Annals of Mathematics and Artificial Intelligence (2005) 44: 269–302

Auction design with costly preference elicitation *

David C. Parkes
Division of Engineering and Applied Sciences, 33 Oxford Street, Harvard University, Cambridge, MA 02138, USA
E-mail: parkes@eecs.harvard.edu

We consider auction design in a setting with costly preference elicitation. Well designed auctions can help to avoid unnecessary elicitation while determining efficient allocations. Careful design can also lead to more efficient outcomes when elicitation is too costly to permit perfect allocative efficiency. An *incremental* revelation principle is developed and used to motivate the role of *proxied* and *indirect* auction designs. Proxy agents, situated between bidders and an auction, can be used to maintain partial information about bidder preferences, to compute equilibrium bidding strategies based on the available information, and to elicit additional preference information as required. We derive information-theoretic elicitation policies for proxy agents under a simple model of costly elicitation across different auction designs. An experimental analysis demonstrates that indirect mechanisms, such as ascending-price auctions, can achieve better allocative efficiency with less preference elicitation than sealed-bid (direct) auctions because they promote better decisions about preference elicitation.

Keywords: computational mechanism design, incremental revelation principle, meta-deliberation, proxy agents, preference elicitation

1. Introduction

As traditional commerce moves on-line and more business transactions are completed in electronic market places there will be an opportunity for agent-mediated transactions, with software agents responsible for dynamic negotiation between multiple, fluidly changing, market participants [16,48]. The success or failure of agent-mediated electronic commerce will depend in large part on the trust that can be placed in agent mediation, which will in turn depend on the level of optimality that can be provided by agents, and thus on the computational complexity of the decision problem facing agents. It is this back drop that suggests that one important role of market design is to design *simple* worlds that can be effectively populated by automated software agents.

Faster optimization technology can scale to clear large and complex winner-determination problems [27] and bidding languages can be designed to allow for useful expressiveness [34]. Some markets can even be designed with simple *dominant-strategy* equilibrium [1,28]. But, the Achilles heel of electronic markets, the part most resistant to simplification, may be the problem of *preference elicitation*. Bidding agents can only

* A preliminary version of this paper appeared in the Proc. of the IJCAI'99 Workshop on Agent Mediated Electronic Commerce.

bid on behalf of market participants if they are well-informed about a bidder's valuation for different allocations. Yet, bidders often find it costly, or even *impossible*, to provide precise and complete information about their preferences.

The London bus route auction presents a good example of the issues surrounding preference elicitation [5]. With thousands of routes it was unreasonable to expect a prospective bus operator to bid her operating cost for all possible combinations of routes. Rather, the London transportation authority chose to implement a multi-round auction to give bidders the chance to refine their bids in the light of feedback about bids from other potential operators. The auction design allowed participants to bid effectively without exact information about their own costs for all routes.

Thus, while agent-mediated electronic markets with automated clearing engines remove many of the transaction costs associated with traditional markets and enable substantial new possibilities for scale and aggregation, electronic markets also serve to bring the problem of preference elicitation to the fore. We consider the role of auction design in mitigating this problem of costly preference elicitation. Although market design cannot simplify the underlying valuation problem facing a bidder (for instance, the scheduling problem facing a supplier of logistics in a transportation setting [46]), a well-designed market can provide feedback to guide a bidder towards the part of the good space on which she is likely to be most competitive.

It is instructive to compare the elicitation properties of a sealed-bid auction with those of an ascending-price auction. Suppose that truthful bidding is a dominant strategy in the sealed-bid auction, and that straightforward bidding, in which an agent bids for the bundle of goods that maximizes utility at the current prices, is an equilibrium bidding strategy in the ascending-price auction. Clearly, an automated bidding agent in a sealed-bid auction will require complete and exact information about the bidder's valuation to follow the optimal (equilibrium) bidding strategy. In comparison, an automated bidding agent in an ascending-price auction only needs enough information on the bidder's valuation (for example bounds on her value for different bundles of goods) to compute a best-response to the current prices.

This paper provides a theoretical framework in which to study the preference elicitation properties of auction mechanisms. We define the family of *incremental-revelation* mechanisms (IRMs), and prove via a revelation principle that it is without loss of generality to focus on these mechanisms even in settings with costly preference elicitation. The bidding language in an IRM allows a bidder to refine information about her value for different outcomes during an auction, but restricts a bidder to make only direct claims about her value. For instance, a bidder can state "my value for bundle AB is at least \$100", but cannot bid directly for the bundle at current prices. The revelation principle also serves to motivate the role of *proxy agents*, that sit between bidders and an auction and maintain a partial model of a bidder's preferences.

We illustrate, with a simple model of costly preference elicitation, that proxied and indirect mechanisms can achieve better allocative efficiency with less preference elicitation than sealed-bid (direct) mechanisms. As an exemplar of the class of IRMs, this provides new motivation for proxied and ascending auctions [40], demonstrating that they have efficiency advantages over sealed-bid auctions because they promote better decisions about preference elicitation. Price feedback guides bidders to the part of the allocation space on which they are likely to be most competitive, given current bids from other bidders.

Preference elicitation is modeled by assuming that bidders can provide *bounds* on their values for items, and refine these bounds at some cost. In a simple posted-price setting we can derive optimal elicitation and bidding policies, that elicit the right amount of information on a bidders value, making a tradeoff between the cost of queries and the possibility of avoiding mistakes by having better value information. Our approach considers the *expected value of information*, and builds on that proposed by Russell and Wefald [44] for meta-deliberation by bounded-rational computational agents. In moving to even simple dynamic auctions, with strategic interactions with other agents, a full game-theoretic equilibrium analysis of the combined elicitation and bidding policies is not possible given current techniques and computational methods. Rather, we develop heuristic elicitation and bidding policies for proxied and indirect auctions for multi-item allocation problems.

Finally, we present an experimental analysis to compare the allocative efficiency and cost of preference elicitation across different auction designs, given this model of elicitation. In section 4, we consider a single-item allocation problem and a model in which bidders incur a cost C for each query. Indirect auctions, with incremental revelation, are shown to implement efficient outcomes with less elicitation cost, and to extend to problems in which sealed-bid auctions fail. In section 5, we consider a multi-item (non-combinatorial) allocation problem and a model in which each bidder has a *budget limit* on the number of queries she will perform. Again, proxied and indirect auctions are shown to dominate sealed-bid and posted-price auctions in terms of allocative efficiency and elicitation properties. We conclude with a discussion of related work and suggest a number of directions for future research.

2. Incremental revelation mechanisms

In this section, we introduce a formal model in which to study the performance of allocation mechanisms in problems with costly preference elicitation. We introduce a variation on the revelation principle, that identifies the role of *incremental-revelation mechanisms* (IRMs) in settings with costly preference elicitation. The strategy space in an IRM restricts a bidder to making statements about her value for goods, and allows a bidder to refine this information during an auction. The incremental revelation principle serves to highlight the role of proxy-based auction design.

We also define *first-best* and *second-best* proxy strategies. In the first-best model, a proxy will assume that the cost of elicitation, while non-negligible (so that elicitation is best avoided if possible) is small enough that a bidder will always be willing to provide information when it can lead to a more accurate bidding strategy. On the other hand, a second-best proxy strategy is one in which a proxy must sometimes bid despite some residual uncertainty about the bidding strategy, because the additional cost of elicitation is not thought justified given the expected value of information.

2.1. Preliminaries: mechanism design

Mechanism design (MD) [19] studies the question of how to implement an outcome with desired properties, such as allocative efficiency, in settings with multiple self-interested agents each with private information about their valuations. Game theory is used to study the properties of mechanisms. Mechanism design embodies a careful separation between the protocol itself, which is under the control of the mechanism designer, and the agent strategies, which are assumed to be out of the designer's control. The second-price sealed-bid, or Vickrey [51] auction, is an example of a simple mechanism. Truthful bidding is a dominant strategy in the Vickrey auction, and the auction is said to be *strategyproof*.

Formally, mechanism design considers a set of choices, K, and a set of agents, I. Let N denote the number of agents. Each agent, i, has a *type*, $\theta_i \in \Theta_i$, which defines its valuation, $v_i(\theta_i, k) \in \mathbb{R}$ across choices $k \in K$. In an auction setting, the choice k can define both the allocation of goods and the payments made by agents. The system-wide goal is represented with a social choice function (SCF), $f : \Theta_1 \times \cdots \times \Theta_N \to K$, that selects choice $f(\theta)$ given type vector $\theta \in \Theta$, where $\Theta = \Theta_1 \times \cdots \times \Theta_N$ is the joint type space. Allocative efficiency is a typical goal in an auction setting, with SCF f defined to select the allocation that maximizes the total value across all agents.

A mechanism, $M = (\Sigma_1, \ldots, \Sigma_N, g)$, defines a set of feasible strategies, Σ_i, for each agent, and an *outcome rule*, $g : \Sigma_1 \times \cdots \times \Sigma_N \to K$, from strategies to a choice. Mechanism M is said to *implement* SCF f if outcome $f(\theta) \in K$ is selected in equilibrium, for all types. Formally, we have $g(s_1^*(\theta_1), \ldots, s_N^*(\theta_N)) = f(\theta)$, for all $\theta \in \Theta$, where s^* is an equilibrium strategy in the game induced by mechanism M. Equilibrium solution concepts adopted in MD include: *Bayesian–Nash* equilibrium, in which each agent's strategy is a best-response in expectation to the strategies of the other agents; and *dominant strategy* equilibrium, in which each agent's strategy is a best-response *whatever* the strategies and private valuations of other agents.

A simple class of mechanisms are the *incentive-compatible direct-revelation* mechanisms (DRMs). In a DRM the strategy spaces, Σ_i, are simply equal to agent type spaces, Θ_i, and the outcome rule, $g : \Theta \to K$, selects a choice based on reported types. A DRM is said to be *incentive-compatible* when truth-revelation is an equilibrium strategy, and $s_i^*(\theta_i) = \theta_i$ for every agent. It follows that the SCF implemented in an incentive-compatible DRM is simply the outcome rule, with $f = g$.

2.2. Failure of the revelation principle with costly preferences

In the absence of computational considerations and without costly preference elicitation, the *revelation principle* can be used to simplify the mechanism design problem. The revelation principle [14,15], states that the outcome of any arbitrarily complex mechanism can be implemented as an incentive-compatible DRM. The intuition behind the revelation principle is simple to state. Given an arbitrary mechanism, M', we can construct an incentive-compatible DRM, M, that asks agents to report their types and commits to simulating the entire system for the original mechanism, both the equilibrium bidding strategies and the rules of the mechanism. As long as M implements the equilibrium strategy profile in M', then every agent should report its true type.[1]

Under classical assumptions, the revelation principle provides a useful simplification to theoretical mechanism design. It can be used both to prove negative results, and also to define incentive-compatible payment rules. However, the revelation principle ignores computational costs [10,25,36]. Consider the following two computational disadvantages of DRMs:

1. All computation (e.g., winner-determination, payments) is centralized in a DRM, while an indirect mechanism can distribute computation to agents. For instance, agents in an ascending-price combinatorial auction must compute the bundle of goods that is a best-response to prices in each round. The combination of the iterative auction with straightforward agent bidding strategies can be viewed as a (decentralized) primal-dual algorithm for solving the underlying allocation problem [11,33,39].
2. Truthful bidding, which is an equilibrium in the incentive-compatible DRMs of the revelation principle, requires that every agent know its complete and exact valuation for all possible (combinations of) goods. This can be infeasible when preference elicitation is costly. In comparison, an indirect mechanism can implement the same outcome as the DRM without agents revealing *or even computing* their exact valuations for all outcomes.

We focus here on the second effect, that of the role of indirect mechanisms in mitigating the cost of preference elicitation. Preference elicitation is not a concern within classic microeconomic theory, where it is standard to assume that every agent can state (at no cost) its value for all possible outcomes [30].

2.3. *An incremental revelation principle*

We now define the family of *incremental revelation mechanisms*, and prove that this space of mechanisms is the appropriate generalization of DRMs to settings with costly preference elicitation.

To develop a revelation principle that is meaningful with costly preference elicitation we introduce an additional requirement to the definition of the equivalence of a pair of mechanisms:

> Mechanisms M and M' are **equivalent in settings with costly preference elicitation** if they implement the same outcome in equilibrium and require the same information revelation by agents.

[1] For instance, the Vickrey auction, or second-price sealed bid auction, is the equivalent direct-revelation mechanism to the ascending-price English auction.

One way to think about this is that we bring the amount of preference elicitation into the definition of the "outcome" of a mechanism, so that two mechanisms are equivalent if they make the same choice in the world and also perform the same elicitation.

In order to define the *information revelation* in a mechanism, we must make explicit the *message space* in the protocol that is used to implement a mechanism. Let X denote the message space, and $\mu_i(s(\theta)) \in X^*$ denote the sequence of messages sent by agent i in mechanism M given strategy profile $s(\theta) = (s_1(\theta_1), \ldots, s_N(\theta_N))$.

Definition 1. The information revelation in mechanism $M = (\Sigma_1, \ldots, \Sigma_N, g)$ by agent i, given type θ_i, types $\theta_{-i} = (\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_N)$, and given equilibrium strategy profile s^*, is defined as:

$$inf_i^M(\theta_i, \theta_{-i}) = \left\{\theta_i' : \mu_i\left(s_i^*(\theta_i'), s_{-i}^*(\theta_{-i})\right) = \mu_i\left(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})\right), \theta_i' \in \Theta_i\right\}.$$

In words, the information revealed by agent i for a particular instantiation of the mechanisms is that the agent identifies itself as one of a set of possible types $inf_i^M(\theta)$. These are all the types that would send the same messages in equilibrium. We can write, $inf^M(\theta) = (inf_1^M(\theta_1, \theta_{-1}), inf_2^M(\theta_2, \theta_{-2}), \ldots, inf_N^M(\theta_N, \theta_{-N})) \subseteq \Theta$, to denote the joint information revealed by agents in a mechanism, given types θ.

Following Parkes [37], let a *query language*, $\mathcal{L}$, define a set of preference queries $Q \in \mathcal{L}$. A query, $Q : \Theta_i \to 2^{\Theta_i}$, is a mapping from the possible types of agent i to a subset of types. Response, $Q(\hat{\theta}_i) \subseteq \Theta_i$ to query Q makes the claim that the agent's type is within the set $Q(\hat{\theta}_i)$. For instance, a query language might ask "is your value for the good at least p?" An agent that wishes to claim that $v_i \geqslant p$ should response "yes" and an agent that wishes to claim that $v_i < p$ should respond "no".[2]

A response to a query is *truthful* if $\theta_i \in Q(\hat{\theta}_i)$, where θ_i is the agent's type and $\hat{\theta}_i$ is the type adopted in providing a response. A sequence of responses, $y^1, \ldots, y^t$, to queries, $Q^1, \ldots, Q^t$, is said to be *valid* if the intersection, $y^1 \cap \cdots \cap y^t \neq \emptyset$. This requires that there is some type that is consistent with the agent's response to every query.

Incremental-revelation mechanisms (IRMs) relax the standard notion of a DRM. In an IRM, an agent's strategy defines the information that it will reveal in response to a query that can be more expressive than the standard "what is your type?" query of DRMs. The outcome rule of an IRM is defined to allow an outcome to be selected before the exact type of agents has been pinned down.

Definition 2 (Incremental revelation mechanisms). An incremental revelation mechanism $M = (\mathcal{L}, \langle g_{\mathcal{L}}, g_K\rangle)$ defines:

1. A **strategy space** Σ_i with strategy $\sigma_i = s_i(\theta_i) \in \Sigma_i$ defining a response $\sigma_i(h, q) \subseteq \Theta_i$ to each query q, given history of queries $h = (q_1, q_2, \ldots)$, with $q_j \in \mathcal{L}$.

[2] Note that we choose to exclude stochastic queries, that provide probabilistic information about an agent's type.

2. A **query rule**, $g_{\mathcal{L}} : \Sigma_1 \times \cdots \times \Sigma_N \to 2^{\Theta}$, that determines the type information $g_{\mathcal{L}}(s(\theta)) \subseteq \Theta$ revealed by agents given query strategies $s(\theta) = (s_1(\theta_1), \ldots, s_N(\theta_N))$.
3. A **choice rule**, $g_K : 2^{\Theta} \to K$, that defines a choice $g_K(\Delta) \in K$ for each possible subset of types $\Delta \subseteq \Theta$.

A query strategy, $s_i(\theta_i) \in \Sigma_i$, defines how an agent will respond to a query given a history of previous queries. Of course, a strategy is only valid if $\bigcap_{j \leqslant m} \sigma_i(h_{<j}, q_j) \neq \emptyset$ for all sequences of queries $\underline{q}$ of length m, with $h_{<j}$ denoting the first j queries in $\underline{q}$ and q_j the jth query. Query rule, $g_{\mathcal{L}}$, determines the final information revealed about types, $g_{\mathcal{L}}(s(\theta)) \subseteq \Theta$, given strategies s. In a realization of an IRM, the query rule can (and should) be history-dependent, with queries chosen adaptively, given the response by agents to previous queries. Choice rule, g_K, defines the final outcome, so that $g_K(g_{\mathcal{L}}(s(\theta))) \in K$ is the choice when agents follow strategy s and have type θ.

An IRM is said to *implement* SCF $f(\theta) = g_K(g_{\mathcal{L}}(s^*(\theta)))$, for all $\theta \in \Theta$, when s^* is an equilibrium strategy profile. Borrowing terminology from that of DRMs, we say an IRM is *incentive-compatible*, or *truthful*, when $\theta_i \in \sigma_i(h, q)$, in equilibrium profile, $\sigma_i = s_i^*(\theta_i)$, for all θ_i, all bidders i, and every history of queries $h = (q_1, q_2, \ldots)$, and every $q \in \mathcal{L}$.

Just as with DRMs, the precise equilibrium solution concept will depend on the model of agents but could (for example) be Bayesian–Nash equilibrium if agents have a distributional model of the types of other agents and play an expected-utility maximizing best-response, or *ex post* Nash equilibrium if the same strategy is an equilibrium whatever the types of other agents as long as they play an equilibrium strategy.

We can now state a revelation principle for IRMs that respects the preference elicitation properties of a mechanism. Consider some mechanism M', and let $\mu_i^z(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})) \in X$ denote the zth message sent by agent i to the mechanism given strategy profile s^* and types θ.

Theorem 1 (Incremental revelation principle). Any SCF, $f(\theta)$, that can be implemented in the equilibrium of some mechanism, $M' = (\Sigma_1, \ldots, \Sigma_N, g)$, with costly preference elicitation, can be implemented in a truthful equilibrium of an IRM, $M = (\mathcal{L}, \langle g_{\mathcal{L}}, g_K \rangle)$, and with the same information revelation as in M'.

Proof. Let $(s_1^*, \ldots, s_N^*)$ denote an equilibrium in M for which $g(s^*(\theta)) = f(\theta)$. We construct an IRM, M, by simulating M' internally, and with a "proxy agent" playing strategy s_i^* for each agent. The proxy agents execute queries on demand, to refine their information about an agent's type when additional information is required to determine the equilibrium strategy s^* for an agent. Concretely, construct the following IRM:

1. Query language, $\mathcal{L}$. Consider step z, when message $\mu_i^z(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})) \in X$ is sent by agent i to the center in M'. Let $\mu_1, \ldots, \mu_h \in X$ denote the *range* of message correspondence, μ_i^z, for different types θ_i. Then $\mathcal{L}$ must contain a query

that asks the agent to place its type into a set j from sets $\{\Delta_1, \ldots, \Delta_h\}$, such that $\mu_i^z(s_i^*(\theta_i'), s_{-i}^*(\theta_{-i})) = \mu_j$ for all $\theta_i' \in \Delta_j$. The IRM language must include a query of this kind for all type profiles, θ, all agents, and all steps z.

2. Query rule, $g_{\mathcal{L}}(s(\theta)) = \mathit{inf}^{M'}(s^*(\theta))$, because the IRM mechanism, M, will execute any query, $Q \in \mathcal{L}$, that is required to simulate strategy profile s^* (but only those queries).

3. Choice rule, $g_K(\Delta) = f(\theta')$, some $\theta' \in \Delta \subseteq \Theta$ when choice $f(\theta)$ is constant for all $\theta \in \Delta$, and $g_K(\Delta) = \perp$ (the *null* choice) otherwise. By construction, this null choice will only occur out of equilibrium.

By construction, if agents reveal truthful information in response to queries then mechanism M will simulate strategy s^* in mechanism M' and terminate with outcome $g_K(g_{\mathcal{L}}(s(\theta))) = g_K(\mathit{inf}^{M'}(s^*(\theta))) = f(\theta)$, where $f(\theta)$ is constant for all θ in $\mathit{inf}^{M'}(s^*(\theta))$ because $f(\theta) = g(s^*(\theta))$ and by the definition of the information set, *inf*. Finally, to see that a truthful query strategy is an equilibrium in IRM M, assume agents $\neq i$ are truthful. Strategy s_i^* was a best-response for agent i in M' (including a consideration of the cost of refining its beliefs about its type), and agent i should respond truthfully so that its proxy agent plays the same strategy within the simulated mechanism. In pursuing its truthful strategy every agent will undertake the same cost in refining its type (and responding to elicitation) as was required to determine its strategy in mechanism M'. □

The proof of this incremental form of the revelation principle assumes that IRM M can simulate equilibrium bidding strategy s^* in M'. Here, we implicitly assume that the center in IRM M can solve the same *meta-level problem* that agents must solve to determine the equilibrium in M'. For instance, the mechanism needs information about the cost for queries and a model for the new information received in performing a query, so that it can make an appropriate tradeoff between this cost and the expected utility from being able to make a more informed decision about how to bid. The easiest way to justify this step is to consider an initial stage in which every agent reports its "meta type", i.e. information about its preference elicitation model, to the mechanism. This is valid under the following assumption.

Definition 3 (Meta-level assumption). It is costless for a bidder to provide her model for preference elicitation, for instance the cost to perform different types of queries and a distribution on the responses to each query, in an initial step of an incremental revelation mechanism.

We believe that the meta-level assumption is more justifiable than the classic assumption that the type information itself can be provided by a bidder at no cost. For instance, it seems reasonable that London bus operators can describe the process by which they will respond to queries about cost and provide an idea of: (a) the cost of this process; (b) a model for the likely response to each query.

2.4. An informational hierarchy

The formal definition of the *information revelation*, $inf^M(\theta)$, of mechanism M suggests a partial order across outcome-equivalent mechanisms. Mechanism M_1 *informationally-dominates* mechanism M_2 if the two mechanisms implement the same outcome in equilibrium, *and* $inf^{M_1}(\theta) \supseteq inf^{M_2}(\theta)$ for all $\theta \in \Theta$, and $inf^{M_1}(\theta') \supset inf^{M_2}(\theta')$ for some $\theta' \in \Theta$.

The following (trivial) lemma observes that DRMs can never informationally-dominate any other mechanism (although there can be settings in which they are worst-case optimal) [35].

Lemma 1. The information revelation in an incentive-compatible DRM, M, is $inf^M(\theta) = \theta$.

Proof. Trivial, since there is a bijection between messages sent by agents to the mechanism and the type of an agent. □

Similarly, the following lemma notes the correspondence between the definition of a *query rule* in an IRM and the information-revelation in the IRM.

Lemma 2. The information revelation in a truthful IRM is $inf^M(\theta) = g_{\mathcal{L}}(s^*(\theta))$, for truthful strategy profile s^*.

The extended example in section 2.6 illustrates that a second-price staged auction can implement an efficient allocation without complete information revelation, and thus informationally-dominates a second-price sealed-bid (Vickrey) auction.

2.5. First-best and second-best proxied and indirect mechanisms

The incremental revelation principle suggests a class of *proxied* and *indirect* mechanisms, with proxy agents that maintain partial information about bidder valuations and query bidders when the cost of elicitation is outweighed by the expected utility of bidding with improved value information. Indeed, *eBay*[3] auctions already require that bidders use proxy agents. eBay provides proxy agents that accept a *lower-bound* on a bidder's value for an item (also thought of as an *upper-bound* on the maximum that a proxy will bid for an item). Proxy agents bid in an ascending price auction, up to these bounds, and then prompt a bidder to refine her lower bound (upwards) whenever she is not the current winner.

The proxy mechanisms we have in mind are a little more sophisticated, in that they will also determine *how many times* to query a bidder, and on *which* items. This is not a problem in current eBay auctions (which are almost always for a single item), but will prove important in scaling agent-mediated electronic markets to settings with thousands

[3] www.ebay.com.

or millions of goods. Thus, a proxy agent's strategy in an IRM combines both *elicitation* and *bidding*. In defining this strategy it is important to distinguish between *first-best* and *second-best* models.

First best. The cost of elicitation, while non-negligible (so that elicitation is best avoided if possible) is small enough that a bidder will always be willing to provide information when it can lead to a more accurate bidding strategy. Proxy agents follow the same bidding strategy as they would follow with complete value information.

Second best. Elicitation is costly enough that a proxy must sometimes bid despite some residual uncertainty about the bidding strategy, because the additional cost of elicitation is not justified given the expected value of this information in improving the proxy's bid.

When proxy agents play first-best strategies the goal is to design auctions that implement a desired social choice function (e.g., allocative efficiency) with minimal elicitation. However, there will be many settings in which even minimal elicitation is too costly and the goal of auction design is instead to best *approximate* a desired social choice function given a model of costly elicitation. This is the second-best model. The incremental revelation principle holds in both first-best and second-best models.

In equilibrium, a second-best proxy strategy defines an elicitation action or bidding action for all possible states of a mechanism, and forms a best-response to the combined elicitation and bidding strategies of the other agents. Unfortunately, a full game-theoretic analysis for this second-best problem appears beyond the scope of current methods (either analytic or computational), even for simple auctions such as an ascending-price auction for a single item.[4] We will instead adopt heuristics to model second-best proxy strategies in the experimental analysis later in the paper.

2.6. Extended example

To understand the preference elicitation advantage that a proxied and indirect mechanism can enjoy over a direct mechanism consider a single-item allocation problem with three bidders, with values $v_1 = 4$, $v_2 = 8$, $v_3 = 12$, respectively. Suppose the bidders have determined the following bounds on their values:

Bidder 1: $v_1 \in [2, 7]$,
Bidder 2: $v_2 = 8$,
Bidder 3: $v_3 \in [11, 15]$.

Now, consider the following two auctions, that we will populate with proxy agents to convert into IRMs. We define the first-best bidding and elicitation strategy for a proxy

[4] Larson and Sandholm [24] have been able to perform some analysis of the combined deliberation and bidding strategy of agents in simple auctions, but stopped short of deriving a full equilibrium (except in a stylized setting with two deliberation actions and a sealed-bid auction [23]). Compte and Jehiel [7,8] and Rezende [42] derive equilibrium strategies, but only for simple models in which agents can refine beliefs at most once. In fact, computing equilibrium in infinite strategy games (such as auctions) is a topic of current AI research [41].

in each auction (i.e., with a proxy that performs queries until it can follow the bidding strategy that is in equilibrium with complete value information).

Second-price sealed-bid auction (Vickrey). Each agent must submit a single bid. The item will be sold to the highest bidder for the second-highest bid price.
First-best Proxy Strategy. Query the bidder for her exact value and then bid this value in the auction.

Staged second-price auction. An agent can increase its bid during the auction. The auction maintains an ask price equal to the current second-highest bid price and a new bid is valid as long as the bid price is at least the current ask price. The auction terminates when no new bids are received, with the item sold to the highest bidder for the second-highest bid price.
First-best Proxy Strategy. Whenever the agent is not the current winner, ask the bidder to refine her lower and upper bounds on value until: (a) the lower-bound is greater than the current ask price, or (b) the upper-bound is less than the current ask price. In case (a), the proxy agent will bid this new lower bound. In case (b), the proxy agent will leave the auction.[5]

The staged second-price auction can be formally mapped to the definition $M = (\mathcal{L}, \langle g_{\mathcal{L}}, g_K\rangle)$ of an IRM. Query language, $\mathcal{L}$, defines questions of the form: "refine bounds on value so that at least one of the following holds: (i) your lower-bound is above the current price; (ii) your upper-bound is below the current price". A valid query strategy must be consistent, requiring the bidder tighten her bounds during the auction. An example (truthful) query strategy, $s_i(\theta_i)$, might be "increase my lower bound to between price and value when price is less than value, and decrease my upper bound to between price and value when price is greater than value". The query rule, $g_{\mathcal{L}}$, defined in M can be constructed by supposing that in quiescence the mechanism selects a (losing) proxy at random and asks it to "bid or leave". At this point the proxy will query its bidder. New bids will advance the state of the auction, perhaps triggering further bids, until the auction returns to quiescence. This process will terminate when only one agent's lower-bound is above the price. Thus, the final type information revealed by bidders will be bounds $[\underline{v}, \overline{v}]$, s.t. there is a single proxy (representing the winner) with a lower bound above the upper bound of the other proxies, and another proxy (representing the second-highest bidder) with exact value information. Finally, the choice rule g_K will allocate the item to the bidder with the highest lower-bound at the final ask price.

In our example, the proxy agents have enough information to bid the price up to $8 + \varepsilon$ in the staged auction, for some small $\varepsilon > 0$. At this point, proxy agents acting for bidders 1 and 2 will drop out of the auction, leaving bidder 3 to win the item for \$8. Neither bidder 1 or bidder 3 needed to know, or report, their exact value to the proxy. In comparison, the proxy agents in the sealed-bid auction must perform additional queries

[5] Bidding at some price between the ask price and the bidder's true value whenever the ask price is less than the value, and leaving the auction otherwise, is an *ex post* Nash equilibrium (i.e. a best-response as long as other agents also follow this strategy, whatever their values). To see this notice that the direct-revelation form of the auction is strategically equivalent to the Vickrey auction.

to implement their first-best equilibrium strategy. Thus, we see that the staged second-price auction informationally-dominates the Vickrey auction.

Briefly, we can also consider how a second-best proxy strategy would differ from this first-best strategy in the second-price staged auction. A second-best proxy might simply stop bidding even though the current price is less than its upper bound (and without performing additional elicitation). Similarly, a second-best proxy might continue to bid without additional elicitation even though the current price is greater than its lower-bound. In both cases, this would suggest that the cost of an additional query is greater than the expected loss in utility from following a suboptimal bidding strategy. We see examples of this kind of analysis in the next section.

3. Modeling costly preference elicitation

To further illustrate the issue of costly preference elicitation and its impact on auction design we will adopt a simple model of preference elicitation. We consider proxy agents that sit between bidders and an auction, and must determine *when*, and for *which* item, to ask a bidder to refine its value. Our analysis considers both single item and multiple item allocation problems, and a variety of different IRM designs. For instance, we determine an optimal (second-best) elicitation and bidding strategy for a proxy agent in a single item posted price auction. We are forced to make some simplifying assumptions to approximate an optimal (second-best) proxy strategy in sealed bid and ascending price auctions.

We adopt the analysis of meta-deliberation proposed by Russell and Wefald [44] for resource-bounded computational agents. The key observation is that additional information about preferences can only be useful to a proxy if it is possible that the information might change an agent's bid. For costly preference elicitation, this requires computing the expected utility of an additional query and comparing this with the cost of elicitation. For budget-limited preference elicitation, at least with a myopic strategy, this suggests continued preference elicitation while within the budget and while the expected utility of additional information remains positive. An interesting effect of costly preference elicitation is that a proxy agent may accept a posted price even when the price is between its bounds on value in cases for which the cost of additional elicitation cannot be justified.

The auction designs are all fairly standard. We define posted-price (PP), ascending-price (AP) and sealed bid (SB) auctions, for both the single item and multi-item allocation problems. One novelty is in the method adopted in the ascending-price auctions to prompt agents to either submit a new bid or leave the auction. This ensures progress in the auction, and is a form of activity rule [31]. In addition, each auction is augmented with mandatory proxy agents that maintain bounds on bidder valuations. It is these proxy agents that make the auctions IRMs. The PP and AP auctions are designed to allow proxy agents to follow *first-best* bidding strategies without complete value information, and to enable effective second-best bidding strategies.

3.1. Model details

In our model, proxy agents maintain a *lower*- and *upper*-bound on a bidder's value for an item. In a single item allocation problem, we denote lower- and upper-bounds on an agent's value as $[\underline{v}_i, \overline{v}_i]$, and $\Delta_i = \overline{v}_i - \underline{v}_i$ denotes the current *uncertainty* about a bidder's true value. The true value is assumed uniformly distributed in this range. With multiple items $j \in G$, we denote the bounds on item j as $[\underline{v}_{ij}, \overline{v}_{ij}]$.

We consider queries of the form "*refine your lower- and upper-bounds*", with a parameterized model for the cost of a query and for its effect on value bounds. In the model, a query response reduces the current uncertainty by a multiplicative factor, $0 < \alpha < 1$, so that the new bounds are $\alpha\Delta_i$ apart, and the mean of the new bounds is uniformly distributed, $U(\underline{v}_i + \alpha^m\Delta_i/2, \overline{v}_i - \alpha^m\Delta_i/2)$, after m queries. The upper and lower limits on this distribution ensure that the new bounds are tighter than the current bounds. We refer to $(1-\alpha)$ as the *effectiveness* of the query. We consider two basic variations in formalizing the cost of elicitation to a bidder:

Costly. Each response to a query incurs a fixed cost, $C > 0$, to the bidder.
Budget-limited. Each bidder will perform a fixed number, $B > 0$, of queries.

The use of both lower- and upper-bounds can be considered a generalization of the lower-bounds on value used within eBay's proxy bidding system.[6] Upper-bound information provides negative evidence, and allows a proxy agent to better focus elicitation on items for which the bidder is competitive. Moreover, in settings such as logistics, in which auctions are used to determine transportation solutions for moving goods in a supply chain, local valuation problems can be formulated and solved as optimization problems [6,46]. Standard methods, such as A* search with admissible heuristics and linear programming-based branch-and-bound maintain lower and upper bounds on the value of the optimal solution.

Many aspects of costly preference elicitation only become apparent with multiple items, and we consider both single-item and multi-item allocation problems.[7] In the multi-item setting, $v_{ij} \geqslant 0$ denotes the value for an agent on item j, if this item is provided in isolation. The value for a bundle of goods is defined for two different models:

Additive. Each bidder has a valuation $v_i(S) = \sum_{j\in S} v_{ij}$ for each bundle $S \subseteq G$.
Unit-demand. Each bidder has a valuation $v_i(S) = \max_{j\in S} v_{ij}$ for each bundle $S \subseteq G$.

We first consider costly elicitation and single-item auctions, and then move to budget-limited elicitation and multi-item auctions.

[6] Recall that eBay users provide a lower-bound on value, up to which the eBay proxy will bid for the user in an ascending-price auction. Often people think of this as an "upper-bound" on the price a user will pay for the item but it corresponds with a lower bound in our framework.

[7] Specifically, with budget-limited elicitation an optimal proxy strategy in a single-item setting is to execute all queries, whatever the auction design. The elicitation problem is more interesting when the proxy must decide which items to query.

3.2. *Costly elicitation and single-item allocation*

We now define the auction rules, first-best and second-best proxy strategies for each auction.

3.2.1. Posted Price. **Choose a fixed price, p, and offer the price to each agent in turn as a take-it-or-leave-it offer, with agents selected in a random order. The item is sold to the first agent (if any) that accepts the price.**

First-best Proxy Strategy. Accept the price if $p < v_i$. Reject otherwise.

Second-best Proxy Strategy. First, we define the optimal bidding strategy for an agent with valuation bounds, $[\underline{v}_i, \overline{v}_i]$, that believes that the true value is uniformly distributed between these bounds and will do no further elicitation. We assume that agents are *risk neutral*, with utility $v_i - p$ for buying an item with value v_i at price p. Then, the agent should *accept* the item at price p if and only if $(\underline{v}_i + \overline{v}_i)/2 \geqslant p$.

Let T denote the current information a proxy has about the valuation (i.e. bounds on value), and let $T.Q^m$ denote the information after an additional m queries (of course, this will be a random variable before the queries are executed). Let b_T^* denote the optimal bid given state T, and let $b_{T.Q^m}^*$ denote the optimal bid after m queries. Let $u_i^T(b_T^*)$ denote the expected utility of the optimal bid given information T, computed as:

$$u_i^T(b_T^*) = \begin{cases} p - \dfrac{\underline{v}_i + \overline{v}_i}{2}, & \text{if } \dfrac{\underline{v}_i + \overline{v}_i}{2} \geqslant p, \\ 0, & \text{otherwise.} \end{cases}$$

Now, consider the expected utility from performing an additional m queries. An agent can only estimate this utility because there is uncertainty about how a bidder will respond to a query, and also about the bidder's true value. The estimate of this expected utility, $\hat{u}_i(Q^m)$, is defined with respect to the query model, as the estimated increase in utility from m queries minus the cost of elicitation:

$$\hat{u}_i\big(Q^m\big) = \hat{u}_i^{T.Q^m}(b_{T.Q^m}^*) - \hat{u}_i^{T.Q^m}(b_T^*) - mC,$$

where $\hat{u}_i^{T.Q^m}(b_{T.Q^m}^*)$ is the expected utility from the optimal bid given information $T.Q^m$, computed with respect to a model of the response of a bidder to the queries. The expected utility from the current bid, b_T^*, is also estimated with respect to the model of the proxy's information about the value for the item once it has performed m more steps of elicitation. As noted in Russell and Wefald [44], this is important to ensure that additional information is always evaluated with positive utility. Otherwise we could have "surprises" when we find that the value of the item is much less than we had assumed when evaluating the expected utility from the current decision given the current model T.

We derive an analytic expression for $\hat{u}_i(Q^m)$, in terms of elicitation parameters, α, C, price, p, and the current bounds on value. An optimal elicitation strategy for the proxy agent is to request additional value information whenever $\hat{u}_i(Q^m) > 0$ for some

$m > 0$. (Notice that we drop the so-called *meta-greedy* assumption in Russell and Wefald [44] and the proxy will consider the effect of m steps.)

Clearly, if $\overline{v}_i \leqslant p$ or $\underline{v}_i \geqslant p$ then there is no possible response to queries that will change the optimal bid and $b^*_{T.Q^m} = b^*_T$, for all Q^m, and $\hat{u}_i(Q^m) = -mC$. So, the interesting case is when $\underline{v}_i < p < \overline{v}_i$. In this case, the utility of elicitation depends on the difference between the price, p, and the mean, $\hat{v}_i = (\underline{v}_i + \overline{v})/2$ of the bounds. It is useful to define:

$$\gamma_i = \frac{2}{\Delta_i}\left|p - \hat{v}_i\right|.$$

This value of γ is between 0 and 1, because $\underline{v}_i < p < \overline{v}_i$. By symmetry, it is sufficient to consider the subcase, $\hat{v}_i \leqslant p < \overline{v}_i$. In this case, the current optimal bid is to *reject* the price, and elicitation is useful when it changes this decision, to *accept*.

Given our model of elicitation (defined in section 3.1), after a sequence of m queries the mean value is distributed uniformly, according to $U(\underline{v_i} + \alpha^m\Delta_i/2, \overline{v_i} - \alpha^m\Delta_i/2)$ where $(1-\alpha)$ is the elicitation effectiveness and $\Delta_i = \overline{v}_i - \underline{v}_i$. Define $m^*(\alpha, C)$ as the smallest number of queries for which $\overline{v}_i - \alpha^m\Delta_i/2 > p$. Solving, and with substitution $\overline{v}_i - p = \Delta_i - (p - \hat{v}_i)$, we have:

$$m^*(\alpha, \gamma) = \left\lceil \frac{\log(1-\gamma)}{\log(\alpha)} \right\rceil,$$

and because the current decision is to reject the price and $\hat{u}_i^{T.Q^m}(b^*_T) = 0$, then the expected utility from elicitation by $m \geqslant m^*(\alpha, \gamma)$ steps, is:

$$\begin{aligned}\hat{u}_i(b^*_{D^m}) &= \int_{\hat{v}_i = p}^{\overline{v}_i - \alpha^m\Delta/2} \frac{(\hat{v}_i - p)}{\Delta_i(1-\alpha^m)}\,\mathrm{d}\hat{v}_i - 0 - mC\\ &= \frac{(\overline{v}_i - \alpha^m\Delta_i/2 - p)^2}{2\Delta_i(1-\alpha^m)} - mC = \frac{\Delta_i(1-\gamma-\alpha^m)^2}{8(1-\alpha^m)} - mC.\end{aligned}$$

The limits on the integral arise because the expected value after m queries is distributed according to $U(\underline{v}_i + \alpha^m\Delta_i/2, \overline{v}_i - \alpha^m\Delta_i/2)$, and the agent will accept the price if the new mean is greater than the price.

Putting everything together, the expected utility to the proxy agent for an additional m queries, is defined as:

$$\hat{u}_i(D^m) = \begin{cases} \dfrac{\Delta_i(1-\gamma-\alpha^m)^2}{8(1-\alpha^m)} - mC, & \text{if } m \geqslant m^*(\alpha, \gamma),\\ -mC, & \text{otherwise.}\end{cases}$$

A quick analysis of the comparative statics of this expression shows that $\partial\hat{u}_i(D^m)/\partial\gamma < f0$, $\partial\hat{u}_i(D^m)/\partial\Delta > 0$, when $p = \hat{v}_i$, $\partial\hat{u}_i(D^m)/\partial\alpha < 0$, when $p = \hat{v}_i$, and $\partial\hat{u}_i(D^m)/\partial C < 0$, i.e. preference elicitation is more useful as the price gets closer to an agent's expected value (because it is more likely to change an agent's bid), as the uncertainty increases, as effectiveness increases, and as cost decreases.

For a particular uncertainty, Δ_i, and (α, C), it can be useful to compute *elicitation bounds*. These bounds, $[\underline{d}_i, \overline{d}_i]$, define the range on prices (perhaps empty) for which the proxy should elicit additional value information before making a decision about whether or not to buy the item. The bounds are centered on an agent's current estimated value, $\hat{v}_i$, and computed in terms of the smallest Δ_i^* for which $\hat{u}_i(D^m) > 0$ for some $m > 0$. When no such Δ exists, then $\overline{d}_i = \underline{d}_i = \hat{v}_i$ and the proxy should choose not to elicit any information for *any* ask price. Otherwise, $\underline{d}_i = \hat{v}_i + \Delta_i^*/2$ and $\underline{d}_i = \hat{v}_i - \Delta_i^*/2$.

Example. Consider a bidder with $\alpha = 0.5$ and $C = 0.05$, and a proxy facing a price, $p = 5$ and with value bounds (a) [4, 7], (b) [3, 8], and (c) [4.5, 5.3]. The expected value is 5.5 in case (a) and case (b), and the elicitation bounds are [5.05, 5.95] and [4.5, 6.5], respectively. The proxy should accept the price without elicitation in case (a), but ask for additional information in case (b). The elicitation bounds are [4.9, 4.9] in case (c), indicating that there is no price for which the proxy should request additional information. Given its expected value of 4.9, the proxy rejects the price.

In summary, the combined elicitation and bidding strategy of a proxy agent in the posted price auction, is to execute preference elicitation queries while the price is within elicitation bounds $[\underline{d}_i, \overline{d}_i]$, and finally buy the item if the price is less than $\hat{v}_i$ and reject the offer otherwise.

3.2.2. Ascending Price. **The auction maintains an ask price, equal to the $\varepsilon > 0$ above the highest bid. The highest bidder is the provisional winner, at its bid price. A list of *active* bidders is maintained, with all agents initially active. In a period of quiescence, with no bids received for a fixed period of time, then the auction picks (at random) one of the agents not winning and asks the agent to improve its bid or leave the auction. The auction terminates when there is only one active agent, which wins at its final bid price.**

First-best Proxy Strategy. Bid while the ask price is less than value, v_i, and not winning. Leave the auction when the ask price is greater than value.

Second-best Proxy Strategy. The preference elicitation problem facing a proxy agent in an ascending price auction is more difficult than in a posted price auction because of price dynamics. For example, while it might be worthwhile to elicit additional information before bidding at the current ask price, if a proxy agent believes that the price will increase further – perhaps to above its upper-bound on value – then the proxy would be better not to query the bidder. A further complication is that a model of price dynamics requires a model of the other bidders.

It is beyond the reach of current game-theoretic analysis to determine the equilibrium, for both elicitation and bidding, in this setting. Instead, we define a heuristic second-best proxy strategy. We assume that proxy agents will determine whether or not to execute queries only when prompted by the auction to bid or leave the auction. Moreover, we assume that the proxy agents act myopically and assume that the current price faced in the auction will be the final price.

To be precise, we make the following assumptions in modeling proxy agent behavior:

A1 Proxy agents only elicit value information from bidders when prompted by the auction to submit a new bid or leave. At that time, proxy agents are myopic and assume the current price is fixed.

A2 Proxy agents will bid whenever the price is less than their lower elicitation bound $\underline{d}_i$.

A3 Proxy agents will leave the auction whenever the price is above their upper elicitation bound, $\overline{d}_i$.

We deviate from a game-theoretic analysis in assumptions (A1) and (A2). The agent bids when it knows it would rather take the current price than query the bidder (A2), even when this might help another agent. The agent queries the bidder when the current price is such that its expected utility from querying is positive (A1), even though prices might increase further. Assumption (A3) is justified because the ask price is *increasing* and will always be greater than the upper elicitation bound. As a special case, when the elicitation bounds $\underline{d}_i = \overline{d}_i$ then the proxy agent will bid while the price is less than its belief about expected value, and leave the auction once the price is greater than its expected value.

3.2.3. Sealed Bid. **Each agent submits a sealed bid, and the item is sold to the highest bidder for the second-highest bid price.**

First-best Proxy Strategy. Bid true value, v_i.

Second-best Proxy Strategy. In the sealed-bid auction, the amount of preference elicitation that a proxy agent should perform depends on the distribution over second price that the proxy will face, which in turn depends on the bids by other agents, which in turn depends on the amount of preference elicitation performed by other agents. Again, we have a complex game-theoretic problem, beyond the limits of current game-theoretic analysis.

We choose to compute an upper-bound on the number of queries performed in a (symmetric) pure Nash equilibrium. We assume that each proxy agent will execute the same number of queries, q, and use a computational method to determine the *maximal* number of queries, $q_{\max}$, that can be in an economy with positive expected utility to each agent. Monte Carlo analysis is used to determine $q_{\max}$. For each level of q, we:

Do β times:
- sample the value distribution;
- perform q queries of each bidder;
- determine the final set of agent bounds;
- determine the outcome of the auction and expected utility.

The number of queries, $q_{\max}$, depends on the number of agents, the query effectiveness $(1-\alpha)$, and the query cost, C. The combined preference elicitation and bidding strategy of a proxy agent in the sealed-bid auction is to execute $q_{\max}$ queries and submit a bid equal to its final belief about the bidder's expected value.

3.3. Budget-limited elicitation and additive values

Budget-limited preference elicitation is interesting in problems with multiple items because the proxy must determine when to query the bidder, and which item(s) to query the bidder about. Again, we describe a posted price, ascending, and sealed-bid auction, and derive heuristic second-best proxy strategies. We assume a lower-bound on uncertainty, $\Delta_{\min} > 0$, that is assumed to correspond with *exact* information. When the proxy has bounds with uncertainty $\Delta_{ij} \leqslant \Delta_{\min}$ the proxy considers this exact information and has no value for continued queries on the value of that item.

3.3.1. Posted Price. **Choose a fixed price, p. Run the single-item posted-price auction for each item in parallel, with the same price p in each auction and agents selected in a different (random) order for each item.**

First-best Proxy Strategy. When offered an item, j, buy the item if $p < v_{ij}$ and reject the offer otherwise.

Second-best Proxy Strategy. The bidding strategy, given bounds on the value of items and no further elicitation is to buy any item offered that has price below the proxy's belief about the expected value of the item. We model a proxy's elicitation strategy as follows:

1. Maintain a set of items with uncertainty, $\Delta_{ij} > \Delta_{\min}$, and for which the price is between the valuation bounds.
2. While this set is non-empty and the number of queries executed is less than the elicitation budget, B, choose an item at random from the set and perform a query.

3.3.2. Ascending Price. **Run the ascending price auction for each item separately, with each auction running in parallel and closing at the same time, when there is only one active agent in each auction.**

First-best Proxy Strategy. For each item, bid while the ask price on that item is less than value and while the agent is not winning. Leave the auction on that item when the price is greater than value.

Second-best Proxy Strategy. The bidding strategy, given bounds on valuation and no further elicitation is to bid on any item with price less than the expected value that the proxy is not currently winning. In addition, the proxy will always bid on an item that is priced below $\underline{v}_{ij}$. As a heuristic elicitation strategy, we model a proxy that only executes elicitation queries on an item when prompted by that auction to submit a new bid or leave. At this point the proxy acts as though the price is a posted price and performs a query on this item until the current price is outside of its valuation bounds (or until $\Delta_{ij} \leqslant \Delta_{\min}$). When its budget is exhausted the proxy adopts expected values $\hat{v}_{ij}$ to guide its bidding strategy.

3.3.3. Sealed Bid. **Run the single-item sealed-bid auction for each item in parallel, with each auction closing at the same time.**

First-best Proxy Strategy. For each item j, bid v_{ij} in the corresponding auction.

Second-best Proxy Strategy. The optimal bidding strategy, given no further elicitation, is to bid the expected value on each item. Notice that a query on an item j with $\Delta_{ij} > \Delta_{\min}$ has positive expected utility during the elicitation phase because it is possible that the highest bid from another agent will have a value between the item's current bounds. With this in mind, we model a proxy agent's elicitation strategy as follows:

1. Maintain a set of items with uncertainty, $\Delta_{ij} > \Delta_{\min}$.
2. While this set is non-empty and the number of queries executed is less than elicitation budget, B, choose an item at random from the set and perform a query.

Once an agent has exhausted its elicitation budget, or the value is known on all items, then it bids the expected value for each item.

3.4. Budget-limited elicitation and unit demand

Finally, we consider the multi-item allocation problem with unit-demand valuations, such that each bidder wants at most one item.

3.4.1. Posted Price. **Choose a fixed price, p, and maintain a set of unsold items, $G' \subseteq G$, that initially contains all the items. Consider the agents in some random sequence. Make at take-it-or-leave-it offer for every item still in set G' to each agent in turn. The agent can purchase any subset of the remaining items, each at price p. Remove the items as they are sold.**

First-best Proxy Strategy. Bid for the item with $\max_{j \in G'}\{v_{ij}\}$, as long as $\max_{j \in G'}\{v_{ij}\} > p$.

Second-best Proxy Strategy. The bidding strategy, given bounds on the value of items and no further elicitation is to buy the item with the greatest *expected* value (as long as the expected value is greater than the price). We model a proxy agent's elicitation strategy as:

1. Maintain the set of items in G' with uncertain values, $\Delta_{ij} > \Delta_{\min}$, that are still undominated by the other items, i.e. for which $\overline{v}_{ij} \geqslant \underline{v}_{ij'}$ for all $j' \neq j$, and for which $\overline{v}_{ij} > p$.

2. While this set contains more than one item, and the number of queries executed is less than B, choose an item at random from the set and perform a query.
3. When this set contains a single item, continue to refine the bounds on this item until its valuation is known ($\Delta_{ij} \leqslant \Delta_{\min}$), or there are no queries left to execute.

3.4.2. Ascending Price. **Run the ascending price auction for each item simultaneously, but with a single *active* status for each agent across all auctions. Agents can bid in multiple auctions. Only when *every* auction is in quiescence will the auction poll an agent and ask it to submit a bid. The auction looks for an agent that is not winning an item in *any* auction, and asks this agent to bid in at least one auction, or leave *all* auctions. Terminate when every active agent is winning an item in one or more auctions.**

First-best Proxy Strategy. While not winning in any auction, bid for the item with $\max_{j\in G}\{v_{ij}-p_j\}$, where p_j is the current ask price on item j, until $\max_{j\in G}\{v_{ij}-p_j\} \leqslant 0$ (at which point leave).

Second-best Proxy Strategy. The optimal bidding strategy, once the proxy will perform no further elicitation, is to bid for the item with the greatest expected utility at the current price (while the proxy is not winning on an item). Before its elicitation budget is exhausted, we model the proxy agent as executing queries only when prompted by the auction to "submit a new bid or leave". At this point, the proxy acts as though the prices in each auction are posted price and follows the elicitation strategy for the posted-price auction.

As a bidding strategy, the proxy will always bid while an item's lower-bound utility (i.e. lower-bound value minus price) dominates the upper-bound utility of the other items. The proxy adopts the expected values to guide its bidding when its elicitation budget is exhausted.

3.4.3. Sealed Bid. **Provide an *exclusive-or* bidding language [34], so that an agent can bid its value for each item but also state that it is interested in winning at most one item. Implement a Generalized Vickrey Auction (GVA) to provide strategyproofness for informed proxy agents.**[8]

First-best Proxy Strategy. Bid the true value for each item, together with an exclusive-or constraint to state that the agent will buy at most one item.

Second-best Proxy Strategy. Notice that a query on an item with uncertain information ($\Delta_{ij} > \Delta_{\min}$) has positive expected utility while it remains possible that more accurate

[8] The allocation is computed as a solution to the assignment problem, $\max_{x_{ij}} \sum_i \sum_j b_{ij}x_{ij}$ s.t. $\sum_i x_{ij} \leqslant 1$ for all j and $\sum_j x_{ij} \leqslant 1$ for all i, and $x_{ij} \in \{0,1\}$, where b_{ij} is the bid from agent i for item j. Given solution x^*_{ij}, then agent i with $x^*_{ij'} = 1$ for item j' is allocated the item, and makes payment $b_{ij'} - (V(I) - V(I \setminus i))$, where $V(I)$ is the revenue from the assignment problem with all agents and $V(I \setminus i)$ is the revenue from the optimal solution to the assignment problem with agent i removed.

value information will change the allocation decision, given the bids from other agents. With this in mind, we model a proxy agent's elicitation strategy as follows:

1. Maintain a set of items with uncertainty, $\Delta_{ij} > \Delta_{\min}$.
2. While this set is non-empty and the number of queries executed is less than elicitation budget, B, choose an item at random from the set and perform a query.

Once an agent has exhausted its elicitation budget, or the value is known on all items, then it submits as its bid the expected value for each item.[9]

4. Experimental results: Single-item auction

In this section, we compare the allocative efficiency of the PP, AP, and SB auctions in a single-item allocation problem and with the costly model of preference elicitation. The experimental analysis adopts the preference elicitation and bidding strategies outlined in the previous section.

Given allocation $x = (x_1, \dots, x_N)$, with $x_i \in \{0, 1\}$ to denote whether agent i gets the item, and $\sum_i x_i \leqslant 1$, the allocative efficiency is computed as $\sum_i x_i v_i / \max_i \{v_i\}$. We compute the average allocative efficiency of each auction across a number of runs. In general, the results show that the AP auction outperforms the PP auction which outperforms the SB auction. The AP auction is better able to control elicitation, with more queries asked of bidders with a high value for the item than of bidders with a low value. Feedback in the auction, via prices, is used to focus elicitation and promote better decisions.

4.1. Experimental set-up

We consider bidders with independent, identically distributed, values drawn from a Uniform distribution, $U(0, 10)$. Every agent is initially ignorant of the bidder's value, with initial valuation bounds [0, 10]. We consider four pairs of preference-elicitation parameters (α, C): [1] (0.7, 0.5); [2] (0.3, 0.5); [3] (0.7, 0.05); and [4] (0.3, 0.05). The initial elicitation bounds (see section 3.2.1) are [5, 5], [4.2, 5.8], [2.9, 7.1], and [1.8, 8.2], for [1], [2], [3] and [4], respectively. Moving from [1] to [4] preference elicitation gets more effective, and less costly.

In simulation we refine bounds to ensure that the bounds are tighter than the current bounds, and to keep the true value uniformly distributed between the bounds with respect to the distribution over possible sequences of refinements. This requires biasing the new range of values towards mean values that are weighted in favor of values towards the *edges* of the current range, because the Uniform distribution is not conjugate with itself,

[9] This is the optimal bidding strategy because the bids from other agents set a price on each item, and the rules of the GVA will purchase the item with greatest reported surplus given the agent's bid, and given these prices. Thus, this maximizes expected utility.

and a Uniform distribution to generate the means of the new range does not generate a final value that is uniformly distributed given the initial bounds.[10]

The price in the posted-price auction is selected (in simulation) to maximize the average revenue to the seller, given the preference elicitation and bidding strategies followed by agents. This is intended to provide a best-case measure of the performance of the PP auction, for a well-informed seller.[11] The posted price increases with the number of agents, and as elicitation becomes cheaper and more effective. For costly or ineffective elicitation the seller should set a price close to the initial expected value of the buyers, to encourage agents to elicit costly value information and also to provide additional surplus to the buyers (and justify additional queries).

The myopic assumption made by the proxy agents when making preference elicitation decisions in the ascending-price auction can *overestimate* the value of queries. For this reason, we set the bid increment, ε, to the minimal increment that leaves bidders with non-negative expected utility from participation.[12] This ensures that any performance benefits that occur in the ascending-price auction do not occur because agents are performing more preference elicitation than can be sustained in the auction given bidder utility.

In the sealed-bid auction, we compute the number of queries as the maximal number of queries that can be sustained with non-negative expected utility in equilibrium.[13] The auction can sustain less preference elicitation as the number of participants increases because there is more competition and the winning agent pays a higher price and receives less surplus. For some settings of (α, C) and for some numbers of agents, *no* elicitation can be supported in the symmetric equilibrium.

4.2. Results: Allocative efficiency

Figure 1 plots allocative efficiency in SB as the number of agents increases for parameters [1], [2], [3] and [4], averaged over 500 trials. SB performs well for small numbers of bidders and with effective or cheap elicitation, but fails with ineffective and costly elicitation, even with small numbers of bidders. In these cases, the expected surplus from participation is less than the cost of performing even a single query to each

[10] Fong [13] fixes this requirement by considering a Gaussian distribution for valuation, which is conjugate with itself.

[11] The optimal price, for (5, 10, 20, 30, 40, 50, 100) agents, is (4.9, 4.9, 4.9, 4.9, 4.9, 4.9, 4.9), (5.8, 5.8, 5.8, 5.8, 5.8, 5.8, 5.8), (6.6, 6.6, 7.0, 7.1, 7.1, 7.1, 7.1), and (6.8, 7.8, 8.2, 8.2, 8.2, 8.2, 8.2), for elicitation parameters [1], [2], [3] and [4], respectively. The optimal ask prices for agents with exact information about values are (6.8, 7.8, 8.6, 8.8, 9, 9.3, 9.5).

[12] The bid-increment, for (5, 10, 20, 30, 40, 50, 100) agents, is set to (1, 1, 1, 1, 1, 1, 1), (0.7, 1, 1, 1, 1, 1, 1), (0.2, 0.2, 0.4, 0.4, 0.5, 0.5, 0.5), and (0.2, 0.2, 0.6, 0.6, 0.6, 0.6, 0.6), for elicitation parameters [1], [2], [3] and [4], respectively.

[13] The number of queries, for (5, 10, 20, 30, 40, 50, 100) agents, are (1, 1, 0, 0, 0, 0, 0, 0, 0, 0), (1, 1, 0, 0, 0, 0, 0, 0, 0, 0), (17, 10, 7, 3, 2, 1, 1, 1, 0, 0), and (16, 10, 6, 2, 1, 0, 0, 0, 0, 0), for elicitation parameters [1], [2], [3] and [4], respectively.

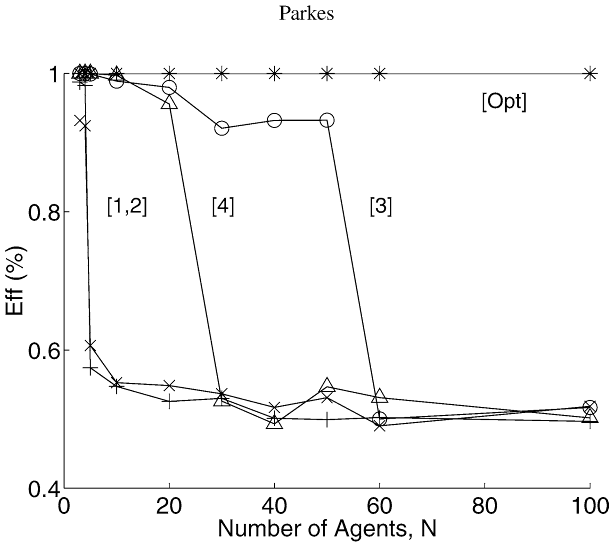

Figure 1. Efficiency in SB for elicitation parameters [1], [2], [3] and [4]. The auction is efficient with agents that have exact value information [Opt].

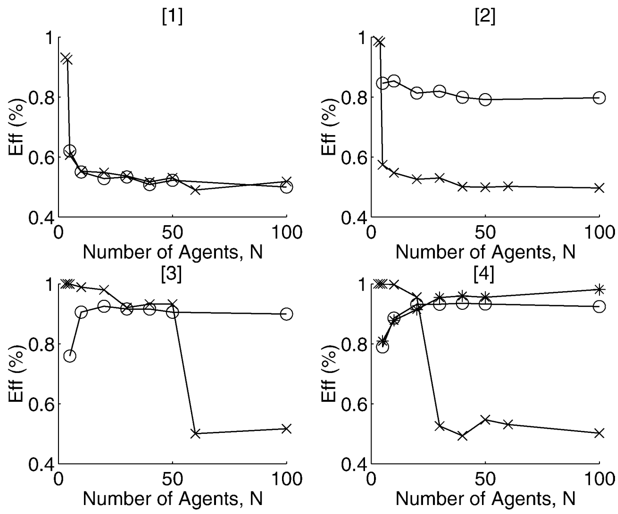

Figure 2. Efficiency in PP ('$\circ$') and SB ('$\times$') auctions, for elicitation parameters [1] to [4]. For comparison, we also plot, in subplot [4], the efficiency of PP with agents that have exact value information (line '$*$').

bidder. For instance, the SB auction cannot sustain an efficient market with 10 or more bidders in elicitation models [1] and [2].

Figure 2 compares the efficiency of SB and PP, for different elicitation parameters. For comparison, we also plot in subplot [4] (line '$*$') the efficiency of PP with agents that have exact value information. All results are averaged over 1000 trials. In general, the PP auction performs better with larger numbers of agents because there is less

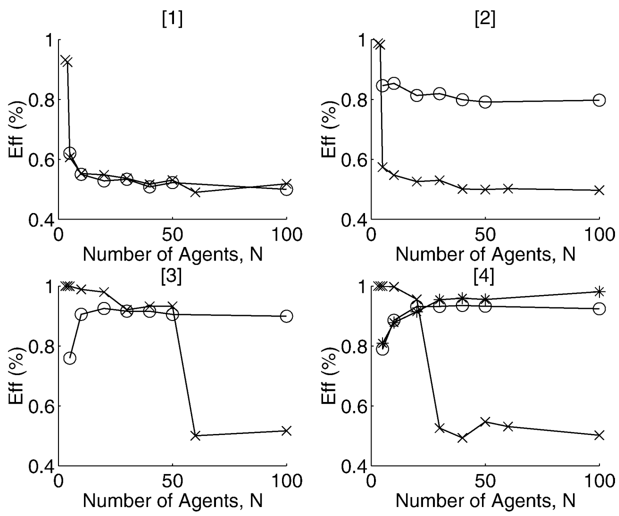

Figure 3. Efficiency in AP ('+') and PP ('∘') auctions, for elicitation parameters [1] to [4].

variance in the maximum value across the agents. However, notice that with costly preference elicitation the efficiency does not approach 100% even as the number of agents gets large. This is because the seller cannot set the initial price above the initial upper elicitation bound, which is 8.2 with parameters [4] (see section 3.2.1). In comparison, the seller in the auction with 100 informed agents can set the price at 9.5 to maximize revenue.

Nevertheless, the PP auction outperforms the SB auction when there are many agents. The posted price, coupled with the take-it-or-leave-it offer that isolates the decision making of one agent from the other agents, allows a proxy agent to follow an optimal elicitation strategy, and without concern as to the strategies of other agents. This leads to the reversal of the comparative efficiency under a classic economic model, in which one would expect the SB auction to be efficient and to outperform the PP auction.

Figure 3 compares the efficiency of AP with PP for different elicitation parameters. These results are averaged over 200 trials. AP matches the efficiency of SB for small numbers of agents, and matches the efficiency of PP for large numbers of agents. Unlike SB, the performance of AP holds up as the number of agents increases because queries are implemented in sequence, with information propagated via prices before other agents execute queries. Only a *single* proxy is called on to perform elicitation and improve its bid when the auction reaches quiescence. Also, we see that AP outperforms PP because it is able to set a higher final price than PP through competition between proxy agents that discover high values through incremental elicitation. Thus, AP seems to dominate both SB and PP auctions.

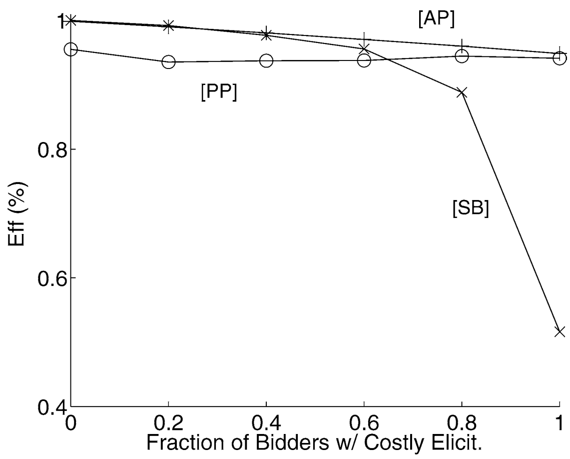

Figure 4. Efficiency of SB, PP and AP, for 30 agents and elicitation parameters [4] as the proportion of bidders with costly preference elicitation increases.

4.3. A mixture of informed and uninformed bidders

It is interesting to consider a mixture of bidders, some with exact value information and costless elicitation, and some with costly preference elicitation problems. We assume a fraction, λ, of bidders with costly elicitation. Figure 4 shows the efficiency of SB, PP and AP,[14] with 30 bidders, and elicitation parameters [4], as the fraction of bidders with costly elicitation is increased from 0 to 100%.

None of the proxies for bidders with costly elicitation choose to execute any queries in SB because the auction is too competitive and the feedback not rich enough to guide elicitation. Thus, the efficiency in SB is due to the bids from the proxy agents with exact value information and costless elicitation (this explains why performance falls as λ approaches 1). Overall, we see that the efficiency of AP dominates both PP and SB, matching that of SB for costless elicitation and remaining efficient as the proportion of bidders with costly elicitation increases.

4.4. A closer look: Distribution of queries

It is interesting to compare the number, and distribution, of queries performed by proxy agents across the different auctions. We compare this distribution with that of the *optimal* elicitation strategy, that asks the minimal number of queries that is required to determine the efficient allocation, given the benefit of hindsight.[15] Bidders must be queried until there is one bidder with a lower bound that is greater than the upper bound of all other bidders. We find that the AP auction allows proxy agents to ask more queries

[14] The bid increments in AP were set to (0.2, 0.3, 0.4, 0.4, 0.4, 0.6) for fractions, (0, 0.2, 0.4, 0.6, 0.8, 1.0), to provide positive expected utility to agents. The ask price in PP was set to (8.8, 8.6, 8.2, 8.2, 8.2, 8.2) for fractions, (0, 0.2, 0.4, 0.6, 0.8, 1.0), to maximize the expected revenue.

[15] This is the nondeterministic verification setting in Nisan and Segal [35].

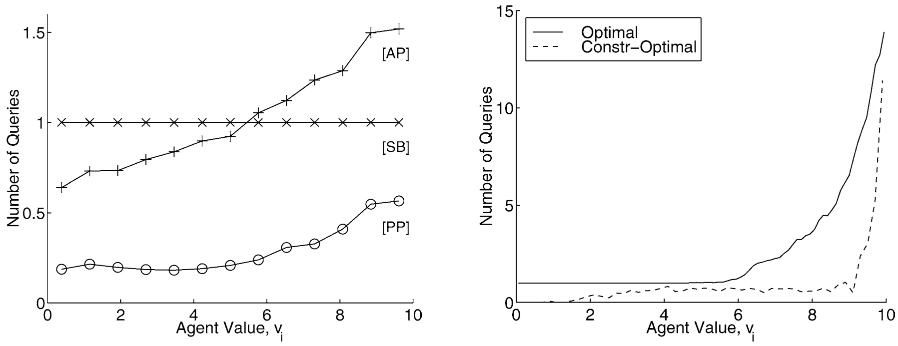

(a) (b)

Figure 5. Query distribution for 30 agents and elicitation parameters [4]. (a) Queries in the AP, SB and PP auctions. (b) Queries with optimal preference elicitation with hindsight, both with and without constraints on the total number of queries.

of bidders with a high value for the item than bidders with a low value, and that the query strategy that emerges in AP shares similar characteristics with that of the optimal strategy.

Figure 5(a) plots the average number of queries executed by agents against the true value of the bidder for the item, for the PP, SB, and AP auctions. These results are averaged over 500 trials, and for elicitation parameters [4]. For this problem, the efficiency of AP is 96.1%, compared to 92.1% for SB and 91.6% for PP. The agents in SB all execute a single query. The agents in AP all execute a single query on average, but 53.4% do not execute any query, including 51.8% of agents with value between 8 and 10.

What is particularly striking is that bidders with a higher value in AP tend to receive more queries than bidders with low values. This is useful, because these are the bidders that matter in determining the efficient allocation. In comparison, the agents in PP executed an average of only 0.29 queries each, and 81.2% of the agents execute no queries, including 80.9% of agents with value between 8 and 10. The problem in PP is that the seller cannot set the initial price above 7.1 (the initial upper elicitation bound), and this price means that the item is sold quite quickly and without being offered to enough high value agents.

Figure 5(b) plots the queries performed in the optimal and constrained-optimal elicitation strategy for 30 agents and elicitation effectiveness as in model [4]. The distribution for the constrained-optimal case is limited to an average of 1 query per agent to match the average number of queries performed in AP. For a particular instance, we compute the optimal number of queries to each bidder by searching across all policies that assign an increasing number of queries to bidders with higher values (this set of policies completely characterizes the set of optimal solutions). The same basic procedure is used to find the constrained-optimal strategy, except that the search is now limited to a finite number of queries.

Bidders with high value receive more queries than bidders with low value in the optimal distribution. For instance, although the average number of queries per-agent is 2.6, the three bidders with the highest values receive an average of 5.8 queries while the bidder with the highest value receives an average of 10.2 queries. To determine the efficient allocation it is more important to have precise information on the bidders with high values, that are in contention for the item, than the bidders with low values that are out of contention. On comparison with the distribution in the AP auction, we can infer that the AP auction is effective because it provides enough information to allow proxies to execute more queries to bidders with high values, which in turn boosts the efficiency of the auction. As the price increases, the agents with high values are incented to continue to refine their information while the agents with low values can leave the auction.

5. Experimental results: Multi-item auctions

In this section, we compare the performance of the PP, SB, and AP auctions in the multi-item allocation problems, considering both linear-additive and unit-demand agent valuations. We adopt the budget-limited model of preference elicitation, and consider the efficiency of auctions with proxy agents that can ask only a finite number of queries of their bidders. In this setting, good auction design can allow proxy agents to make better decisions about which item to query.

5.1. Experimental set-up

In all problems we consider elicitation parameter, $\alpha = 0.7$, and each query costs one unit of a total budget limit B, which is varied from 1 to the total number of queries required for exact information about a bidder's valuation. We consider problems with 10 items and define the value on an item in isolation, v_{ij}, as identically and independently distributed across bidders according to $U(0, 10)$. The proxy agents are initially completely uncertain about a bidder's value for each item, with initial bounds [0, 10]. We set the minimal uncertainty, $\Delta_{\min} = 0.5$, and agents view the value information as *exact* for bounds closer than this. Given this, the number of queries required to compute an exact value for a single item, denoted C_{exact}, is 9. Given allocation $S = (S_1, \ldots, S_N)$, with $S_i \subseteq G$, efficiency is defined as $\sum_i v_i(S_i) / \max_{S'}\{\sum_i v_i(S'_i)\}$, for feasible allocations S', with values as defined in section 3.1.

5.2. Additive values: Efficiency and correctness

Figure 6 illustrates the efficiency and correctness for the additive-value allocation problem with 10 items and 20 agents, averaged over 30 trials.[16] Figure 6(a) plots efficiency against elicitation budget, with the elicitation budget normalized with respect to

[16] The price in PP is set to 8.4, and the minimal bid increment in AP is set to $\varepsilon = 0.1$, to maximize allocative efficiency over a range of budget limits.

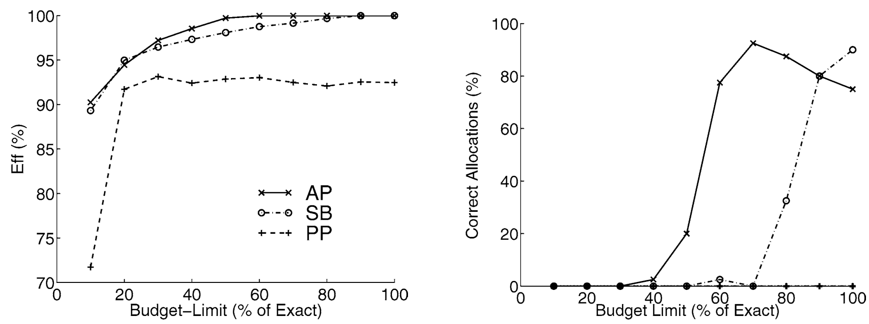

(a) Allocative Efficiency

(b) Correctness

Figure 6. Additive-value problem and budget-limited queries: efficiency and correctness as the elicitation budget is increased.

the number of queries required for exact value information (i.e. 90 queries). Figure 6(b) plots the percentage of correct allocations, and provides a more nuanced view in cases in which both the SB and AP auctions have reasonable allocative efficiency. The PP auction never produces a correct allocation.

We see that the AP auction performs better than the SB auction for agents with the same computation budget, at least until the agents have almost exact information about preferences. The AP auction uses the same amount of query information more effectively, as prices guide agents as to *which* items to query.[17] The performance of the PP auction is dominated by the AP and SB auctions, largely because it is difficult to set a single price for every item that will support an efficient allocation.

Figure 7 plots the allocative efficiency against the number of queries performed by proxy agents in each auction, with the data generated by varying the elicitation budget between 0 and $10C_{\max}$.[18] The proxy agents in AP will often not use the full elicitation budget, when the bidding strategy is well-defined with partial information. For instance, only 31% of the complete elicitation queries are required for allocative efficiency in AP with 5 agents, and only 49% is required with 20 agents. In comparison, the SB auction needs almost 100% of the complete information queries to achieve allocative efficiency. The PP auction is able to outperform the AP auction when the elicitation budget is very small (less than 10% of $C_{\max}$) and as the number of bidders increases (i.e. with 20 but not with 5 bidders). However, the PP auction is limited in its ability to achieve high efficiency (above around 93%) because it lacks dynamic pricing, and the AP auction dominates for all elicitation budgets with smaller numbers of bidders.

[17] Notice that although the AP and SB auctions approach 100% efficiency as the budget-limit approaches C_{exact} the number of correct allocations remains less than 100%. This is a result of the residual error inherent in adopting an acceptable uncertainty level of $\Delta_{\min}$.

[18] Results are averaged over 80 trials with 5 agents and 20 trials with 20 agents, the posted price is 5.6 and 8.8 for 5 and 20 agents, respectively.

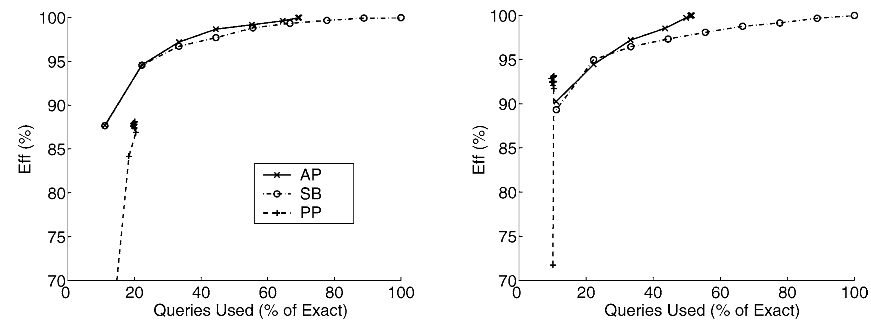

(a) 5 agents (b) 20 agents

Figure 7. Additive-value problem: efficiency vs. queries executed.

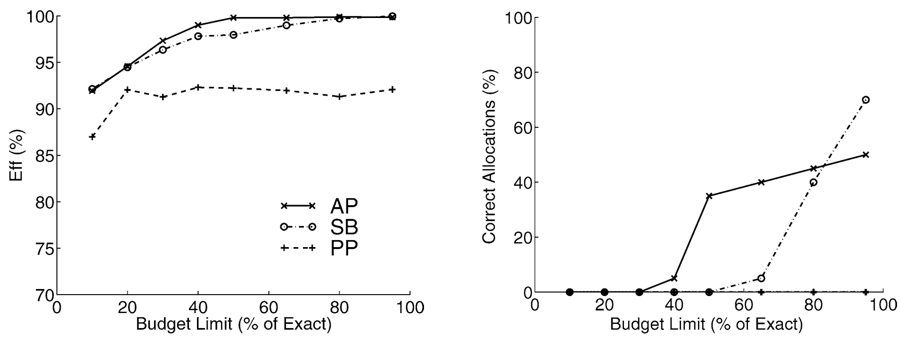

(a) Allocative Efficiency (b) Correctness

Figure 8. Unit-demand problem and budget-limited queries: efficiency and correctness as the elicitation budget is increased.

5.3. Unit-demand: Efficiency and correctness

Figure 8 illustrates the efficiency and correctness of each auction in the unit-demand allocation problem, with 20 agents, and as the elicitation budget varies between 0 and $10C_{\max}$.[19] The auction properties are very similar to the properties in the multi-item allocation problem with additive values. The AP auction is more efficient than the SB auction except at very low or very high budget limits, and agents in the AP auction can compute optimal bidding strategies (and the auction is efficient) with only 51% of the queries required for complete information.

5.4. Discussion

The proxied AP auction is more efficient than the SB and PP auctions for a budget-limited number of queries. The information feedback provided by prices in AP allows better decisions about how to use each query. In fact, we can determine the efficient

[19] Results are averaged over 20 trials and the price in PP set to 6.5, and the bid increment in AP to $\varepsilon = 0.1$.

allocation in AP with around 50% of the number of queries that are required to reveal full information on bidder preferences.

6. Related work

6.1. Mechanism design with bounded-rational agents

An earlier study of Vickrey auctions in computational settings demonstrated that the strategyproofness of an auction can break when agents have approximate values for items and options to continue computation or submit bids [47]. Sandholm's analysis shows that an agent can make a better decision about whether or not to perform further computation about the value of an item when it is well informed about the bids from other agents. Larson and Sandholm [22,24] have modeled agent deliberation in situations of strategic interdependence, and in particular when agents must make explicit decisions about whether to deliberate about their own values or the values of other agents. The authors model a *deliberation equilibrium*, including both deliberation actions and base-level (strategic) actions and are able to state a number of facts about the equilibrium. However, they are unable to derive the full game-theoretic equilibrium, and their work places less emphasis on allocative efficiency. Recent work adopts computational methods to determine an equilibrium for a sealed-bid auction and agents considering a sequence of two deliberation actions [23].

Contemporaneously with this work, Compte and Jehiel [7,8] and Rezende [42] have proposed stylized equilibrium models for costly information acquisition by agents in sealed-bid and ascending-price auctions. The models assume agents can make a single decision about whether or not to acquire information. This simplification makes equilibrium analysis tractable. Both models provide support for the analysis in this paper, and suggest that ascending auctions have better economic properties than sealed-bid auctions because they promote better decisions about information acquisition.

6.2. Auction models with costly participation

A number of economic models consider the costs associated with *participation* in an auction, for example costs of bid preparation and information acquisition [20,26,29, 45,49]. However, these models assume that all participation decisions are made as a *one-shot decision* before an auction starts, and none can distinguish between iterative and single-stage auctions, because agents decide whether or not to enter the auction *before* the auction starts. In one of the few models to allow agents to enter sequentially, Ehrman and Peters [12] compare the performance of different auctions for agents with one-shot participation costs. The authors show that a *sequential posted-price auction* is useful in settings with high costs of participation because it limits competition.

There is also a subtle, but interesting relation between Milgrom and Weber's [32] seminal work on auction design with affiliated values and the design of auctions with private values but costly elicitation. Milgrom and Weber show that the English auction outperforms other auctions in an affiliated value model. Information exchanged between agents during the auction allows agents to refine their estimates of value. In our setting, iterative auctions are preferred because bids from other agents provide information that improves an agent's preference elicitation decisions, and thus the accuracy of the final valuations. By analogy to Milgrom and Weber's *linkage principle*, which states that a seller should reveal any information that will help bidders to value a good, this suggests the importance of helping to mitigate the elicitation problem facing bidders in our model.

6.3. *Preference elicitation in auctions*

Parkes [37] and Nisan and Segal [35] (in a more general model) derive a lower bound on the information-revelation requirements for efficiency in combinatorial auctions. Nisan and Segal use this to emphasize that an indirect mechanism cannot outperform a DRM in the worst case, because one can always construct instances for which *exact* information is required to determine the efficient allocation (for example two bidders with almost identical values). This is a worst-case result, and does not preclude the many benefits that proxied and indirect mechanisms can enjoy in typical cases. Lahaie and Parkes [21] use this to emphasize the role of prices, and derive polynomial elicitation results for a computational learning theory model of elicitation. Many other papers [2,9,18,39,40] have considered the preference-elicitation problem in combinatorial auctions, but all for a *first-best* model in which the goal is to elicit just enough information to determine the efficient allocation. A number of studies have considered the role of proxied and indirect auctions in mitigating the cost of preference elicitation in multiattribute auctions [38,50].

Within a second-best model, Holzman et al. [17] characterize the properties of a family of VCG-based mechanisms with less communication complexity than the full-revelation mechanism, and Ronen [43], considers mechanism design for limited preference-revelation languages. In recent work, Blumrosen et al. [3,4] consider auction design with severely limited communication and are able to design single-item auctions for $\log(k)$ bits of communication with an optimal loss in efficiency and revenue. Iterative auctions are shown to have better performance, but only up to a factor of 2 in the amount of communication needed [4]. Unlike our work, these papers on preference-constrained settings work with an *a priori* fixed limit on the amount of elicitation.

7. Conclusions

Proxied and indirect mechanisms can outperform sealed-bid mechanisms when there is costly preference elicitation. For instance, carefully designed iterative auctions can allow proxy agents to follow optimal bidding strategies without complete preference elicitation. The dynamic price feedback in an ascending-price auction provides aggregate information about the preferences of other agents, and allows a proxy agent to elicit value information on parts of the outcome space that matter in the efficient allocation.

We have identified the central role of *incremental revelation* mechanisms, that allow a bidder to refine her preference information during an auction.

The ultimate goal in auction design for costly preference elicitation should be to develop incremental-revelation mechanisms that are *provably optimal*, given a particular model of costly elicitation. This will require methods to explicitly structure and coordinate preference elicitation decisions across a system of agents, so that the right level of elicitation emerges endogenously as a tradeoff between elicitation cost and value of information. Determining an informational hierarchy for classes of IRMs should be another interesting direction for future work. Finally, given that it appears difficult to determine equilibrium elicitation and bidding strategies, even in simple single-item auctions, we should work to design mechanisms with well-defined (and computable) second-best equilibria.

Acknowledgements

Thanks to Lyle Ungar, who provided early inspiration for this work, and to participants in the Agent Mediated Electronic Commerce and DIMACS market design workshops, where earlier versions of this work were presented. This work is supported in part by NSF grant IIS-0238147.

References

[1] A. Archer and E. Tardos, Truthful mechanisms for one-parameter agents, in: *Proc. 42nd IEEE Symp. on Foundations of Computer Science* (2001).

[2] A. Blum, J. Jackson, T. Sandholm and M. Zinkevich, Preference elicitation and query learning, Journal of Machine Learning Research 5 (2004) 649–667.

[3] L. Blumrosen and N. Nisan, Auctions with severely bounded communication, in: *Proc. 43rd Annual Symposium on Foundations of Computer Science* (2002).

[4] L. Blumrosen, N. Nisan and I. Segal, Multi-player and multi-round auctions with severely bounded communication, in: *Proc. 11th Annual European Symposium on Algorithms* (2003).

[5] E. Cantillon and M. Pesendorfer, Auctioning bus routes: The London experience, in: *Combinatorial Auctions*, eds. P. Cramton, Y. Shoham and R. Steinberg (MIT Press, 2006) Chapter 22.

[6] C. Caplice and Y. Sheffi, Combinatorial auctions for truckload transportation, in: *Combinatorial Auctions*, eds. P. Cramton, Y. Shoham and R. Steinberg (MIT Press, 2006) Chapter 21.

[7] O. Compte and P. Jehiel, On the virtues of the ascending price auction: New insights in the private value setting, Technical report, CERAS and UCL (2000).

[8] O. Compte and P. Jehiel, Auctions and information acquisition: Sealed-bid or dynamic formats?, Technical report, CERAS and UCL (2002).

[9] W. Conen and T. Sandholm, Preference elicitation in combinatorial auctions, in: *Proc. 3rd ACM Conf. on Electronic Commerce (EC-01)* (2001) pp. 256–259.

[10] V. Conitzer and T. Sandholm, Computational criticisms of the revelation principle, in: *Proc. AAMAS Workshop on Agent Mediated Electronic Commerce (AMEC V)*, Melbourne, Australia (2003).

[11] S. de Vries, J. Schummer and R. V. Vohra, On ascending Vickrey auctions for heterogeneous objects, Technical report MEDS, Kellogg School, Northwestern University (2004).

[12] C. Ehrman and M. Peters, Sequential selling mechanisms, Economic Theory 4 (1994) 237–253.

[13] K. Fong, Multi-stage information acquisition in auction design, Master's thesis, Harvard University (2003).

[14] A. Gibbard, Manipulation of voting schemes: A general result, Econometrica 41 (1973) 587–602.

[15] J. Green and J.-J. Laffont, Characterization of satisfactory mechanisms for the revelation of preferences for public goods, Econometrica 45 (1977) 427–438.

[16] A. Greenwald and J.O. Kephart, Shopbots and pricebots, in: *Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI-99)* (1999) pp. 506–511.

[17] R. Holzman, N. Kfir-Dahav, D. Monderer and M. Tennenholtz, Bundling equilibrium in combinatorial auctions, Games and Economic Behavior 47(1) (2004) 104–123.

[18] B. Hudson and T. Sandholm, Effectiveness of query types and policies for preference elicitation in combinatorial auctions, in: *Proc. 3rd Int. Joint. Conf. on Autonomous Agents and Multi Agent Systems* (2004) pp. 386–393.

[19] M.O. Jackson, Mechanism theory, in: *The Encyclopedia of Life Support Systems* (EOLSS Publishers, 2000).

[20] C.D. Kolstad and R.M. Guzman, Information and the divergence between willingness-to-accept and willingness-to-pay, J. Env. Econ. Manag. 38(1) (1999) 66–80.

[21] S.M. Lahaie and D.C. Parkes, Applying learning algorithms to preference elicitation, in: *Proc. ACM Conf. on Electronic Commerce* (2004) pp. 180–188.

[22] K. Larson and T. Sandholm, Bargaining with limited computation: Deliberation equilibrium, Artificial Intelligence 132(2) (2001) 183–217.

[23] K. Larson and T. Sandholm, Experiments on deliberation equilibria in auctions, in: *Proc. 3rd Int. Joint. Conf. on Autonomous Agents and Multi Agent Systems* (2004) pp. 394–401.

[24] K. Larson and T.W. Sandholm, Costly valuation computation in auctions, in: *Proc. Theoretical Aspects of Rationality and Knowledge VII* (2001) pp. 169–182.

[25] J.O. Ledyard, The design of coordination mechanisms and organizational computing, Journal of Organizational Computing 1 (1993) 121–134.

[26] T.K. Lee, Competition and information acquisition in first price auctions, Economic Letters 19 (1985) 129–132.

[27] D. Lehmann, R. Muller and T. Sandholm, The winner determination problem, in: *Combinatorial Auctions*, eds. P. Cramton, Y. Shoham and R. Steinberg (MIT Press, 2004) Chapter 12.

[28] D. Lehmann, L.I. O'Callaghan and Y. Shoham, Truth revelation in approximately efficient combinatorial auctions, Journal of the ACM 49(5) (2002) 577–602.

[29] D. Levin and J.L. Smith, Equilibrium in auctions with entry, American Economic Review 64 (1994) 585–599.

[30] A. Mas-Colell, M.D. Whinston and J.R. Green, *Microeconomic Theory* (Oxford University Press, Oxford, 1995).

[31] R.P. McAfee and J. McMillan, Analyzing the airwaves auction, J. Econ. Perspect. 10 (1996) 159–175.

[32] P. Milgrom and R. Weber, A theory of auctions and competitive bidding, Econometrica 50 (1982) 1089–1122.

[33] D. Mishra and D. Parkes, Ascending price Vickrey auctions for general valuations, Technical report, Harvard University (2005).

[34] N. Nisan, Bidding languages for combinatorial auctions, in: *Combinatorial Auctions*, eds. P. Cramton, Y. Shoham and R. Steinberg (MIT Press, 2006) Chapter 9.

[35] N. Nisan and I. Segal, The communication requirements of efficient allocations and supporting prices, Journal of Economic Theory (2005) to appear.

[36] D.C. Parkes, Iterative combinatorial auctions: Achieving economic and computational efficiency, Ph.D. thesis, Department of Computer and Information Science, University of Pennsylvania (2001).

[37] D.C. Parkes, Price-based information certificates for minimal-revelation combinatorial auctions, in: *Agent Mediated Electronic Commerce IV: Designing Mechanisms and Systems*, Lecture Notes in Artificial Intelligence, Vol. 2531 (2002) pp. 103–122.

[38] D.C. Parkes and J. Kalagnanam, Models for iterative multiattribute procurement auctions, Management Science 51(3), Special Issue on Electronic Markets (2005) 435–451.

[39] D.C. Parkes and L.H. Ungar, Iterative combinatorial auctions: Theory and practice, in: *Proc. 17th National Conference on Artificial Intelligence (AAAI-00)* (2000) pp. 74–81.

[40] D.C. Parkes and L.H. Ungar, Preventing strategic manipulation in iterative auctions: Proxy agents and price-adjustment, in: *Proc. 17th National Conference on Artificial Intelligence (AAAI-00)* (2000) pp. 82–89.

[41] D.M. Reeves and M.P. Wellman, Computing equilibrium strategies in infinite games of incomplete information, in: *Proc. of Game-Theoretic and Decision-Theoretic Workshop at AAMAS'03* (2003).

[42] L. Rezende, Mid-auction information acquisition, Technical report, Department of Economics, Stanford University (2001).

[43] A. Ronen, Mechanism design with incomplete languages, in: *Proc. 3rd ACM Conf. on Electronic Commerce (EC'01)* (2001).

[44] S. Russell and E. Wefald, Principles of metareasoning, Artificial Intelligence 49 (1991) 361–395.

[45] W.F. Samuelson, Competitive bidding with entry costs, Economic Letters 17 (1985) 53–57.

[46] T.W. Sandholm, An implementation of the Contract Net Protocol based on marginal-cost calculations, in: *Proc. 11th National Conference on Artificial Intelligence (AAAI-93)* (1993) pp. 256–262.

[47] T.W. Sandholm, Limitations of the Vickrey auction in computational multiagent systems, in: *Second International Conference on Multiagent Systems (ICMAS-96)* (1996) pp. 299–306.

[48] T.W. Sandholm, Automated negotiation, Communications of the ACM 42(3) (1999) 84–85.

[49] M. Stegeman, Participation costs and efficient auctions, Journal of Economic Theory 71 (1996) 228–259.

[50] A.V. Sunderam and D.C. Parkes, Preference elicitation in proxied multiattribute auctions, in: *Fourth ACM Conf. on Electronic Commerce (EC'03)* (2003) pp. 214–215.

[51] W. Vickrey, Counterspeculation, auctions, and competitive sealed tenders, Journal of Finance 16 (1961) 8–37.