data set is large enough: when both K and P data are pooled together, as well as with the largest of the two data sets, the K set (N = 115);
- reveal an influence of risk attitudes on bidding behaviour;
- reveal only a weak effect of environmental attitudes on bids.

The negative sign of the environmental attitudes coefficient suggests that more conservation-oriented bidders tender lower bids on average. Although this is only a weak effect, it does meet one's expectations, given that the experiment was conducted within a resource-conservation context. The two other trends suggested by Table 2 also meet our sign expectations: higher risk-

averse bidders tend to bid lower, and bids tend to increase, on average, with repetition. These indicate the existence of a learning effect to the extent that costs do not increase over rounds; as a result, bidder information rents increase instead.

Table 2 about here

Table 2: Statistical test for the quality of bid prediction by the BC model in K and P (all three rounds combined)

*Equations: Actual bid = A*Optimal bid + B*Round dummy + C*Risk + D*Env*

Treatment	Regression coefficient (t - statistics)				Model R ²
	Optimal bid	Repetition effect	Risk proneness	Conservation oriented	
K+P (n = 190)	0.86*** (24.1)	10.60*** (3.81)	0.30*** (5.21)	-4.46** (-2.13)	0.96
K (n = 115)	0.89*** (20.6)	9.90*** (3.14)	0.24*** (3.22)	NS	0.97
P (n = 75)	0.82*** (14.3)	NS	NS	NS	0.94

Legend :
 NS = non significant
 *** = significant at 1% level
 ** = significant at 5% level

By un-pooling the data and considering each round successively, it is possible to see whether the model's capacity to predict the experimental bids is maintained or deteriorates. As Table 3 shows, the results are not consistent across the two replicates. The K data set reveals a steady decline in the capacity of the model to accurately predict the experimental bids: indeed, the regression coefficient increases its distance from the maximum value of 1 (by 0.07; 0.16 and 0.22). Given that the regression constant increases, we are seeing the experimental bid curve gradually becoming more horizontal than the predicted bid curve: the model increasingly under-estimates low-cost bids. The P data show no such trend.

Table 3 about here

Table 3: Quality of bid prediction: $b = f(b^*)$: $b = u(b^*) + cst$ using individual β_i 's

	K1	K2	K3	P1	P2	P3
Coefficient	0.93	0.84	0.78	0.79	0.83	0.83
t statistic	14.4	9.0	12.7	6.3	9.3	8.3
Regr. constants	25	62	66	NS	NS	NS

Notes: All linear coefficients are significant at the 1% confidence level.

The regression constants in the P data set are not statistically significant.

The predictive capacity of the model can be envisaged from another angle, by considering the dispersion of bid predictions over the three bidding rounds. Table 4 (rows 1, 2 and 3) first compares the coefficients of variation of observed and computed bids. If the difference between them increases over rounds, then the predictive capacity of the model can be said to deteriorate. Again, the K and P data do not tell the same story. The K data show a deteriorating trend while the P data do not. Row 4 measures the mean of the absolute difference between the computed and the observed bids. It does not reveal any consistent increase over the three rounds, in neither of the two replicates. From this perspective, it does not appear that the BC model loses much of its predictive capacity with repetition. Bid computations were of course done in the spirit of the LH 1997 model: in each round bidders were asked to provide their estimate of the bid cap (thus β_i^1 , β_i^2 and β_i^3) and these, as well as their individual costs (c_i^1 , c_i^2 and c_i^3), were used to compute the predicted bids in each round.

Table 4 about here

Table 4 : Dispersion of bid predictions over the three bidding rounds

			K data			P data		
#		K1	K2	K3	P1	P2	P3	
1	CV (observed bids)	40%	29%	26%	43%	32%	36%	
2	CV (computed bids)	43%	35%	36%	44%	35%	38%	
3	CV(obs) – CV(cmp)	3%	6%	10%	1%	3%	2%	
4	mean (cmp. – obs.)	€3 (\$37)	€7 (\$59)	€1 (\$49)	\$35	\$21	\$24	

Legend: CV = coefficient of variation SD = standard deviation
 cmp = computed bids obs = observed bids

Note: 1 € = \$1.60 approx. at the time of writing

This lack of deterioration in the model’s predictive capacity is somewhat surprising, given the evidence that bidders are learning the implicit reserve price. How can this be explained? The only possibility, given the fact that bids depend here only on costs and expected bid caps, is that some adjustment is happening in how this dependence operates. A natural hypothesis to make is that, as bidders learn, their bids should depend less on their costs and more on their bid cap expectation. Is this the case?

As Table 5 shows, this is indeed the case for the K data set, at least regarding the falling dependence on costs. This trend is also visible in the P data, though less clearly. The corresponding rise in dependence on the bid cap expectations is also apparent, though less marked than that on costs. In fact, the main changes happen after the first round, with the changes between the second and third rounds being less marked. Nevertheless, it can be said that bids do indeed show a tendency to form by gradually shifting the weight from cost information to bid cap expectations. A greater number of experimental repetitions would probably clarify and settle this point.

Table 5 about here

Table 5 : Linear coefficients u and v of $b = f(c, \beta) = u.c + v.\beta + cst$

(t-statistics given between brackets)

	K data			P data		
	K1	K2	K3	P1	P2	P3
Cost coeff.	0.64 (8.86)	0.30 (3.80)	0.26 (6.23)	0.81 (14.05)	0.47 (4.92)	0.62 (7.03)
Beta coeff.	0.25 (3.30)	0.72 (6.65)	0.60 (9.58)	0.12* (2.57)	0.34* (2.59)	(NS)
Constant	42 (3.21)	(NS)	37 (3.34)	42 (2.67)	47** (2.00)	76 (3.22)

Results with an * mean significant only at the 5% confidence level.

Results with an ** mean significant only at the 10% confidence level.

Results with no asterisk are significant at the 1% confidence level.

The next step in understanding the experimental results is clear: the results above beg a closer analysis of how bid cap expectations are formed – something about which, we may remember, the LH (1997) model has nothing to say: it provides no ‘ β -theory’. Do the experimental results provide us with any insights? Do the expected bid caps evolve in any specific way? As shown in Table 6, the average absolute difference between individual expectations and the actual cut-off bid (the marginal bid) falls from the first to the third round, by 45% and 58% in the K and P replicates, respectively. This fall shows that bidders in both replicates adjust their expectations of the bid cap (the maximum acceptable bid), reflecting some learning of its position.

Table 6 about here

Table 6: Average absolute difference between expected and actual cut-off bids in % of the actual cut-off bid

Replicate	Round 1	Round 2	Round 3
K	29%	20%	16%
P	43%	26%	25%

This raises a question regarding the behaviour of the expected β_i over successive rounds, as observed in the experiments. Figure 3 tells an interesting story. It shows, first, that their individual

distribution depends on the cost-related information available to bidders, the cost quartile to which each bidder knows he or she belongs. On average, high cost bidders expect the bid cap to be higher than low cost bidders. That is, expected bid caps are an increasing function of one's own known costs. Secondly, the bids appear to approximate a normal distribution around the quartile mean. Note that this is totally independent of the BC model's assuming a uniform distribution of where the bid cap might lie, which holds for an individual bidder, not across bidders. Thirdly, the variance of the bid cap expectations falls with higher known costs, and remains true across repeated rounds. Fourthly, the variance of bid cap expectations falls with repetition. These last two results can be observed by comparing the two panels in Figure 3.

Whereas this fourth result can easily be attributed to bidder learning, the third result can not. The decrease in variance with knowledge of higher costs is simply due to the smaller margin between one's known cost and the maximum acceptable bid which appears most likely to the bidder, given her knowledge of the budget constraint and the number of bidders. This also explains why β_Q 's increase with costs on average.

Figure 3 about here




Figure 3 : Influence of repetition on the distribution of bid cap expectations

Can we go any further? How exactly do the β 's depend on costs? To answer this question,

the relationship between individual β_i 's and the corresponding individual costs c_i was investigated.

In round 1, the K replicate yielded the following relationship:

$$\beta = 0.53 c + 133 \tag{6}$$

and the P replicate yielded

$$\beta = 0.44 c + 180 \tag{7}$$

In both cases, the statistical fit was quite poor (low R^2 and low statistical significance of the c coefficient), but the two samples do yield similar relationships. The value of the constants 133 and 180 are somewhat larger than the average costs of about 125. Across the two samples, an average β would equal 200 and 235 respectively, a difference of 8%.

However, the quality of these relationships is also subject to how the cost coefficient behaves over repeated rounds. Table 7 yields bad news. For one to have some faith in this relationship, one would like the trend in the cost coefficient to be consistent across the two replicates and, secondly, to show a monotonic trend over repetitions. Neither of the two happens. The K replicate exhibits an imperfect downward trend and the P replicate exhibits no trend at all.

Table 7 about here

Table 7 : Linear fit of $\beta = f(c) = u.c + cst$, using individual experimental data

$\beta = u.c + cst$	Round 1	Round 2	Round 3
K data (t statistic)	0.53c + 133 (4.69)	0.38c + 170 (3.46)	0.39c + 156 (4.48)
P data (t statistic)	* 0.44c + 180 (1.94)	0.49c + 157 (4.39)	0.43c + 163 (4.04)

The * indicates significance at the 6% confidence level only. All other results are significant at 1%.

The conclusion is clear. With the information available from our experiments, it is not possible to produce a β -theory; that is, a theory describing the formation of bid cap expectations by bidders, based on their imperfect knowledge of costs. For the time being, we only have at our disposal some empirical relationships. If we interpreted the results of Table 7 as indicating invariance of β formation given the parameters of the tender (budget constraint, number of bidders

and cost distribution), then, at the very most, one could formulate an average relationship, formed by taking the average of all the coefficients over both replicates, yielding:

$$\beta = 0.44 c + 160 \quad (8)$$

Based on the empirical relationships of Table 7, we can construct Table 8 which will be needed in the next section for examining the capacity of the BC model to predict the institution's performance. Since the policy maker will be primarily interested in knowing in advance whether to *initiate* the process or not, we shall focus on first round results only. Columns (1), (2), and (3) in both data sets K and P represent actual experimental data, costs, bids and expected bid caps respectively. Column (4) is directly computed using the relationships in Table 7 and the costs in column (1). Column (5) shows the quartile-specific average of the individually computed bids, using equation (3''). Column (6) computes an average quartile-specific bid using columns (1) and (4) and equation (3').

Table 8 about here

Table 8: Computations for assessing the model's predictive performance, given limited information on bidder costs

	(1)	(2)	(3)	(4)	(5)	(6)		(1)	(2)	(3)	(4)	(5)	(6)
K1	Ave	Ave	Ave		Ave		P1	Ave	Ave	Ave		Ave	
Quartiles	c_i	b_i	β_i	β_Q	b_i^*	b_Q^*	Quartiles	c_i	b_i	β_i	β_Q	b_i^*	b_Q^*
Q1	24	112	157	146	92	85	Q1	24	68	204	191	114	107
Q2	88	131	162	179	125	134	Q2	89	143	198	219	144	154
Q3	159	195	213	217	186	188	Q3	160	181	206	250	183	205
Q4	228	254	251	254	240	241	Q4	240	261	292	286	266	263
Ave K1	123	174	196	198	160	161	Ave P1	122	157	223	234	172	178

How well does the BC model predict the institution's performance?

The previous results form the groundwork necessary to examine whether the BC model is sufficiently reliable to be used for guiding environmental policy when a tendering system appears to be a reasonable option. In this respect, the following results add to those obtained by Schilizzi and Latacz-Lohmann (2007). We analyse the performance of the tendering institution in the same

terms as these authors, with the additional criterion of ‘% N abated’, a measure of the quantity of environmental service provided (see fourth criterion in Table 9). This percentage is defined as the amount abated by bidders accepted into the scheme relative to the amount that would have been abated had all bidders participated. Also, we limit ourselves here to round 1 since the policy maker will be primarily interested in knowing in advance whether to *initiate* the process or not.

Table 9 first provides the performance estimates as a function of the information available on bidders’ costs (both panels A for replicates K and P). This information on costs decreases from left to right. The first two columns represent the information known to the experimenter while the third and fourth columns represent that known to the policy maker. Panels B (for both replicates) take the institution’s performance measured with actual experimental bids (column 1) as a benchmark for assessing the performance estimated with the three computed bids (columns 2 to 4). Panels C and D consider a counterfactual aspect of the model’s predictive performance: whereas columns (3) and (4) use the average or quartile cost information for evaluating performance, as known to the policy maker, columns 5 and 6 evaluate performance using the real costs (and β_i ’s) known to the experimenter, as if evaluating the policy-maker’s own assessment. That is, to what degree would the assessment of the tender be ‘off’ compared to what it would be had the underlying costs been known. Although hypothetical in nature, since such knowledge of costs would undermine the need for a tender in the first place, this, as we shall see, is an informative counterfactual.

The model’s predictive capacity from the point of view of the experimenter. In terms of the model’s predictive capacity using the experimentally observed β_i ’s, both replicates K and P perform roughly as well (Table 9, columns 1 and 2 in panels B): the modelled predictions are never off by more than 10% on any of the performance criteria. The results are all statistically significant at least at the 5% confidence level. There seems to be no systematic trend of over- or under-estimation of model performance across the two replicates. Given our experimental data, the BC

model, implemented with bids computed using individually known β_i 's, predicts auction performance with an error of less than 10% compared to what it would have predicted if real (ex-post) bids were used.

The model's predictive capacity from the point of view of the policy maker. By contrast, when simulating the policy maker's limited information on costs, the K replicate yields 'better' results than the P replicate, particularly in the QCS scenario (Table 9, column 3 in panels B). In the K replicate, estimating the institution's performance using only a four-point estimate of bidder costs leads to surprisingly accurate predictions when column one (real bids) is taken as a benchmark: they are never off by more than 2%. On the other hand, the predictions in the P replicate can be out by as much as +42% for bidder information rents (the amount they are paid over and above their participation costs), in this case grossly over-estimating them at the rate of 2.09 relative to the 'real' measure of 1.47.

Column (4) reveals that if information on bidders' costs is limited to a single (average) point estimate, then model predictive performance is unreliable. In both replicates, results can be off by more than 20% on any of the performance criteria.

Assessing the policy-maker's assessment of the tender's performance. When discussing the results for the 'average' and 'quartile' scenarios in panels A, we were doing so in terms of what we may call their 'apparent' performance; that is, in terms of the performance criteria as measured using cost information assumed known to the policy maker. Panels C examine predictive performance using the individual costs known to the experimenter but not to the (simulated) policy maker, and panels D take the corresponding measures in panels A as a benchmark for panels C. Thus the budget cost-effectiveness of 2.69 in panel C of K (column 5) compares to the corresponding 'quartile' measure of 2.66 in panel A of K (column 3) to the ratio of 0.99, shown in panel D. When considered from this vantage point, model performance (as opposed to the institution's performance) using quartile cost information is remarkably accurate in both replicates (never off

by more than 6%), while the model using the single-point ‘average cost’ information performs poorly on all criteria, being off by at least 30% (budgetary cost-effectiveness in K) and by as much as 61% (information rents in P). This means that if only single point cost estimates are known, the performance predicted by the model will appear to be much higher or much lower than it really is. By contrast, such over- or under-estimates of institution performance are small if quartile cost estimates are known, at least on the basis of our experimental results.

Conclusions regarding model capacity to predict the institution’s performance

From the results described above, three things stand out. The BC model predicts well, as already discussed in an earlier section, when individual β_i 's, as expressed by the bidders themselves, are known. Table 3 showed the fit to the 45 degree perfect prediction line to be close to 1 with high statistical significance. Second, the K replicate shows that, when cost quartiles are known, and thereby the corresponding four β_Q 's, the model predicts nearly as well as when individual β_i 's are known, whereas model performance deteriorates significantly when only one cost estimate (and therefore only one β_A) is known. Thirdly, the P replicate does not yield as good measures of model predictive performance as the K replicate. This may be due to specific bidder behaviour or to the size of the bidding populations, or both. The K sample was 63% larger than the P sample (44 versus 27) and more homogenous: all students were of the same class, whereas the P sample included students with different backgrounds and academic levels. The K results are likely to be more reliable than the P results.

From the experimenter's (and theoretician's) point of view, if one provides the BC model with information on individual bidder costs and bid cap estimates, as per the model proposed by LH in 1997, then the model predicts experimental bids reasonably well, unless bidders have a high degree of risk aversion. In that case, the risk-neutral optimal bids are no longer a good enough approximation.

The model predicts reasonably well only to the extent that it uses information unknown to the client, the policy maker. Individual costs and bid cap estimates are unknown. However, if the policy maker has quartile information on bidders' costs, the errors he will make in estimating the institution's performance are likely to be small: this estimate will not be very far from that which he would have obtained had he known individual costs more accurately.

This encouraging conclusion does not carry through to single-point cost estimates, which lead to errors in performance predictions that can differ significantly from those which would have been obtained with better knowledge of bidders' costs. To make matters worse, our results do not allow us for the time being to specify the direction of these large errors: they could as easily be over- or under-estimates.

These conclusions can be summarized by saying that it is worth the policy maker's effort (and cost) to gather a certain amount of information on the *variability of costs* across potential bidders before running a conservation tender: too little information will not allow him to assess whether using a tendering mechanism is a reasonable policy given existing alternatives, and too much information will make the use of a tendering mechanism redundant, in addition to incurring high collection costs. Determining the amount of information that needs to be collected – quartiles, quintiles or deciles, for example – is a question beyond the scope of the present paper, but would depend, among other things, on the costs of collecting that information, which is likely to be highly specific to the policy region. However, on the basis of our results, it is reasonable to recommend as a rule of thumb that the procurer obtain a three or four point distribution of cost-estimates in the target region.

VI. CONCLUSIONS

Summary of results

In a paper published in 1997 in the *American Journal of Agricultural Economics*, Latacz-Lohmann and Van der Hamsvoort (LH) proposed a model for budget-constrained auctions to fill a gap existing in the literature. While auction theory has mostly focused on target-constrained auctions, government-funded tenders for allocating conservation contracts to landowners have nearly always been carried out under a budget constraint. The LH model subsequently inspired conservation-oriented tendering programs, the BushTender program in Victoria, Australia, being a notable example (Stoneham et al., 2003). Because of its ‘advisory’ success, the present paper investigates how reliable the LH model really is for guiding environmental policy.

The first requirement is that an economic model make reasonable assumptions about people’s behaviour, in this case, bidding behaviour. Accordingly, controlled experiments with university students were carried out in two different countries that simulated the conditions of the conservation-oriented budget-constrained (BC) tendering mechanism. In addition, three consecutive rounds were carried out to see whether repetition induced any noticeable effects. Given the information assumed known in the model by bidders, and reproduced in the experiments, the first question was whether bids computed using the model predicted the experimental bids to a satisfactory degree. This was found to be true not only in the first round, but also in subsequent rounds. To the extent that evidence of bidder learning was found to be happening, the BC model appears to maintain its capacity to predict actual bidding behaviour at least over three repetitions, and this, in spite of the fact that, as shown in Schilizzi and Latacz-Lohmann (2007), the performance of the institution itself deteriorates.

This result was somewhat intriguing and was therefore further investigated. In particular, the way bidders collectively adjust their bidding strategies was studied, and revealed that, in the terms of the BC model, bidders progressively shift the emphasis in bid formation from using information based on their own (roughly known) position in the cost distribution to that based on their bid cap expectations. Given that the cost distribution remained unchanged across the three

rounds, this naturally led to investigate how bid cap expectations evolve. Here, however, the experimental setup, and the BC model itself, showed its limits: no consistent pattern was discernible regarding the formation of bid cap expectations, at least not in terms of the only piece of information allowed, the bidders' own roughly known cost positioning (cost quartiles). However, a clear set of empirical relationships was established, and was used for the next step in the study.

A theoretical model by itself can only be useful for guiding policy in broad qualitative terms. This study investigated whether combining a theoretical model with results from appropriately designed experiments could lead to improved information for policy guidance. The answer seems to be yes. The difference between the experimenter and the policy maker is the nature and quality of the information available. The experimenter can help the policy maker by comparing experimental results with and without certain pieces of crucial information. In this case, the crucial information refers to bidders' participation costs, as well as their bid cap expectations. Whereas the policy maker only has rough average estimates of the first and no information on the second, the experimenter has accurate information on both. The question was, could the BC model be used to predict the performance of a BC tendering mechanism given that other competing policy mechanisms might also be available?

It was found that insufficient information about bidders' costs led to performance predictions that differed markedly from those made with full experimental information. It appears a certain amount of information needs to be collected before being able to decide whether a BC tendering mechanism will perform satisfactorily or not, using performance criteria such as budget cost-effectiveness or the rate of bidder information rents. At the same time, however, having too much information on costs defeats the purpose of a tender since contracts can then be priced as a direct function of the known costs.

Implications for policy

Some clear conclusions result from this study for policy purposes. This paper does not explicitly study how a budget-constrained tendering mechanism compares with other institutional mechanisms. Schilizzi and Latacz-Lohmann (2007) addressed one aspect of this question by making the comparison with traditional fixed price schemes. The present work does however provide information as to what a policy maker would need to do in order to know whether going ahead with a specific type of tender was likely to be worth it or not.

If the procurer has decided to use a budget-constrained tender, provided the stakes are high enough to warrant the associated expenses, it is suggested he use the skills of the appropriate specialists to simulate, in controlled lab experiments, the conditions of the projected tender: ratio of budget to number of eligible bidders, number of repetitions projected, if any, and appropriate cost distributions. He should invest some (limited) effort in collecting some information about the distribution of bidders' likely participation costs, which will usually be done on a mixture of science-based analysis of the environmental works needed and of information on landholders' properties already available. Estimates of minimum and maximum possible participation costs would help define boundaries to the cost distribution in the target region. In the future, with sufficient experience and controlled experiments, it may even be possible to apply a sort of 'benefit transfer' approach and not need to run further specifically designed experiments: knowledge of the key tender parameters may suffice for predicting the value of running a tender.

In the meantime, the use of the BC model designed by LH in 1997, coupled with results from its appropriately designed experimental implementation, should provide useful guidance for knowing whether a specific BC tender is likely to perform satisfactorily or not, using performance criteria such as those proposed by Schilizzi and Latacz-Lohmann (2007). The key insight from this study is that the value of running such a tender depends critically on the knowledge of bidders' participation costs.

Limitations and further research

We end this paper by critically reviewing the key results and highlighting their limitations. The BC model proposed by Latacz-Lohmann and Van der Hamsvoort (LH) in 1997 appears to either under-estimate or over-estimate the experimental bids. It would appear, given the information collected from experimental subjects, that risk aversion could be a determining factor. The version of the 1997 BC model used was the risk-neutral formulation; it would be interesting to make explicit the general formulation of the risk-averse version and check whether risk attitudes do indeed explain the discrepancies observed using risk-neutral optimal bids. An interesting extension of the experimental setup would be to sort a large enough population of participants into groups with different levels of risk aversion, previously measured. However, this is probably an issue of secondary importance, as the discrepancies were small.

A second extension would be to computerise our pencil-and-paper experiment and examine a greater number of repetitions, in particular to see whether a specific pattern in the evolution of bid cap estimates emerges, which might suggest the ‘ β -theory’ the BC model presently lacks. What information do bidders use when trying to guess the bid cap? Several exploratory statistical analyses of the experimental data were carried out, but to no avail; probably because of the way the experiments were designed. They aimed at testing the existing model, not at developing it further. The existing model only suggests costs and one’s own past bids as relevant information; but there could be other sources of relevant information.

Thirdly, the BC model gives equal weight to both the known costs and the expected bid caps, but it must be remembered that it formulates optimal bids rather than learning bidders. It would therefore be necessary, in a future study, to know if after many repetitions where, in each round, bidders express their bid cap expectations, the relative weights tend towards a fifty-fifty split or not. If such were not the case after controlling for the risk aversion effect, it would, *a posteriori*, invalidate the assumption of uniform distribution made by the model about *individual* bid cap expectations.

The fact that the BC model formulates optimal bids rather than ‘learning bids’ and yet manages to predict experimental bids satisfactorily must be set against the fact that it maintains its

predictive capacity over several repetitions where evidence of bidder learning is manifest (Schilizzi and Latacz-Lohmann, 2007). Does this lead to a contradiction? How can there be learning going on if the model, computing optimal “fully learned” bids predicts well in the first, inexperienced round? This would mean bidders bid optimally from the start and had nothing to learn. Is it that ‘learning’ is not the correct interpretation of what is going on? Or are there different types of learning, as clearly described for instance by Brenner (1999), and bidders in the experiments do not ‘learn’ in the terms allowed by the formulation of the BC model? As often happens in science, it could be those terms in which the problem is framed that may be creating the problem.

Paul Klemperer, a widely cited auction specialist, is known to have stated that auctions are an excellent institution for studying all kinds of economic behavior (Klemperer, 2004: chapter 2). The depth of perspective provided by our experimental results seems to confirm Klemperer’s judgment. In spite of extensive analysis, we cannot say to have exhausted the possibilities made available by our data.

REFERENCES

- Cox, J.C., Smith, V. L. and Walker, J.M. 1984. “Theory and Behavior of Multiple-Unit Discriminative Auctions.” *The Journal of Finance*, 39 (4): 983-1010.
- Cason, T., Gangadharan L. and Duke C. 2003. “A laboratory study of auctions for reducing non-point source pollution.” *Journal of Environmental Economics and Management* 46: 446-71.
- Cason, T. and Gangadharan, L. 2005. “A laboratory comparison of uniform and discriminative price auctions for reducing non-point source pollution”, *Land Economics* 81 (1): 51-70.
- Cassady, R. Jr. 1967. “*Auctions and Auctioneering*.” University of California Press, Berkeley, CA.
- Cummings, R.G., C.A. Holt, and S.K. Laury. 2004. “Using Laboratory Experiments for Policymaking: An Example from the Georgia Irrigation Reduction Auction.” *Journal of Policy Analysis and Management*. 22 (2): 341-63.
- Engelbrecht-Wiggans, R. 1980. “Auctions and bidding models.” *Management Science* 26: 119-21.

- Guala, F. and Mittone, L. 2005. "Experiments in economics: External validity and the robustness of phenomena." *Journal of Economic Methodology* 12 (4): 495 - 515
- Hailu, A. and Schilizzi, S. 2004. "Are auctions more efficient than fixed-price schemes when bidders learn?" *Australian Journal of Management* 29 (2): 147-68.
- Hailu, A., Schilizzi, S. and Thoyer, S. 2005. "Assessing the performance of auctions for the allocation of conservation contracts: Theoretical and computational approaches." Selected Paper for the American Agricultural Economics Association Annual Meeting, Providence, Rhode Island, July 24-27, 2005. Available at: http://agecon.lib.umn.edu/cgi-bin/pdf_view.pl?paperid=16034
- Harris, M. and Raviv, A. 1979. "Allocation mechanisms and the design of auctions." *Econometrica* 49: 1477-99.
- Johansson, R. 2006. "Participant Bidding Enhances Cost Effectiveness." USDA ERS Economic Brief Number 3 (March), 6pp. www.ers.usda.gov/Publications/EB3/. Last accessed August 2008.
- Kagel, J.H. and Roth, A.E. (eds) 1995. "*The Handbook of Experimental Economics*." Princeton University Press: Princeton, N.J.
- Klemperer, P. "Auction theory: a guide to the literature", *Journal of Economic Surveys*, 13 (3), 1999, 227-86.
- Klemperer, P. 'What really matters in auction design', *Journal of Economic Perspectives*, 16(1), 2002, 169-89.
- Klemperer, P. 2004. "*Auctions: Theory and Practice*." Princeton University Press: Princeton, N.J.
- Latacz-Lohmann, U. and Van der Hamsvoort, C. 1997. "Auctioning conservation contracts: a theoretical analysis and an application." *American Journal of Agricultural Economics*. 79: 407-18.
- Latacz-Lohmann, U. and Van der Hamsvoort, C. 1998. "Auctions as a means of creating a market for public goods from agriculture." *Journal of Agricultural Economics* 49 (3): 334-345.
- Laury, S. 2002. "Enhancing and improving designs for auction mechanisms that can be used by the EPD for irrigation auctions." Water Policy Working Paper #2002-012, Georgia State University - Andrew Young School of Policy Studies.
- McAfee, R.P. and McMillan, J. 1987. "Auctions and bidding." *Journal of Economic Literature* 25 (June): 699-738.
- Milgrom, P. 1985. "Auction theory." In Truman, B. (ed.). *Advances in Economic Theory: Fifth World Congress*, Cambridge University Press.
- Milgrom, P. 1989. "Auctions and bidding: a primer." *Journal of Economic Perspectives* 3 (3): 3-22.

- Müller, K. and Weikard, H.-P. 2002. "Auction mechanisms for soil and habitat protection programmes." In Hagedorn, K.: *Environmental Co-operation and Institutional Change*, Cheltenham, UK: Edward Elgar, pp. 200-11.
- Plott, C. and Smith, V.L. 2008. "*Handbook of Experimental Economics Results*." Edward Elgar: Cheltenham, UK.
- Riechelderfer, K. and Boggess, J. 1988. "Government decision making and program performance: the case of the conservation reserve program." *American Journal of Agricultural Economics* 70: 1-11.
- Rothkopf, M.H. and Harstad, R.M. 1994. "Modeling competitive bidding: a critical essay." *Management Science* 40 (3): 364-84.
- Schilizzi, S. and Latacz-Lohmann, U. 2007. "Assessing the performance of conservation auctions: an experimental study." *Land Economics*, 83 (4): 497-515.
- Schramm, A. 2005. "Artificiality: The tension between internal and external validity in economic experiments." *Journal of Economic Methodology* 12(2): 225-37.
- Stoneham, G., Chaudri, V., Ha, A. and Strappazon, L. 2003. "Auctions for conservation Contracts: An Empirical Examination of Victoria's BushTender Trial." *Australian Journal of Agricultural and Resource Economics* 47 (4): 477-500.
- Vickrey, W. 1961. "Counterspeculation, auctions and competitive sealed tenders." *Journal of Finance* 16: 8-37.
- Vickrey, W. 1962. "Auctions and bidding games." In *Recent Advances in Game Theory*, Princeton University Press, pp.15-27.
- Wilson, R. 1992. "Strategic Analysis of Auctions." In Aumann, R.J. and Hart, S. (eds). *Handbook of Game Theory*, Vol. 1, pp. 228-79.

Table 9: Comparative predictions of auction performance for K and P, round 1

		Column number	(1)	(2)	(3)	(4)	(5)	(6)	
K1			Actual	Predicted (apparent) performance			Real but unknown predicted performance		
Performance criteria	Measure	Panel A	bi	Bi	Beta-Q	Ave Beta	Panel C	Beta-Q	Ave Beta
			Real bids	Indiv Beta's	4-bid estim	1-bid estim	4-bid estim	1-bid estim	
Budg. Cost-Effectiveness	Payts / kg N		2,72	2,56	2,66	2,68	2,69	3,83	
Econ. Cost-Effectiveness	Opp Costs / kg N		1,67	1,71	1,68	2,06	1,65	1,41	
Information Rents	Payts / Opp Costs		1,62	1,50	1,58	1,30	1,64	2,71	
Environ. Effectiveness	% N abated		54%	58%	55%	55%			
			Using real bid performance as benchmark				Apparent / real ratio		
Performance criteria	Measure	Panel B	bi	Bi	Beta-Q	Ave Beta	Panel D	Beta-Q	Ave Beta
			Real bids	Indiv Beta's	4-bid estim	1-bid estim	4-bid estim	1-bid estim	
Budg. Cost-Effectiveness	Payts / kg N		1,00	0,94	0,98	0,99	0,99	0,70	
Econ. Cost-Effectiveness	Opp Costs / kg N		1,00	1,02	1,00	1,23	1,02	1,46	
Information Rents	Payts / Opp Costs		1,00	0,92	0,98	0,80	0,97	0,48	
Environ. Effectiveness	% N abated		1,00	1,07	1,02	1,00			
P1			Actual	Predicted (apparent) performance			Real but unknown predicted performance		
Performance criteria	Measure	Panel A	bi	Bi	Beta-Q	Ave Beta	Panel C	Beta-Q	Ave Beta
			Real bids	Indiv Beta's	4-bid estim	1-bid estim	4-bid estim	1-bid estim	
Budg. Cost-Effectiveness	Payts / kg N		2,49	2,49	3,15	3,03	3,21	4,70	
Econ. Cost-Effectiveness	Opp Costs / kg N		1,69	1,86	1,51	2,07	1,45	1,24	
Information Rents	Payts / Opp Costs		1,47	1,34	2,09	1,46	2,22	3,78	
Environ. Effectiveness	% N abated		58%	56%	44%	48%			
			Using real bid performance as benchmark				Apparent / real ratio		
Performance criteria	Measure	Panel B	bi	Bi	Beta-Q	Ave Beta	Panel D	Beta-Q	Ave Beta
			Real bids	Indiv Beta's	4-bid estim	1-bid estim	4-bid estim	1-bid estim	
Budg. Cost-Effectiveness	Payts / kg N		1,00	1,00	1,27	1,22	0,98	0,64	
Econ. Cost-Effectiveness	Opp Costs / kg N		1,00	1,10	0,89	1,23	1,04	1,66	
Information Rents	Payts / Opp Costs		1,00	0,91	1,42	0,99	0,94	0,39	
Environ. Effectiveness	% N abated		1,00	0,97	0,76	0,83			

APPENDIX I


Pages 2 and 4 of the Budget-Constrained Auction sheet in P

(Page 1 provided the 'story' and the motivation.)

Individual farm data (page 2)

to work out the costs of your participation in our P River protection program.

Suppose you are a horticulturalist and producing vegetables for P. Output as a function of NP fertiliser use is given by the following graph:



The optimal fertiliser amount maximises value of output minus cost of inputs (NP fertilisers).

This results in the following:

	With NP = 80	With NP*	Difference
Net revenue (\$/ha)			

My costs of participation are \$/ha (= the income difference)

Important:

- Your costs of participation are known only to you and your private adviser; they are not known by the environmental authority, or anyone else.
- Your competitors all have different participation costs. So that you may have a better idea of how you compare relative to your competitors, we give you the following information: you are in one of the following four quartiles:

--	--	--	--

lower quarter

second quarter

third quarter

upper quarter

(Page 3 provided "some advice from your private consultant")

Bidding sheet (page 4)

Now it is time you put in your bid. Please first write in your full name. We shall need it to pay you your gains if you are among the winners.

Name:

1) First please write down the highest possible bid you believe will be accepted. This must be your best guess:

Highest acceptable bid (best guess): \$/ha

2) Now please write in the amount we must pay you so that you accept to participate in our P River protection program:

Your bid: \$/ha

The selection of participants will be made on the ground of their bid in \$/ha. The lowest bid will be selected first, then the second lowest, then the third lowest, and so on until the available budget of \$2300 is exhausted.

For paying the winners in **real money**, the following rules hold:

- The successful bidders will be paid, not their bid, but the gains from their participation in the program, that is: bid minus participation costs.
- Unfortunately, because of limited research funds, we cannot pay out the full value of the gains, but only a fixed percentage of the gains. This percentage will be calculated after the end of the bidding session. Of course, the higher your gains, the higher your proportional payment. For this session the funds we have available for payment to this group total an amount of approx. \$300.

APPENDIX II : Raw experimental data from both replicates (c_i , β_i and b_i)

Data are ordered by bids, with indication of cut-off line.

#	K1			K2			K3			P1			P2			P3		
	c_i	β_i	b_i	c_i	β_i	b_i	c_i	β_i	b_i	c_i	β_i	b_i	c_i	β_i	b_i	c_i	β_i	b_i
1	18	50	48	5	60	45	39	55	50	13	275	25	39	200	43	5	109	20
2	15	300	60	54	61	60	103	110	109	9	100	50	56	120	98	33	180	63
3	31	85	61	65	113	103	18	200	130	18	400	55	18	150	100	13	100	85
4	54	80	63	27	250	145	59	137	136	33	148	60	87	260	140	65	300	100
5	5	75	75	6	155	154	44	150	140	5	100	65	5	180	145	39	150	130
6	11	100	85	39	160	159	5	175	150	39	80	69	116	160	156	9	145	138
7	77	105	100	11	180	160	6	151	150	49	130	70	103	160	160	77	190	149
8	35	250	100	114	175	164	81	165	151	56	200	100	103	270	160	56	170	150
9	59	125	100	103	210	170	35	200	160	87	160	119	137	200	160	98	250	150
10	81	120	100	35	174	173	11	180	160	108	190	128	77	200	160	103	300	154
11	27	120	109	144	180	173	65	170	165	27	400	150	13	163	163	125	165	160
12	49	135	130	44	200	175	137	168	167	137	155	154	125	200	165	116	190	160
13	98	140	130	150	180	176	114	210	169	65	160	160	27	201	180	49	200	160
14	39	150	133	131	185	179	119	200	169	103	180	160	157	200	185	147	260	175
15	108	150	140	9	195	180	15	200	170	157	85	162	186	191	190	27	200	180
16	44	145	144	31	220	180	131	170	170	171	250	180	164	250	190	179	200	195
17	137	148	148	81	190	185	54	180	175	164	250	186	171	200	195	137	300	200
18	65	300	150	71	190	185	108	180	178	116	300	190	147	290	200	186	210	210
19	119	175	150	59	200	192	77	200	180	186	195	191	108	250	200	210	250	218
20	144	170	160	164	200	199	31	195	181	179	210	191	196	201	201	196	250	220
21	150	188	166	18	220	200	85	200	190	125	300	200	210	225	225	203	223	223
22	6	200	170	49	310	200	125	195	190	237	245	245	203	255	230	221	225	225
23	131	180	170	125	200	200	150	195	194	203	500	253	237	270	250	229	250	250
24	114	178	177	157	200	200	164	200	194	229	400	260	221	251	251	249	260	260
25	186	195	194	177	170	203	22	195	195	249	280	264	229	260	259	258	300	265
26	171	200	198	108	250	210	27	230	195	258	150	268	258	300	280	264	270	270
27	103	250	200	191	212	211	182	200	200	264	175	275	264	400	325			
28	125	200	200	119	215	214	98	218	204									
29	177	210	200	182	220	215	177	217	207									
30	216	219	219	186	220	219	205	215	214									
31	9	275	225	171	230	225	210	220	220									
32	210	140	230	205	230	225	216	220	220									
33	221	235	233	216	170	230	224	240	240									
34	224	250	235	210	240	239	226	250	249									
35	205	240	239	237	195	240	186	230	250									
36	191	250	240	200	350	250	237	255	250									
37	234	246	245	231	250	250	249	215	250									
38	157	256	255	234	200	260	255	215	255									
39	182	350	260	221	300	284	191	400	300									
40	255	270	264	255	350	300	261	185	300									
41	249	279	274	15	345	345												
42	237	295	283	249	190	350												
43	261	290	285															
44	200	295	290															

ENDNOTES

¹ For the sake of clarity, we will use the term ‘tender’ rather than ‘auction’ to refer to buying mechanisms as opposed to selling mechanisms.

² This is not an essential assumption and could be relaxed to include risk aversion, as done by LH in their 1997 paper. However, it would not add much to the present argument and might confuse matters unnecessarily.

³ The optimal bid formula in the single-unit selling case is $B(v) = \frac{n-1}{n}v$ (Wolfstetter, 1996), while the corresponding optimal bid in the single-unit procurement case is given by $B(v) = \frac{n-1}{n}v + \frac{1}{n}$ (Hailu, Schilizzi and Thoyer, 2005).

⁴ By contrast, uniform (second-price) sealed-bid auctions should in theory lead to bidding one’s true opportunity costs; but they have not to date been used in conservation contracting programs.

⁵ This budget constraint of 3900€ was in “nominal” lab euros, which reflected the production functions underlying the costs imposed by reduced nitrogen applications. This was clearly distinguished from the limited funds available for each session of the experiment (300€). Saliency was preserved through the fixed proportionality rate between gains in nominal lab euros and payments in hard currency.