A Study of Limited-Precision, Incremental Elicitation in Auctions

Alexander Kress
Department of Computer Science
University of Toronto
Toronto, ON, M5S 3H5
akress@cs.toronto.edu

Craig Boutilier
Department of Computer Science
University of Toronto
Toronto, ON, M5S 3H5
cebly@cs.toronto.edu

ABSTRACT

We investigate the design of iterative, limited-precision mechanisms for single-good auctions with dominant strategy equilibria. Our aim is to design mechanisms that minimize the number of bits required to determine approximately optimal allocations by sequentially asking bidders to reveal their valuations with increasing precision, and limiting participation to those bidders who might win. We prove several necessary conditions that severely restrict the space of mechanisms satisfying our criteria. We also study empirically the optimization of the parameters of our sequential mechanisms, and how number of bidders and cost of communication impact expected amount of communication, expected loss in welfare, and other measures. Finally, we show that incremental limited-precision mechanisms offer advantages over fixed, single-shot mechanisms.

Categories and Subject Descriptors

J.4 [**Computer Applications**]: Social and Behavioral Sciences—*economics*; I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*intelligent agents, multiagent systems*

Keywords

auctions, approximation, mechanism design, revelation principle

1. INTRODUCTION

Day-to-day business transactions have come to rely increasingly on the computer networks that link market participants by providing fast, seamless communication and negotiation channels. This migration to online negotiation has led to the development of more and more sophisticated software agents that mediate such transactions. However, since the interests of the parties on whose behalf such agents act generally conflict, ideally such agents should reason strategically according to the well-studied principles of game theory and economics. As such, recent research in computer science and economics has focused on the design of economic agents and the mechanisms through which they interact.

Mechanism design [10] has played a central role in much of this research. Key results in mechanism design, such as the revelation principle, have had a strong influence on the direction taken by research at the intersection of the economics and computer science. However, recently, limitations of standard approaches to mechanism design have been identified, and are starting to be addressed. Chief among these is the computational complexity of the problems faced by software agents who must interact. For instance, mechanisms based on the revelation principle must reveal their type (generally, their utility function) accurately. This presents a problem in circumstances where utility functions are large and difficult to communicate effectively and/or hard to compute accurately. Recent research has begun to examine methods involving limited or incremental elicitation of types to circumvent some of these difficulties (see, e.g., [1, 2, 14, 12, 13, 8, 3], much of this in the context of (single-good or combinatorial) auctions.[1]

In this paper, we pursue the same line of research. Specifically, in the context of single-good auctions, we develop and analyze mechanisms that allow bids with *limited precision* and that incrementally elicit bids by allowing bidders to sequentially refine their bids. We propose various natural constraints on such incremental mechanisms—such as an activity constraint that removes any bidder from the auction if they provably do not have the highest valuation—and show that any mechanism satisfying these constraints, and having dominant strategy equilibria, must have a very restricted form. Specifically, it must take the form of an ascending-style auction (much like a Japanese auction) in which the space of bidders' actions can be safely restricted to just two actions. We then study one class of mechanisms of this type in detail. Our *adaptive, symmetric, incremental auction (ASIA)* is parameterized by a (conditional) sequence of price levels, but has simple dominant strategies based on any such thresholds. Because of this, the thresholds can be optimized with respect to specific priors over bidder types, using a straightforward Markov decision process formulation. In this respect, we combine the spirit of traditional mechanism design (where typically, though not exclusively, dominant-strategy inducing mechanisms exist independent of the priors), optimal auction design [11] (where the mechanism varies in a parameterized way with the prior), and automated mechanism design [5] (where the mechanism is optimized using the specific priors). Importantly, we can optimize the mechanism to account for cost of communication (or computation) as well [13]. Finally, we show that by relaxing some of the original constraints, a stochastic variant of ASIA (STASIA) emerges. This mechanism is interesting for two reasons. First, it maintains the benefits of the original mechanism, but with much improved efficiency and revenue properties. Second, it can

Draft. Comments welcome.
ACM-EC 2005 Vancouver BC

[1]For other criticisms of the revelation principle, see [4].

be shown that, under certain conditions, STASIA is superior to any threshold-based, single-shot auction.

We begin in Section 2 with a discussion of relevant background and a brief overview of related work on mechanism design for auctions that adopts limited precision or incremental bidding. We discuss a simple one-shot, limited-precision mechanism with a *take-it-or-leave-it* opportunity in Section 3 in order to introduce simple limited-precision mechanisms and the optimization of their parameters. Our main results are found in Sections 4 and 5. In Section 4 we describe an incremental mechanism based on the same limited-precision principles. We show that, under certain assumptions, all mechanisms must have a very restricted "ascending" form. We also describe a model for the optimization of price thresholds, how it reduces the expected amount of communication, and the impact of communication cost. In Section 5 we present STASIA, a stochastic version of the ASIA mechanism, and compare its performance to the deterministic version. We also discuss the advantages of STASIA relative to the one-shot, threshold-based mechanisms of Blumrosen and Nisan [1].

2. INCREMENTAL ELICITATION

Mechanism design deals with the problem of designing a game—in which a collection of self-interested agents interact—so as to optimize some objective on the part of the designer [10]. More formally, we suppose we have a collection of agents (e.g., potential buyers of some good) and some set of outcomes O (e.g., the allocation of a good to a particular agent at a specific price). Each agent i has a type $t_i \in T_i$ known only to itself and utility function u_i over types and outcomes, where $u_i(t_i, o)$ reflects the utility of outcome o to agent i if its type is t_i. For example, the type of the agent might dictate its valuation for the good being auctioned, and its utility for any outcome o in which it obtains the good as given by its valuation less the price paid. We generally assume some commonly known prior over T_i for each i. We suppose the mechanism designer has some *social choice objective* f it wishes to optimize, where $f(\vec{t}, o)$ denotes the objective value (to the designer) of outcome o when the agents (jointly) have type vector $\vec{t} \in \times T_i$. For example, the designer may wish to maximize social welfare by ensuring the good goes to the agent with the highest valuation; thus f would denote the total welfare of outcome o given type profile $\vec{t}$. A *(deterministic) mechanism* comprises a set of strategies Σ_i for each agent i, and an outcome rule $g : \Sigma \rightarrow O$ mapping strategy profiles $\Sigma = \times\Sigma_i$ into outcomes. As such a mechanism induces a Bayesian game among the agents. A mechanism *implements* a social choice objective f iff, in equilibrium, the outcome of the game is $o \in \arg\max f(\vec{t}, o)$ whenever the agents types are given by $\vec{t}$.[2] The mechanism design problem then corresponds to finding a mechanism that implements f using the desired notion of equilibrium (commonly, Bayes-Nash or dominant strategy equilibrium).

The *revelation principle* makes the mechanism design problem somewhat simpler by noting that if a mechanism exists that implements f, then a direct, incentive-compatible mechanism exists for f as well. In other words, we let strategies correspond to types (hence agents directly reveal their types), and in equilibrium, each agent will *truthfully* reveal its type [10]. The revelation principle has led to an almost exclusive focus on direct, incentive-compatible mechanisms.

Direct mechanisms have several drawbacks. Most importantly, the requirement that agents report their true type accurately and completely imposes a severe, and often unnecessary, burden on agents in terms of computation (e.g., to accurately compute their valuation for a good) and communication. In general settings, simply communicating a full utility function to the mechanism can be prohibitive (e.g., in combinatorial auctions) [3]. But even when utility function can be represented by a single number, computing this accurately may be costly [13] and the level of precision required to maximize f may not require accurate reporting (e.g., if the agent with the highest valuation for a good has a valuation that is very different from all others, then this can be determined with bids of very limited precision). Furthermore, communication costs for single valuations may be relevant in high volume settings, such as auctions for packet routing [1].

In the context of single-good auctions, significant research has tackled these problems. Among these, the work of Blumrosen, Nisan and Segal [1, 2] bears the closest relation to ours, studying in great theoretical detail many of the issues we focus on here. In their work, bidders for a single good can offer one of k distinct bids, requiring communication of $\lg(k)$ bits. Because the possibility of "ties" can have an effect on revenue or welfare (since it becomes impossible to distinguish valuations to arbitrary precision), they introduce tie-breaking rules that use a fixed ordering of bidders, giving priority to bidders higher in the ordering. They show that such *priority games* give rise to dominant *threshold strategies*, where each agent adopts a set of threshold values and associates a specific bid with each induced interval. In the terms we describe below, we view this as a *limited-precision* mechanism, since bidders can be interpreted as revealing their valuation directly (and "truthfully" since we can assign *a priori* a meaning to each bid based on the dominant threshold strategy of any agent), but with limited precision. As such, the mechanism can be designed with the thresholds given first.

The results of [1, 2] show that, for any prior over valuations, there exists a priority game that is welfare-optimal among those mechanisms allowing at most k distinct bids, and that the welfare loss relative to an optimal unlimited communication auction (e.g., a Vickrey auction [16]) is bounded by $O(\frac{1}{k^2})$ for a fixed number of players (and this bound is tight in the case of uniformly distributed valuations). In [2], a *sequential* mechanism is considered in which bidders reveal their bits in sequence, but with full knowledge of all previously revealed bits by other players. While such mechanisms can do better, the gain is limited in the sense that a one-shot mechanism can achieve the same welfare with only twice as much communication.

The sequential mechanisms we consider in this paper have a very different flavor to those above. In spirit, they are more reminiscent of ascending auctions, and the work of Parkes [13], who considers the cost of preference elicitation. Parkes's model admits agents with uncertain valuations, who must decide whether to refine bounds on their current valuation in order to bid in various forms of auctions. A computational cost is assumed, and the focus is on the computational strategies of the bidding agents themselves (rather than the mechanism design problem).[3]

The bisection auction of [7] is sequential as well, refining the mechanism's estimates of agent valuations up to some limited precision, but otherwise duplicating a Vickrey auction. In order to determine payment, losing bidders must continue to answer queries until the second-highest price can be determined to the required degree of precision. Survival auctions [6] bear some relation in motivation to our work as well—however these auctions do allow

[2]More commonly, one speaks of implementing a social choice function f_c [10]. These views are equivalent if one defines $f_c(\vec{t}) = \arg\max f(\vec{t}, o)$.

[3]Larson and Sandholm [8, 9] study similar phenomenon in auctions and bargaining settings, deriving equilibria that account for cost of computation.

for revelation of full precision bids. Survival auctions combine aspects of sealed-bid and ascending auctions to speed up ascending auctions by eliminating bidders at each round and setting minimum bids at each round. Finally, the model of rational computation proposed in [15] bears some relationship to our work, considering the question of the communication complexity of computing certain functions through auction-like mechanisms.

Our model incorporates various aspects of all of the work described above. Our general view can be characterized as follows. Rather than consider direct mechanisms in which an agent's type is revealed fully, we are interested in *limited-revelation mechanisms*, in which the set of strategies is limited. For example, if only a finite number of strategies are made available to each agent in the case where the type space is continuous (e.g., the space of valuations is the interval $[0,1]$), strategy choice cannot be used to fully distinguish type. Note that the amount of communication required to signal a choice to the mechanism is $\lg(k)$ bits if only k strategies are allowed. A *direct* limited revelation mechanism is one in which each strategy is interpreted as corresponding to a subset of possible types. All of the aforementioned work can be viewed as direct in this sense. A *limited-precision* mechanism is a direct mechanism in which each strategy choice can be associated with a unique utility *interval* (or hyperrectangle in the multidimensional case), and can be viewed as conveying the agent's type with limited precision.

We also consider *incremental* mechanisms, iterative mechanisms in which the types of agents are revealed sequentially. Intuitively, we require that each move by an agent *refines* the space of possible valuations for that agent. Incremental mechanisms have a certain intuitive appeal, since for many mechanism optimization criteria, a sequential approach where valuations are refined allows certain agents to be "ruled out" early, thus obviating the need for them to compute or communicate their valuations with maximum precision. This can be exploited to great effect [13, 7].

Finally, we wish to consider the use of optimization criteria in the mechanism design problem *that account for properties of the mechanism itself.* For example, when maximizing social welfare, if we assume that each bit communicated to the mechanism has a certain cost, it makes sense to incorporate the expected cost of communication in the objective function when designing the mechanism. While much of the work above motivates limited revelation by appealing to cost of computation or communication, none (explicitly) considers the problem of designing mechanisms in which this factors into the objective. We will *explicitly* consider this (albeit in a somewhat restricted way).

3. A ONE-SHOT TIOLI MECHANISM

We assume a single good is to be auctioned to a set of n players, each player i having valuation $v_i \in [0,1]$ for the good, where v_i is drawn from some commonly known prior f_i. We begin by describing an especially simple one-shot, limited-precision mechanism that is similar is many ways to priority games [1, 2], specifically in its use of thresholds. This sets the stage for discussion of our main results on incremental mechanisms. Our *limited-precision, take-it-or-leave-it (LP-TIOLI)* auction works as follows: we fix a set of k prices, $0 \le p^1 < p^2 < \ldots < p^k \le 1$. Each bidder announces one of these prices to the mechanism; let B be the set of such bids. If only one bidder announces the highest bid in B, the good is allocated to this bidder at the second highest bid (from among the k possible bids). If the highest bid is offered by two or more bidders, a random highest bidder is selected and is made a take-it-or-leave-it (TIOLI) offer for the good at the highest bid price (see [14] for more on TIOLI offers). If the offer is rejected, the good goes unallocated. Thus, LP-TIOLI can be viewed as a limited-precision variant of a second-price auction (the highest bidder is selected and offered the good for the second-highest bid) with the addition of a TIOLI option offered to the winner in the case of ties. The mechanism differs from priority games in that players are treated symmetrically, and in the added TIOLI component.[4]

It is not hard to see that LP-TIOLI has a dominant strategy equilibrium: each agent should bid the least price p^l at least as great as her valuation v_i; that is, bid any p^l such that $p^{l-1} \le v_i \le p^l$. It should be clear that if the highest bid excluding that of i is greater than p^l or less than p^{l-1}, then bidding p^l is optimal (much as in a Vickrey auction). If the highest bid is p^l, then i bidding p^l offers a utility of 0 (since she will refuse the offer at price p^l if chosen); but any bid above or below p^l has utility less than or equal to 0.

We note that the TIOLI component is critical to the existence of dominant strategies.[5] If we removed the TIOLI component and simply considered a limited-precision second-price auction, there are cases in which bidding above her valuation is an agent's best move, and cases in which bidding just below her valuation is optimal. The best choice depends on the prior over other agent valuations, thus requiring a Bayes-Nash equilibrium analysis. Interestingly, one can derive *threshold-based* Bayes-Nash equilibria.

Note that in the dominant strategy equilibrium, the good will be allocated in the case that there is a unique highest bidder, and optimal welfare will result. The case of multiple highest bidders results in the good remaining unallocated, leading to a loss in efficiency and revenue. One can optimize this mechanism by setting price thresholds appropriately and optimizing for either expected revenue or social welfare, exploiting specific distributional information. One appeal of this mechanism is that the dominant strategy equilibrium has the same form regardless of the specific prices. Thus, optimization of price thresholds has no impact on the strategic reasoning of the bidders (e.g., approximation of the optimal parameters won't change a bidder's strategic deliberations).

As an exercise, we used gradient ascent to derive (locally) optimal price levels for LP-TIOLI using both expected revenue and expected welfare (we omit the expressions for expected revenue and welfare for LP-TIOLI as a function of the prices p^i). Figure 1(a) shows the expected revenue as a function of the number of possible bids for two bidders, each with valuations drawn from a (truncated) half-Gaussian density function. The unoptimized prices (uniformly set at $0, 1/k, 2/k, \ldots$) fare much worse than the optimized prices.

Note that the revenue exceeds that of a Vickrey auction by a considerable margin and approaches that of the optimal (Myerson) auction [11]. Even though the good remains unallocated in many cases, the fact the second-highest bid lies above the second-highest valuation ensures that more surplus is extracted from the winner when the good is allocated (relative to Vickrey). Figure 1(b) shows similar results for five bidders with uniformly distributed valuations. Here the probability of leaving the good unallocated outweighs the benefit of using limited precision.

A rough approximation of loss in efficiency, measured as the probability of the good not being allocated (recall that by the rules of LP-TIOLI if the good is allocated, it must go to the bidder with the highest valuation), is shown for the five-bidder, uniform valuation case in Figure 1(c). Finally, we illustrate the (locally) optimized price levels found by gradient ascent for the two-bidder, uniform case in Figure 2. Note the implicit setting of a "reservation value" of 0.5 in the case of revenue maximization (exactly as in the Myerson auction).

[4] Blumrosen and Nisan [1] show that symmetric mechanisms, in the two-player case, are not as efficient as priority games.

[5] Simply leaving the good unallocated in the case of a tie would have the same effect.

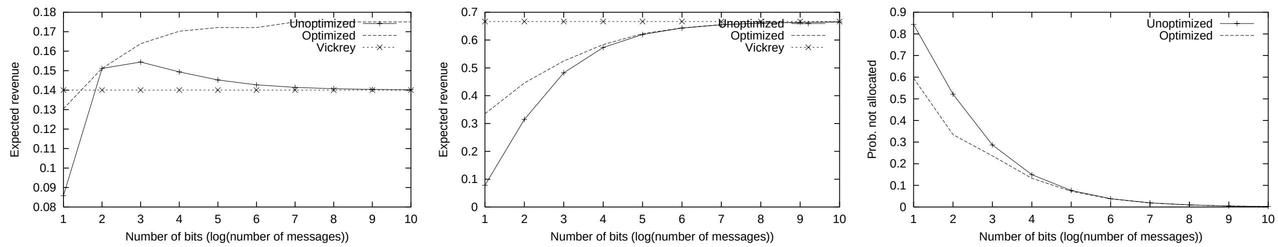

Figure 1: (a) Expected revenue (truncated half-Gaussian, $\mu = 0, \sigma = 0.3$); (b) Expected revenue (uniform distribution); (c) Probability of good unallocated (uniform distribution).

The approach to optimization is closely related to optimal auction design [11] and automated mechanism design (AMD) [5] in the sense that we wish to optimize the mechanism using specific distributional information; however, we restrict our attention to a class of mechanisms with specific parameterized dominant strategies, and simply optimize the parameters, rather than leaving the whole mechanism "up for grabs." The optimization task need not be limited to the traditional objectives, but can incorporate "mechanism properties" such as expected amount of communication as well. In the next section on incremental mechanisms, we will study optimization with respect to welfare that accounts for the cost of communication.

4. AN INCREMENTAL, LIMITED PRECISION MECHANISM

While LP-TIOLI achieved the objective of limiting both communication and revelation, it is not very flexible in the following sense. Many agents may reveal much more about their valuations than is required to determine the desired outcome. For example, even though k bits of precision may be used, we may easily be able to rule out many (or most) bidders as potential winners with far fewer than k bits. Doing so allows the mechanism to engage in precise elicitation of valuation only of a few players, reducing overall communication or computation cost. This leads to a natural motivation for sequential, or *incremental* mechanisms [13, 7]. For example, if one wants to maximize social welfare where communication cost is included, an incremental mechanism offers greater potential than a one-shot mechanism.

In this section, we introduce a class of incremental mechanisms for single good auctions by proposing a set of intuitive restrictions on how the mechanism proceeds. We show that all such mechanisms with dominant strategy equilibria have a very specific "ascending" nature, and describe how to optimize the messages (or prices) for one such mechanism.

4.1 General Assumptions

We consider iterative mechanisms for the allocation of a single good, in which bidders make moves sequentially (possibly with indirect knowledge of the prior moves of other players). Furthermore, using the terminology introduced above, we will consider *direct, incremental* mechanisms, where each move is viewed as providing some information about the player's valuation, and each successive move *refines* the information revealed earlier. The mechanism proceeds in iterations, with each player making a move each time.

We restrict the class of mechanisms further by imposing some restrictions that seem rather natural in the space of single-good auctions. As before we assume each agent i has valuation $v_i \in [0, 1]$.

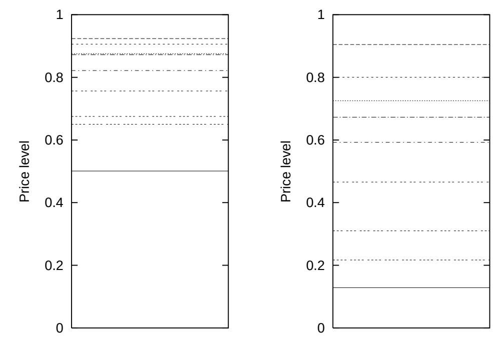

Figure 2: Optimized prices (revenue, welfare)

- Agents reveal one of a finite set M^t of messages at each round t (possibly with activity constraints).
- Finite sequences of messages are comparable. In other words, there exists a total order $\leq$ such that either $m \leq m'$ or $m' \leq m$ for any length t message sequence m and length t' message sequence m'. As a consequence, for any t, there is a minimum and maximum sequence of length t, and a minimum and maximum "extension" of any such length t sequence to length $t + k$. This allows message sequences to be interpreted as bids, and the mechanism to be interpreted as limited precision.[6]
- The auction is fully deterministic in a sense that a good is never allocated randomly. To formalize this we assume that the auction terminates at iteration t with an allocation of the good only if t is such that some bidder has specified a unique greatest message sequence. The good is allocated to that bidder. This takes advantage of the sequential nature of the mechanism in an intuitive and powerful way. Note that the auction may terminate at some finite stage without allocating the good if no bidder has offered a unique greatest message sequence.
- We assume quasi-linear utility and *ex post* rationality. As such, only the winner makes a payment.

[6] This is not to say that this is the only way of realizing incremental elicitation. Arbitrary query languages (e.g., asking agents to communicate upper and lower bounds on valuations) are certainly possible [13]. We consider only mechanisms that can be viewed as allowing "limited-precision bids."

A strategy σ_i for agent i in such an incremental mechanism requires a choice of message m^t at each round t as a function of its type v_i and its history h_i^{t-1} up to that point. For simplicity, we assume that i knows only what she has bid, not what other players have bid (and when we discuss limiting participation, whether or not she is active). But the results below do not rely on this. We let $\sigma_i(v_i)$ to denote i's choice of action as a function of history when her valuation is v_i, and σ_{-i} to denote the strategy profile of all agents except i.

We are interested in mechanisms satisfying the restrictions above in which dominant strategy equilibria exist. It turns out that all such mechanisms have a very specific form, in particular, they are *increasing price mechanisms*:

Defn 1 An iterative mechanism is an *increasing price mechanism* if the following condition holds for any player i. Let v_i be i's valuation and let s_{-i} be any sequence of moves of all other players (with $s_{-i}[k]$ denoting the length k initial segment). Let t be the least iteration at which i can win the good playing against sequence s_{-i} and pay a price $p \leq v_i^{\max}$, and s_i a sequence of moves that achieves this. Then for any other sequences $\hat{s}_{-i}$, $\hat{s}_i$ where $\hat{s}_{-i}[t] = s_{-i}[t]$, if i wins playing $\hat{s}_i$ against $\hat{s}_{-i}$ at some stage $k > t$ with price $\hat{p}$, then $\hat{p} \geq p$.

Intuitively, an increasing price mechanism has the following property for any player i: if we fix the moves of her opponent, and let t be the earliest round at which i could win, then the price paid by i at round t (against these fixed opponents), should she choose a strategy that wins at t, must be no more than the price paid if i wins at any later round against the same opponent moves.

Proposition 1 *All single-good, incremental auctions satisfying the conditions above and having dominant strategy equilibria are increasing price mechanisms.*

Proof sketch: Suppose the incremental mechanism is not an increasing price mechanism. Then for some player i we can find an opponent strategy profile σ_{-i} where i can win with positive utility at round t (i.e., pay less than her valuation), but can win with greater utility at a later round k (by paying less). Thus the optimal σ_i would choose moves that ensure winning at some round $k > t$. However, against such a σ_i, we can easily construct an alternative profile σ'_{-i} where i loses by playing σ_i but could win (with positive utility) using an alternative strategy. To see this, consider some σ'_{-i} where at every iteration after t at least one opponent of i bids at least as high as the bid dictated by σ_i; this ensures that i will never be able to win after iteration t. However, by assumption i could win at t with positive utility were she to follow a different strategy. Therefore, there does not exist a dominant strategy for i. ◄

We now consider another important restriction on the class of incremental mechanisms that is facilitated by the sequential nature of the elicitation process. We say that an incremental mechanism *limits participation* iff no player is allowed to participate once her utility function has been refined to the extent necessary to permit optimization of the mechanism's objective. In our limited-precision auction setting, this corresponds to the following activity rule: a player remains *active* as long as her message sequence is at least as great as that of any other bidder. This implies that if the auction has not yet terminated, then i is active iff her bids at all prior rounds have been tied with other highest bids. Limiting participation is one of the main factors that allow us to achieve a reduction in both communication and revelation. Intuitively, it ensures that we deal only with players whose valuations are potentially high enough to win and to remove everyone else from the auction.

The following lemmas describe necessary conditions that must hold for any incremental auction with limited participation to have dominant strategy equilibria.

Lemma 1 *Let player i have dominant strategy σ_i and valuation v_i. Let σ_{-i} and v_{-i} be such that if $\sigma_i(v_i)$ is played against $\sigma_{-i}(v_{-i})$, the mechanism terminates at iteration t with i winning and paying $p_i < v_i$. Then for any $v'_i \geq v_i$ and any dominant strategy σ'_i, if i plays $\sigma'_i(v'_i)$ against $\sigma_{-i}(v_{-i})$, we must have: (a) i wins and pays p_i (as with $\sigma_i(v_i)$); and (b) the mechanism terminates at iteration t (as with $\sigma_i(v_i)$).*

Proof sketch (a) Assume the claim is not true. Then we can show that there exists a strategy σ'' s.t. for some valuation $v''_i \geq v_i$ player i would chose to play $\sigma''_i(\hat{v}_i)$ for some $\hat{v}_i$ instead of his dominant strategy $\sigma'_i(v''_i)$. This shows that i does not have a dominant strategy. (b) Assume the claim is not true. Then by claim (a), we can show that there exists a player without a dominant strategy, using the same approach as in proof of Proposition 1. ◄

Intuitively, Lemma 1 shows that with limited participation, if i has valuation v_i and wins at iteration t with price p using a dominant strategy, then for any greater valuation, it must win at the same iteration for the same price using any dominant strategy, against fixed opponents.

Defn 2 The *Last Profitable Iteration* for i, given history h^t and valuation v_i is defined as follows: we say $LPI(v_i, h^t) = now$ if there exist moves of other players such that i can profitably win at round $t+1$, but cannot profitably win at any future round; $LPI(v_i, h^t) = future$ if there exist moves of other players such that i can profitably win at some round later than $t+1$; $LPI(v_i, h^t) = past$ otherwise.

Note that it is weakly dominant to avoid winning after iteration $t+1$ if $LPI(v_i, h_i^t) = now$, since if *i* wins at any iteration greater than *t* she would make a payment equal to at least her valuation.[7]

Lemma 2 *Let player i have dominant strategy σ_i and valuation v_i. Suppose there exists some σ_{-i}, v_{-i} such that, if i plays $\sigma_i(v_i)$ against $\sigma_{-i}(v_{-i})$ (inducing history h_i^t), then $LPI(v_i, h_i^t) = future$ and player i remains active at iteration $t+1$. Then for any other $\widehat{v}_i \geq v_i$ and any dominant strategy $\widehat{\sigma}_i$ it must be the case that if i plays $\widehat{\sigma}_i(\widehat{v}_i)$ against $\sigma_{-i}(v_{-i})$ (inducing history $\hat{h}_i^t$ up to iteration* t*), we must have $h_i^r = \hat{h}_i^r$ and $\sigma_i(v_i, h_i^r) = \widehat{\sigma}_i(\widehat{v}_i, \hat{h}_i^r)$, for all $r \leq t$.*

Proof sketch: Since $LPI(v_i, h_i^t) = future$, it is possible for *i* to win against some σ_{-i}^d with positive utility at some iteration $k > t+1$, where σ_{-i}^d is identical to σ_{-i} for all histories up to length *t*. Lemma 1 ensures that for all $\widehat{v}_i \geq v_i$ and any dominant strategy $\widehat{\sigma}_i$, $\widehat{\sigma}_i(\widehat{v}_i)$ wins against σ_{-i}^d with the same payment and at the same iteration as $\sigma_i(v_i)$. Then, assuming limited participation, we can use an inductive argument to show that all histories and bid sequences will be the same. ◄

The results above lead to some interesting conclusions. If player *i*, with valuation v_i, is facing a history h_i^t where $LPI(v_i, h_i^t) = future$, then any dominant strategy must choose the same actions (against the same opponent moves) at all stages prior to t for *any* valuation $v'_i \geq v_i$. Thus if $LPI(v_i, h_i^t) = future$, i is not required to "bid" so much as she must simply signal her willingness to participate in the future. The only time an "interesting" message is

[7] Any player can avoid winning by always offering the "least" message $minbid$ at each iteration.

proposed is when $LPI(v_i, h_i^t) = now$. This severely restricts the space of limited-precision mechanisms that admit dominant strategies. The next lemma describes another useful restriction.

Lemma 3 *Suppose i has a dominant strategy and that for some history h_i^t and valuation v_i, $LPI(v_i, h_i^t) = now$. Then there is a dominant strategy σ_i in which $\sigma_i(v_i, h_i^k) = minbid$ for any $k > t+1$ and history h_i^k s.t. $h_i^k[t] = h_i^t$.*

Proof Sketch: Since $LPI(v_i, h_i^t) = now$, if i wins at any iteration later than $t+1$, she pays a price at least as high as her valuation. Therefore it is weakly dominant in this case to ensure the object is never won. Since the object is allocated only to a *unique* highest bidder, bidding $minbid$ is optimal, since it ensures i never wins. ◄

Lemma 3 implies that, should all agents follow the suggested dominant strategy, then if there is a tie among the highest bidders, and each is at their LPI, then the object will never be allocated. Hence, extra communication from this point will be wasted.

From this section we would like the reader to retain the effect the necessary conditions have on the structure of any auction fitting the given requirements. To reiterate, in defining the auction it is sufficient to provide the players with just two possible actions during all iterations (except possibly the last one), also to reduce communication costs the players should be removed from the auction as soon as they are found to be suboptimal.

Finally, in what follows, we will assume mechanisms are symmetric. Thus the mechanism determines its outcomes (i.e., winners and payments) in a way that is independent of a player's identity.

To summarize, given the assumptions and requirements discussed, the structure of the auctions one need consider are rather restricted. Specifically, in defining the auction it is sufficient to provide players with just two possible actions (except possibly the last iteration). Furthermore, to reduce communication costs the players should be removed from the auction as soon as they are fail to remain among the "high bidders."

4.2 Examples of Incremental Mechanisms

We now define two symmetric, incremental elicitation mechanisms that satisfy the properties derived above.

The first mechanism is a simple *adaptive, symmetric, incremental auction (ASIA)* with a TIOLI flavor. The rules of this mechanism are in fact very similar to the rules of the *Japanese auction* with discrete bid levels. ASIA operates as follows. Initially, all players are active. At iteration *t*, the mechanism announces price p^t to all active players, with $p^t \geq p^{t-1}$. All active players reveal either 1, indicating a willingness to purchase the good for p_t (intention to participate), or 0, indicating a desire to become inactive. The mechanism terminates when: only one player bids 1, in which case that player receives the good and pays the last announced price; or all active players bid 0, in which case the good is not assigned and no payments are made.

Proposition 2 *In the ASIA mechanism, it is a weakly dominant strategy of any player i with valuation v_i to bid 1 at iteration t as long as $v_i > p^t$, and to bid 0 otherwise.*

The ASIA mechanism is basically an ascending-price auction with bid levels set by the mechanism. As long as the sequence of announced prices is non-decreasing, Proposition 2 will hold no matter what prices are announced. Furthermore, as the mechanism is aware of how many players are currently active, an opportunity for optimization of the prices exists. We discuss this in the next section. It is worth noting that the auction may not terminate if there are two or more bidder valuations that lie above the "highest price" that a specific instance of the ASIA mechanism could announce (e.g., if there is some $v', v'' \in [0,1]$ for two distinct bidders s.t. $v' > p^t$, $v'' > p^t$ for all t). This can be prevented of course if the prices are set so that $p^t = v_{\max}$ for some finite t. Note also that early termination (e.g., stopping at round t before a highest bid is determined) does not affect incentives or the dominant strategies above.

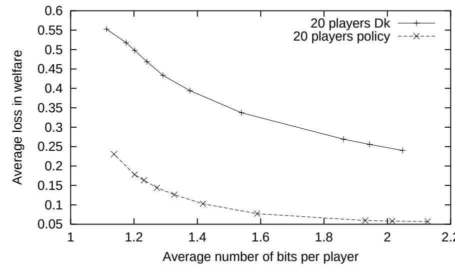

Figure 3: Comparison of loss in welfare in ASIA under optimized policy vs. ASIA under *Divide(k)* price update rule

The *second-price, symmetric incremental auction (SPSIA)* incorporates the LPI notion more directly than ASIA. The mechanism operates as follows. At each iteration *t*, a price interval $[a_t, b_t]$ is announced to all active players, where $a_t < b_t$. We require only that $a_k \geq b_{k-1}$ for any $k > 1$. All active players then announce 0 or 1, where as before 1 indicates "Intention to participate" and 0 indicates that the player wishes to become inactive. The auction continues as long as two or more players bid 1. If only one player bids 1 or if all players bid 0 at iteration *t*, then the mechanism announces n prices in the interval $[a_t, b_t]$ and runs the LP-TIOLI auction for the active players. SPSIA is, in essence, a combination of ASIA and LP-TIOLI. This is reflected in dominant strategy equilibrium for the game.

Proposition 3 *It is a weakly dominant strategy of any player i in SPSIA to bid 1 ("intention to participate") at iteration t as long as $v_i > b_t$ (which means that $LPI(h_i^t, v_i) = future$). Otherwise, the player should bid 0. If at any point the player is asked to choose from some set of prices she should pick the (announced) price which is just above her valuation.*

4.3 Optimization of Prices

As with the LP-TIOLI, the incremental mechanisms described above can easily be optimized to account for specific distributions over valuations. As well, like LP-TIOLI, the existence of simple dominant strategies that are independent of the price thresholds allow straightforward sequential optimization to be applied. In this section, we describe a Markov decision process (MDP) model that allows optimization of the prices for ASIA, and examine its performance empirically.

We assume a finite horizon T, after which round the auction will terminate if no unique highest bidder has been determined. Note, that this does not affect the dominant strategy equilibrium, since if the auction is forced to terminate at round T it will do so without an allocation. Assume a prior density f over valuations (for simplicity we assume this is the same for all bidders). We wish to set prices to optimize social welfare, possibly including the cost of communicating (or computing bids). We do this by formulating the optimization as an MDP as follows: states are pairs $\langle p, m \rangle$, where p

is the prior price threshold and m is the number of active bidders.[8] We use $c(m)$ to denote the communication cost as a function of the number of *active* bidders m.[9] At any stage $k < T$, and at any state $\langle p^{k-1}, m\rangle$ where $m > 1$, we can set any price $p \geq p^{k-1}$. If $m = 0$ or $m = 1$, the auction terminates. We define the optimal k-stage value function (reflecting expected welfare) EW^k and Q-function as follows:

$$EW^k(p^{k-1}, 0) = 0$$
$$EW^k(p^{k-1}, 1) = E_f(v|v > p^{k-1})$$
$$EW^k(p^{k-1}, m > 1) = \max_{p^k > p^{k-1}} Q^k(p^{k-1}, m, p^k) - c(m)$$
$$Q^k(p^{k-1}, m, p^k) = \Pr(1bid|p^{k-1}, m, p^k)EW^{k+1}(p^k, 1) + \sum_{2\leq n\leq m} \Pr(nbids|p^{k-1}, m, p^k)EW^{k+1}(p^k, n)$$

Here $\Pr(nbids|p^{k-1}, m, p^k)$ is the probability of exactly n bidders, among m bidders with valuations greater than p^{k-1}, having valuations greater than p^k.

We computed optimal policies for varying numbers of bidders and communication costs, assuming uniformly distributed valuations and using simple price discretization to keep the action space of the MDP finite.[10] These policies were examined, and the performance of ASIA was measured by simulating 100,000 runs of each optimized auction, using randomly drawn valuations for each bidder.

Some general observations verify expected behavior: for example, price thresholds within the same stage of a given auction are set higher when more bidders are active.

To demonstrate the benefit of computing a policy versus using a fixed price selection rule, we also simulated the performance of ASIA under the *Divide(k)* price update rule. We briefly state the definition of *Divide(k)* here; we defer motivation and further discussion to the next section.

Defn 3 Define the price update rule *Divide(k)* as follows. Given any fixed $k > 1$, an iterative mechanism operating under this rule would announce the price $p_0 = \frac{1}{k}$ at iteration 0. At any other iteration $t > 0$ the mechanism would announce $p_t = p_{t-1} + \frac{1-p_{t-1}}{k}$.

The parameter *k* was set in such a way as to force ASIA with *Divide(k)* to use the same average number of bits per player as ASIA with an optimized policy (assuming zero communication cost and uniform priors). We then compared the resulting loss in welfare for each price selection strategy. Figure 3 demonstrates that under the optimized policy, ASIA achieves a much smaller loss in welfare.[11]

Figure 4(a) shows the relation between the cost of communication (per bit) and the expected amount of communication per player until mechanism completion (each auction is optimized for the specific communication cost and maximum number of players). As expected, the amount of communication decreases as cost increases. Furthermore, per-player communication is lower with a greater number of players, since (in expectation) more players can be eliminated from the auction earlier. Figure 4(b) demonstrates the relation between communication cost and the loss in welfare (relative to the welfare optimal outcome). While this suggests that ASIA fares worse with increasing numbers of players, we can recast our results by considering loss in welfare as a function of the amount of per-player communication. Figure 4(c) demonstrates that, in fact, with more players ASIA requires *less* per-player communication to achieve the same loss in welfare. This gain is due to the fact that with a greater number of players, the policy can initially be more "aggressive," eliminating more players right away. Once only a few active players remain, the policy can become more conservative, which results in a higher chance of allocating the good. The sequential mechanism also appears to offer considerable improvement over the one-shot, LP-TIOLI mechanism with respect to welfare loss per bit of communication.

5. USING STOCHASTICITY TO GAIN EFFICIENCY

Blumrosen and Nisan [1] showed that, in the context of single-shot auctions, symmetric mechanisms—those that treat all players equally—are suboptimal relative to asymmetric mechanisms Their *priority games* incorporate tie-breaking rules based on a fixed ordering of the players. Essentially, this implies that the mechanism embodies a preference ordering over players, so that in the event that multiple players state the same "limited-precision desire" for the good, the mechanism always allocates it to the most preferred player. Thus the mechanism exhibits a certain unfairness.[12]

We have thus far constrained our attention to the mechanisms that are nondiscriminatory, with this notion taken to the extreme: whenever the mechanism ends without determining a unique highest bidder, the good remains unallocated. This is true for both *LP-TIOLI* and the incremental auctions.[13] For incremental auctions, this sort of allocative decision arises due to the *determinism* assumption made in the previous section. Since we only want to allocate the good to a player with the highest valuation, we have to disallow any sort of randomness in this decision. Although determinism (along with our other assumptions) imposes strict constraints on the form of incremental mechanisms, it clearly leads to a potential loss in efficiency and revenue.

5.1 Stochastic ASIA

To alleviate this problem we need to relax the assumption of allocating only to the unique highest bidder (given limited precision it might not always be possible to find this player). There are a number of ways of doing this. One would be to follow the style of *priority games* and introduce discrimination between players. An alternative approach is to relax the determinism assumption, but force the mechanism to be "fair" in expectation. This is the approach we consider here.

[8]The number of bidders is a sufficient characterization of state assuming identical priors. Bidder identities could be used in more general circumstances.

[9]In what follows we model only bidder communication costs, not those incurred by announcing prices. We could alternatively view these communication costs as a crude surrogate for the *computational costs* incurred by bidders when determining or refining their valuation to the required precision. More general forms of the function $c(m)$ are also possible.

[10]The discretization (with 1000 possible prices over $[0, 1]$) is largely for simplicity: methods for continuous state and action spaces could easily be adopted. Note also that approximation methods for solving MDPs will not impact the strategic properties of the mechanism produced.

[11]However, this fact should be considered in the context of the general performance of *Divide(k)*. As we show in the next section, *Divide(k)* exhibits some good properties which make it a feasible alternative to an optimized policy.

[12]Of course, should the priority ordering be determined randomly (and fairly) in advance, some of this concern may be mitigated.

[13]Although the rules of *LP-TIOLI* state that in case of a tie the good is offered randomly, any rational player will reject this offer; therefore it is outcome equivalent to forgo making any offers and not allocate the good.

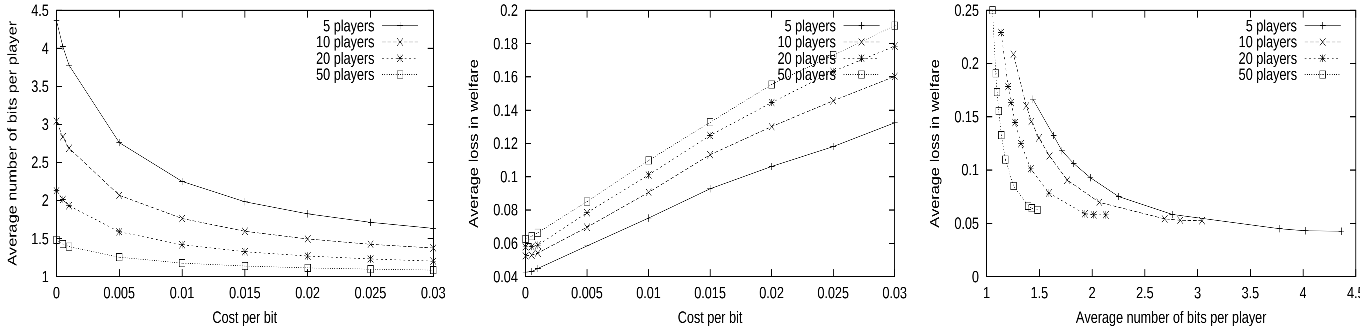

Figure 4: (a) Required bits per player as costs vary; (b)Average welfare loss as costs vary; (c) Average welfare vs. bits per player.

We propose the following variant of ASIA, which we call *Stochastic ASIA* or *STASIA*. STASIA has a dominant strategy equilibrium, but, as opposed to ASIA it always allocates the good. This mechanism operates as follows.

- As in ASIA, initially all the players are active. Denote the set of active players by A.
- At each iteration t, the mechanism randomly (and uniformly) *holds out* one of the active players, say, player i. It then announces price p^t to all active players, with $p^t \geq p^{t-1}$.
- All active players *except* the holdout player i are asked to reveal either 1, indicating a willingness to purchase the good for p_t (intention to participate), or 0, indicating a desire to become inactive.
- If all players in $A - \{i\}$ reveal 0, the good is allocated to *i* for price p^{t-1}.
- If at least one of the players in $A - \{i\}$ reveals 1, then *i* is also asked to reveal either 0 or 1 (with the same interpretation):
  - If only one active player (including the holdout) reveals 1, then this player receives the good and pays p^t.
  - If more than one active player (including the holdout) reveals 1, then the game moves into the next iteration (with only players revealing 1 remaining active).

Proposition 4 *In the STASIA mechanism, it is a weakly dominant strategy for any player i with valuation v_i (if given the opportunity bid) to bid 1 at iteration t as long as $v_i > p^t$, and to bid 0 otherwise.*

Proof sketch: Suppose player i is held out at stage t. Then the only time i is asked a to bid is when at least one other player has bid for the good at the current price p^t. If $v_i > p_t$, reporting 1 ensures a payoff of 0 or more; while reporting 0 will always result in payoff of 0. Alternatively, if $v_i \leq p^t$ reporting 0 ensures a payoff of 0, while reporting 1 will give a payoff less than or equal to 0. The case when i is not held out is similar. ◀

As was mentioned above, without the determinism assumption, the necessary conditions for the mechanism to have a dominant strategy equilibrium, described in the previous section, no longer hold. However, it is easy to see that STASIA belongs to the space of increasing price mechanisms. Furthermore, as with ASIA, the form of the dominant strategy is independent of the precise values of the price thresholds, which allows us to use a similar approach to mechanism optimization. Before proceeding to the performance results of STASIA with optimized price policies, we provide a general comparison of STASIA to the one-shot auctions.

5.2 Stochastic ASIA versus One-shot Auctions

In this section we analyze and compare the performance of two auction types under the assumption that the players' valuations are independently distributed according to the uniform distribution. The purpose of this section is to show that, with large and variable number of players, incremental auctions, specifically STASIA, can be significantly better than limited-precision, fixed-structure, threshold-based, one-shot auctions. Even though the analysis is performed under an unrealistic assumption of IID uniform distribution, it is sufficient for demonstrating that incremental auctions can be better than even the optimal one-shot auctions of Blumrosen and Nisan.

Consider the behavior of a *fixed structure* (the auction structure is known before the number of participants in announced) one-shot, limited-precision, threshold auction as the number of participants is increased. As more players enter the auction the expected highest valuation of this group rises. However, the expected welfare of a fixed structure one-shot auction will eventually start to fall or will stabilize at some fixed level (below the expected welfare of the optimal auction with unrestricted communication). This is very easy to see in the case of *LP-TIOLI*. As the number of players increases so does the probability of having a tie, which implies that eventually the probability of not allocating the good will become large enough to outweigh any increase in the expected highest valuation. This problem is not as pronounced in the case of priority games. However, even a priority game has the highest fixed price threshold, therefore as the number of players increases and the highest valuation of the group reaches this price threshold, the expected loss in welfare will be bounded above by a distance between this price threshold and its closest neighbor, a fixed positive constant. The actual expected loss is less than this upper bound (e.g., under the uniform distribution, it is nondecreasing in the number of players and equal to about half of this interval.)

Suppose we were to implement STASIA with *Divide(k)* price update rule.

Defn 4 Let m be an increasing price mechanism. We say m uses the *Divide(k)* price update rule, if for some $k > 1$: at iteration 0, it announces price $p_0 = \frac{1}{k}$; and at any other iteration $t > 0$, it announces $p_t = p_{t-1} + \frac{1-p_{t-1}}{k}$.

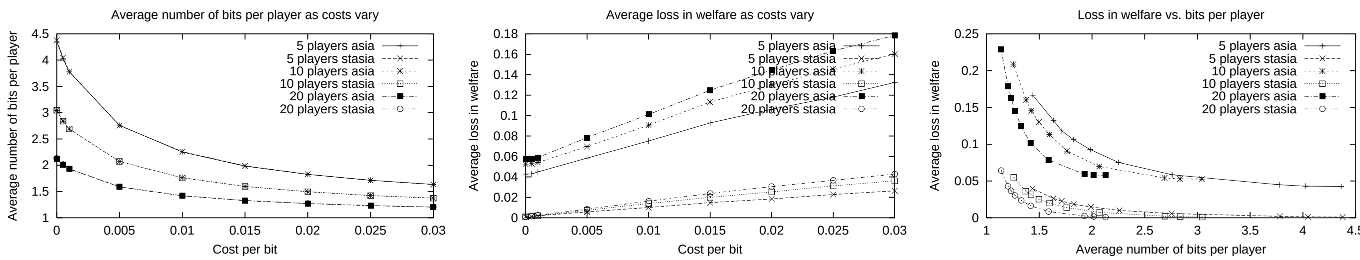

Figure 5: ASIA vs. STASIA under the same price policies (a) Required bits per player as costs vary; (b) Average welfare loss as costs vary; (c) Average welfare loss vs. bits per player.

One important feature of this update rule is the fact that it does not depend on the number of participating players (i.e., the mechanism structure is "fixed"). A second attractive property is that the distance between consecutive prices decreases with each iteration. Finally, the update rule also has another interesting property:

Proposition 5 *Suppose that all player valuations are drawn from a uniform density. The expected total number of bits sent by each player participating in STASIA, with an unrestricted number of iterations, is* less than k *when the* Divide(k) *price update rule is used.*[14]

Intuitively, implementing STASIA with *Divide(k)*, for any fixed k, and allowing it to run to termination (by finding the unique highest bidder or announcing a price above all valuations) would require each player to submit less than k bits in expectation, independent of the number of participants.

With this price update rule, if STASIA terminates at iteration t, then the loss in welfare is bounded by $p_t - p_{t-1}$. From the definition of *Divide(k)*, we have that $p_t - p_{t-1} < p_{t-1} - p_{t-2}$ and $\lim_{t\to\infty}(p_t - p_{t-1}) = 0$. This means that, as the number of players increases, the expected loss in welfare will fall. This follows from the fact that the expected value of the highest valuation increases with the number of players, and therefore the ties, if any, are expected to occur at later rounds. Note, that this occurs despite the fact that the expected number of bits submitted by each player remains unaffected.

As noted above, the loss in social welfare of any "fixed-structure" limited-precision, single-shot auction has a fixed, non-zero lower bound, independent of the number of players. Therefore, as the number of players is increased the expected loss of welfare of a fixed structure single-shot threshold-based mechanism will become greater than the expected loss of welfare of STASIA, while the total expected communication of STASIA is no greater than the total communication of the single-shot mechanism. This demonstrates that, even without price optimization, the types of mechanisms presented here offer advantages over one-shot mechanisms, even the asymptotically optimal mechanisms of [2].

5.3 Evaluating the Performance of STASIA

In this section we empirically evaluate the performance of STASIA. We begin by comparing the performance of STASIA to ASIA,

[14] While *Divide(k)* is "designed" for uniformly distributed valuations, it can be modified to work with any density, by selecting prices so that the distance between the each price "removes" $\frac{1}{k}$ of the remaining probability mass.

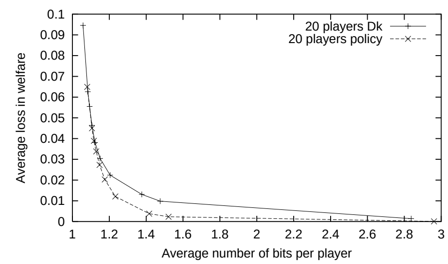

Figure 6: Comparison of loss in welfare in STASIA under optimized policy vs. STASIA under *Divide(k)* price update rule

with both auctions operating under the same price policies, specifically, the policies optimized for ASIA (the ones used in the previous section). We ran each mechanism using ten different price policies, with each policy generated by optimizing social welfare using a different "cost per bit" value. Each mechanism-policy pair was run 100000 times on randomly drawn player valuations.

Figure 5(a) demonstrates that the average number of bits sent by each player is approximately the same for both mechanisms. Since both mechanisms are using the same policies this is expected. Figure 5(b) shows that STASIA achieves a much lower loss in social welfare than ASIA. Figure 5(c) demonstrates the same result by presenting the loss in social welfare as a function of the average number of bits per player. These results stems from the fact that STASIA always allocates the good, and therefore the maximum loss in welfare it can incur is $p^t - p^{t-1}$ (when terminating at iteration t).

We use the same approach as in the previous section to produce new price policies optimized for STASIA. Figure 6 compares the performance of such optimized policies to the performance of STASIA under *Divide(k)* update rule. We see that while the optimized policy performs better, the difference is not as pronounced as with ASIA. This is primarily because being myopic under STASIA is not as bad as under ASIA, since the penalty of making an "incorrect" choice under ASIA is not allocating the good.

Finally, we compare the performance of STASIA under the price update policies optimized for ASIA to STASIA with price update policies optimized for STASIA. Figure 7(a) demonstrates that the policies are indeed different, with policies optimized for STASIA tending to use fewer bits of communication for any fixed cost per bit. Figure 7(b) shows that the policies optimized for STASIA do perform better than the policies optimized for ASIA.

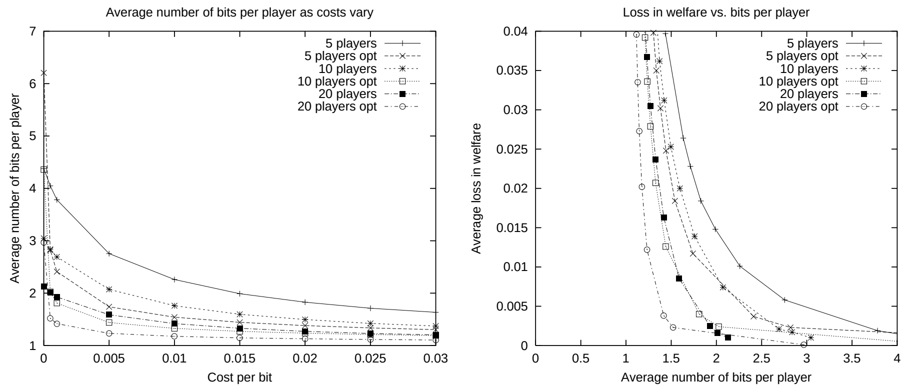

Figure 7: STASIA under different price policies (a) Required bits per player as costs vary; (b) Average welfare loss vs. bits per player.

6. CONCLUDING REMARKS

We have proposed a class of incremental, limited-precision auctions that extends earlier work on costly preference elicitation and communication. The main motivation behind this work was to propose a set of properties and then derive conditions that all mechanisms in possession of these properties must satisfy. By deriving these conditions we were able to demonstrate that all mechanisms belonging to this class have a similar structure. The proposed mechanisms have simple dominant-strategy equilibria, and use price thresholds that can be optimized to maximize social welfare, revenue or some other objective given a bound on communication, and can also be optimized to directly account for communication cost (and indirectly computational cost or related measures). Empirical results suggest that such mechanisms can find near-optimal allocations with very little communication (confirming the theoretical results of [1], though with a somewhat different mechanism); we also show that incremental mechanisms offer advantages over one-shot auctions.

We hope to extend the study of limited and incremental revelation mechanisms to more complex settings, such as multi-attribute and combinatorial auctions, and multi-attribute bargaining problems. Indeed, it is in these complex settings where limited, incremental revelation will be critical.[15] We further plan to study more deeply the design of mechanisms that incorporate cost of communication and computation in their design objectives.

Acknowledgements

Thanks to Moshe Tennenholtz and Tuomas Sandholm for extremely valuable discussions on this and related topics. Thanks also to Vincent Conitzer, David Parkes and Yoav Shoham for their comments. This research was supported by the Institute for Robotics and Intelligent Systems (IRIS) and the Natural Sciences and Engineering Research Council (NSERC).

7. REFERENCES

[1] Liad Blumrosen and Noam Nisan. Auctions with severely bounded communication. In *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science (FOCS 02), Vancouver, Canada*, pages 406–416, November 2002.

[2] Liam Blumrosen, Noam Nisan, and Ilya Segal. Multi-player and multi-round auctions with severely bounded communication. In *11th Annual European Symposium on Algorithms (ESA 03), Budapest, Hungary*, September 2003.

[3] Wolfram Conen and Tuomas Sandholm. Partial-revelation VCG mechanism for combinatorial auctions. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence*, pages 367–372, Edmonton, 2002.

[4] Vincent Conitzer and Tuomas Sandholm. Computational criticisms of the revelation principle. Unpublished manuscript, 2004.

[5] Vincent Conitzer and Tuomas Sandholm. Complexity of mechanism design. In *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence (UAI 02), Edmonton, Canada.*, pages 103–110, August 2002.

[6] Yuzo Fujishima, David McAdams, and Yoav Shoham. Speeding up ascending-bid auctions. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, pages 554–563, Stockholm, 1999.

[7] Elena Grigorieva, Jean-Jacques P. Herings, Rudolf Müller, and Dries Vermeulen. The private value single item bisection auction. Research Memorandum 51, METEOR, Maastricht, 2002.

[8] Kate Larson and Tuomas Sandholm. Costly valuation computation in auctions. In *Proceedings of the Eighth Conference on Theoretical Aspects of Rationality and Knowledge*, pages 169–182, Sienna, 2001.

[9] Kate Larson and Tuomas Sandholm. An alternating offers bargaining model for computationally limited agents. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 135–142, Bologna, 2002.

[10] Andreu Mas-Colell, Micheal D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.

[11] Roger Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

[12] David C. Parkes. iBundle: An efficient ascending price bundle auction. In *Proceedings of the ACM Conference on Electronic Commerce, Denver (EC 99)*, pages 148–157, November 1999.

[13] David C. Parkes. Auction design with costly preference elicitation. *Annals of Mathematics and AI*, 2004. to appear.

[14] Tuomas Sandholm and Andrew Gilpin. Sequences of take-it-or-leave-it offers: Near-optimal auctions without full valuation revelation. In *Proceedings of the AAMAS-03 Workshop on Agent-Mediated Electronic Commerce V (AMEC V), Melbourne, Australia*, July 2003.

[15] Yoav Shoham and Moshe Tennenholtz. Rational computation and the communication complexity of auctions. *Games and Economic Behavior*, 35:197–211, 2001.

[16] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[15] Note that ASIA can be extended to multi-unit auctions (multiple units of a single good) in a straightforward fashion, while STASIA can also be applied with slightly increased communication.