Envy Quotes and the Iterated Core-Selecting Combinatorial Auction*

Abraham Othman and **Tuomas Sandholm**
Computer Science Department
Carnegie Mellon University
{aothman,sandholm}@cs.cmu.edu

Abstract

Using a model of agent behavior based around envy-reducing strategies, we describe an iterated combinatorial auction in which the allocation and prices converge to a solution in the core of the agents' true valuations. In each round of the iterative auction mechanism, agents act on *envy quotes* produced by the mechanism: hints that suggest the prices of the bundles they are interested in. We describe optimal methods of generating envy quotes for various core-selecting mechanisms. Prior work on core-selecting combinatorial auctions has required agents to have perfect information about every agent's valuations to achieve a solution in the core. In contrast, here a core solution is reached even in the private information setting.

Introduction

The *Vickrey-Clarke-Groves mechanism (VCG)* is ubiquitous in theoretical mechanism design. In the standard private-values setting, it is the revenue-maximizing mechanism among all incentive compatible individually rational efficient mechanisms [Krishna and Perry, 1997]. Unfortunately, in addition to being unwieldy to implement in practice, VCG suffers from a number of pathologies [Rothkopf et al., 1990, Sandholm, 2000, Ausubel and Milgrom, 2006, Rothkopf, 2007]. These include revenue non-monotonicity, in which adding another bidder can lower the seller's revenue, and receiving an arbitrarily small fraction of the revenue achievable by posting prices.

To obtain higher revenues than VCG, one can explore inefficient mechanisms, which leads to combinatorial generalizations of the revenue-maximizing single-item auction [Myerson, 1981]. Revenue-maximizing mechanisms are unknown even for the (unrestricted) two-item setting, and in general a concise description of the revenue-maximizing combinatorial auction cannot exist (unless P=NP) because that design problem is NP-complete [Conitzer and Sandholm, 2004]. Some work has been done on automated mechanism design for finding high-revenue combinatorial auctions, but those approaches have not been used for large numbers of items [Likhodedov and Sandholm, 2004, 2005]. Even simple revenue-enhancement approaches like setting reserve prices require good knowledge of a prior distribution over agent valuations, which may or may not be available depending on the setting.

A different approach is to relax incentive compatibility. One recent stream of research in non-incentive compatible mechanisms has involved *core-selecting combinatorial auctions*. Mechanisms of this class were used in the recent successful spectrum license auction in the United Kingdom [Cramton, 2008a,b, Day and Cramton, 2008]. These mechanisms mitigate the poor revenue properties of the VCG mechanism without necessarily subscribing to a first-price mechanism, which motivates significant underbidding. Selecting an outcome in the core yields a host of desirable properties that VCG lacks [Parkes, 2002], such as revenue monotonicity and resistance to bidding using multiple pseudonyms [Yokoo, 2006].

The word *iterative* has taken on a confusing double meaning in combinatorial auction research. On the one hand, the auction process itself can be iterative, in which bids are solicited in a series of rounds until a termination condition is reached (for instance, no agent submits a new bid). This is the concept we explore in this paper. On the other hand, given a set of bids, a solution may be produced iteratively, for example, by raising the price of bundles in ascending rounds in a specific way until reaching a point in the core. Examples of this latter process include Parkes [1999], Ausubel and Milgrom [2002], Wurman et al. [2004], and Hoffman et al. [2006]. Unfortunately, these techniques tend to be too slow to be used in an explicitly multi-round auction, so these mechanisms work by inputting valuations into proxy agents that bid on behalf of the auction participants. Thus, price increases are a function of iterating on the bids of these proxy agents rather than multiple iterative rounds of buyers changing their valuations. As a consequence these auctions are essentially one-shot. Thus they do not achieve the main benefit of auctions that elicit the bidders' offers via multiple interactions: multi-interaction mechanisms in effect give feedback to each bidder regarding what information is (not) needed from her, thus reducing the bidders' deliberation effort in determining their own valuations. For a review of preference elicitation in combinatorial auctions see Sandholm and Boutilier [2006].

One attempt to get around the inability to solve for a core solution quickly is the clock-proxy auction [Ausubel et al., 2006]. It maintains (fast-to-compute) linear prices over items through a number of bidding rounds before solv-

*This work was supported by NSF grants IIS-0427858 and IIS-0905390.

ing a single final core-selecting round.

It would be ideal, however, to conduct a series of fully core-selecting rounds. Recent computational work by Day and Raghavan [2007] has shown that constraint generation can be used to clear large-scale core-selecting combinatorial auctions orders of magnitude faster than previously possible. We argue that this makes multi-interaction core-selecting combinatorial auctions practical. In this paper we introduce and study such a mechanism.

Prior work by Milgrom [2006] and Day and Milgrom [2008] has studied some of the revenue and strategic properties of core-selecting combinatorial auctions. That literature has examined the case where bidders have perfect information about their own valuations and the valuations of others. They analyze a mechanism under which a bidder-optimal (i.e., revenue-pessimal) point in the core is selected. They emphasize that, with perfect information, this will be the outcome if side-payments between players are allowed in any core-selecting mechanism. They show that agents will adopt a set of strategies that involve *truncation*, that is, each agent shaves all her bids by a certain individual value.

Unfortunately, the perfect information assumption is seldom realistic in practice. One of the primary reasons for running an auction in the first place—as opposed to, for example, simply posting prices—is that there usually is considerable uncertainty about agents' valuations. It is a poor modeling choice to assume the auctioneer has no prior information about the agents' valuations while at the same time assuming that every agent exactly knows the private valuation of every other agent. A contribution of our paper is to provide a mechanism with all the benefits of a core-selecting combinatorial auction in the more plausible setting where bidders' valuations are private information.

Core-selecting combinatorial auctions are not (generally) incentive compatible. This is easy to see because they are efficient and individually rational, yet promise higher revenue than the VCG (but the VCG is the highest-revenue efficient individually rational incentive compatible mechanism [Krishna and Perry, 1997]).[1] (In parallel to our work, Goeree and Lien [2009] have shown non-incentive compatibility of core-selecting combinatorial auctions using a different argument.) So, in lieu of traditional incentive compatibility, a different model of bidder behavior is needed.

In this paper we study core-selecting combinatorial auctions where bidders use envy-reducing strategies. We begin by formally establishing the properties of envy-freeness and envy-reduction, and show that a straightforward approach to applying these techniques to an iterated core-selecting combinatorial auction fails to produce an outcome in the core. We then introduce the driving concept behind this paper, *envy quotes*, which serve as estimates of the current prices of bundles the agent is losing. Employing these envy quotes in an iterated setting, we prove that when agents act to reduce their envy of their envy quotes, the solution converges to the core of the agents' true valuations. Then, we discuss what happens when agents do not behave according to our prescribed model, showing that our mechanism has desirable safeguards against low-revenue outcomes. Finally, we discuss how to generate these envy quotes in practice, showing a general method that works with any core-selecting combinatorial auction, a computationally inefficient optimal (in a sense defined later) technique for any core-selecting combinatorial auction, an optimal technique for any bidder-optimal core-selecting combinatorial auction, and an optimal technique for the bidder-pessimal core-selecting combinatorial auction (which is equivalent to a first-price mechanism).

[1] In special cases of agent valuations that cause the VCG outcome to be in the core, and where a bidder-optimal core-selecting mechanism is used, the mechanism is, of course, incentive compatible.

Theory of envy-freeness

There exists a set of n agents and k items, and therefore $2^k - 1$ bundles. As is normal, we assume that agents have quasilinear utility, so that an agent's utility for receiving allocation a and paying π is $u(a) - \pi$.

Definition 1. *A (feasible)* outcome *is a set of allocations $a_1, \ldots, a_n$ and payments $\pi_1, \ldots, \pi_n$, where $a_i \cap a_j = \emptyset$ for $i \neq j$.*

Definition 2. *A* blocking coalition *is a group of agents G who can propose an alternate outcome (with allocations $a'_1, \ldots, a'_n$ and payments $\pi'_1, \ldots, \pi'_n$) in which only members of G win a bundle, such that for all $i \in G$, $u_i(a'_i) - \pi'_i \geq u_i(a_i) - \pi_i$ where the inequality is strict for at least one i, and $\sum_i \pi'_i > \sum_i \pi_i$.*

Definition 3. *An outcome is in* the core *if it induces no blocking coalitions.*

Definition 4. *An outcome is* efficient *if no other feasible outcome has a higher social welfare (sum of utilities) for the participants.*

It follows that every outcome in the core is efficient, because any inefficiencies would yield a blocking coalition (c.f. Shapley and Shubik [1971], Day and Raghavan [2007]).

Definition 5. *Agent i (with allocation a_i and payment π_i)* envies *agent j (with allocation a_j and payment π_j) if $u_i(a_i) - \pi_i < u_i(a_j) - \pi_j$, where $(u_i(a_j) - \pi_j) - (u_i(a_i) - \pi_i)$ is the* amount of envy *(or just* envy*).*

Definition 6. *An agent plays a (myopic)* envy-reducing strategy *if, given a set of reports of the other agents, she modifies her type report to reduce her envy of some agent without lowering her utility. A group of agents play a (myopic)* group envy-reducing strategy *when no agent in the group lowers their utility and at least one agent changes her bid to reduce her envy of another agent.*

Definition 7. *An outcome is* envy-free *if no subset of agents prefers the allocation-payment pair of any other subset of agents to its own allocation-payment pair. If both subsets are restricted to consist of only individual agents, then we call that set* individually envy-free.

The set of individually envy-free points is at least as large as the set of envy-free points because the concept is less restrictive.

Definition 8. *An* envy-free fixed point *is a fixed point of a system where groups of agents follow envy-reducing strategies.*

Corollary 1. *Every envy-free fixed point is in the core (with respect to* true *valuations).*

Corollary 2. *Every envy-free fixed point is efficient (with respect to* true *valuations).*

Lemma 1 (Leonard [1983])**.** *The revenue from VCG does not exceed that of the bidder-optimal (i.e., revenue-pessimal) outcome in the core.*

Corollary 3. *Every envy-free fixed point delivers at least as much revenue as the VCG mechanism.*

The path forward seems straightforward: Simply have agents iterate in an envy-reducing manner towards an envy-free solution in the core. This would yield desirable revenue properties while only requiring agents to have private information. However, as we show in the next section, the combinatorial nature of the problem complicates a simple model of individual envy-reduction.

Individual envy-reduction is insufficient

The following example shows how individual envy reduction can be insufficient for reaching a fixed point in the core.

Example 1. *Consider a two-item three-bidder problem, where bids (b_i) and true valuations (v_i) are given by the following table:*

Bundle	*Bid*	*Valuation*
A	$b_1 = 5, b_2 = 0, b_3 = 0$	$v_1 = 9, v_2 = 0, v_3 = 0$
B	$b_1 = 0, b_2 = 5, b_3 = 0$	$v_1 = 0, v_2 = 9, v_3 = 0$
AB	$b_1 = 5, b_2 = 5, b_3 = 15$	$v_1 = 9, v_2 = 9, v_3 = 15$

Every core solution awards item A to bidder one and item B to bidder two at a total price of between 15 and 18. But with these valuations and bids, every core-selecting combinatorial auction awards both items to the third agent at a price in [10,15]—a solution that is not in the true core. However, no agent individually *envies the allocation of any other: neither losing bidder would prefer to pay 10 for AB.*

Iterative combinatorial auctions with envy quotes

The problems illustrated by the above example arise because we are not properly expressing to an agent how her bid impacts the combinatorial nature of the allocation. As we have discussed, envy-free dynamics in the combinatorial setting only work when groups of agents work together. But this is undesirable, because it encourages collusion among bidders (which could lead to bad outcomes) or could be illegal, which is the case in many public goods auctions [Day and Raghavan, 2007]. In this section, we explore how to present *envy quotes* to agents regarding the clearing prices of bids they have lost, such that the envy quotes meaningfully reflect what actual prices are. In effect, we are changing the target of agents' envy from distinct bidders to the winners' prices on bundles in which the agent is interested. Furthermore, we show that the fixed point of individuals reducing their envy on these quotes is in the core with respect to true valuations.

Our iterative scheme

We propose the following process:

1. Solicit bids from agents.
2. Compute current winners and payments according to some core-selecting combinatorial auction.
3. For each of her losing bids, an agent receives an envy quote, $p(S)$, in the form of "The bundle S is currently going for price $p(S)$".
4. Repeat steps 1 through 3 until no new bids are received.

When we provide an envy quote to an agent on the price of a bundle, we have competing objectives. On the one hand, the envy quote has to be low enough so that it does not cut into the core (which could lead to agents not envying outcomes they should legitimately envy). On the other hand, an envy quote should not be lower than the agent's bid on a bundle, in order to reflect the core-selecting nature of the mechanism. This leads us to the following definitions:

Definition 9. *An agent's* core support $c(S)$ *for a bundle S she is losing is the highest bid she could submit and not change the current allocation or prices.*

It follows that the core support $c(S)$ is always less than the agent's quote on S, that is, the amount she would need to win the bundle.

Definition 10. *Let an agent bid $b(S)$ for bundle S and not win that bundle. The* envy quote $p(S)$ *satisfies $b(S) \le p(S) \le c(S)$.*

Now we can define what envy means in the envy quote context.

Definition 11. *Let an agent currently be winning bundle w at price π_w. She* envies the envy quote $p(S)$ *on a bundle she is losing, S, if $u(S) - p(S) > u(w) - \pi_w$.*

Proposition 1. *If the current solution in the iterated core-selecting combinatorial auction is not in the core (with respect to true valuations), then some agent has a bid that reduces her envy of the envy quote she receives on at least one bundle.*

Proof. Assume we are in a non-core state such that allocations are given by $a_1, \ldots, a_n$ and prices by $\pi_1, \ldots, \pi_n$. We will show that some agent has envy in this state, and that she has a bid to reduce that envy.

Since the solution is not in the core, there exists some blocking coalition in which at least one member of the coalition is strictly better off. Call that new solution $a'_1, \ldots, a'_n$, with prices $\pi'_1, \ldots, \pi'_n$. Without loss of generality, let agent 1 be strictly better off. We have $u_1(a_1) - \pi_1 < u_1(a'_1) - \pi'_1$, where the left hand side is non-negative, and $a'_1 \neq \emptyset$ by the restriction that no agent bids above her valuation. We will show that in the initial state, agent 1 is presented with an envy quote that induces envy. Let the envy quote of a'_1 in the original state be p. To show agent 1 has envy for a'_1, we must have $u_1(a_1) - \pi_1 < u_1(a'_1) - p$, for which it is sufficient to show that $p \le \pi'_1$, which holds because envy quotes are always less than quotes. As an example of an envy-reducing strategy, the agent can increase her bid on a'_1 to $p + \epsilon$, where

$0 < \epsilon < (u_1(a_1') - p) - (u_1(a_1) - \pi_1)$, because either the agent's next envy quote on a_1' must be higher so her envy of it is reduced, or she wins the item at a price of at most $p + \epsilon$, which gives her more utility. ■

Envy-reducing dynamics converge

In this section, we show that if agents respond to their envy quotes on items they are not winning, then prices to converge to a fixed point. Furthermore, it suffices that agents select such envy-reducing actions with positive probability. This kind of convergence result is standard in the matching market literature (c.f. Roth and Vande Vate [1990]), and holds regardless of the path taken to the current set of prices.

Proposition 2. *Assume that bids must be from a finite set of discrete levels, where the difference between consecutive levels is at most ϵ (e.g., bids are in integer dollars). If at every state of the auction at least one agent has positive probability of selecting an action among those that reduce her envy of an envy quote on a bundle she is not winning, then any iterated core-selecting combinatorial auction converges to a fixed point with probability 1. In this fixed point, each agent's envy of the envy quote on any bundle is at most ϵ, each coalition's payoff is no less than its pessimal core payoff minus $|coalition| \cdot \epsilon$, and revenue is no less than that of the revenue-pessimal point of the core minus $n \cdot \epsilon$.*

Proof. The proof of Proposition 1 can be extended trivially to show that in the fixed point, no agent can have more than ϵ envy on any bundle, and that an ascending bid exists otherwise. Furthermore, the largest envy quotes on bundles of some core outcome are no smaller than the agents' bids in that core outcome minus ϵ. If this were not the case, some agent could reduce her envy by bidding ϵ higher for some bundle. Since some agent has an ascending bid at every non-fixed point, by assumption the agent will select such an action with positive probability. Since there are only a finite number of (agent, bundle, bid level)-triples, we can construct a finite number of steps to reach a state in which no agent has envy greater than ϵ of her envy quote on any bundle. Thus, these envy-reducing dynamics converge to such a fixed point with probability 1.

Suppose some coalition's payoff is less than its lowest core payoff minus $|coalition| \cdot \epsilon$. Then some agent in the coalition has at least ϵ envy of the core state. Because envy quotes are smaller than quotes, it follows that the agent has at least ϵ envy of the envy quote of the bundle she would receive in that state. Thus we are not in a fixed point, which contradicts our premise. Therefore, each coalition's payoff is at least its lowest core payoff minus $|coalition| \cdot \epsilon$.

Let r denote the revenue in the fixed point we reach. Let r_c denote the revenue from the core solution that we are near (not necessarily the revenue-pessimal core solution), and let $\overline{p_i}$ represent the largest envy quote received by any agent for agent i's bundle in the core solution we are near. Because the mechanism is core-selecting with respect to reported bids, we have $r_c \geq \sum \overline{p_i}$. Since we are in a fixed point, we must have $r \leq \sum(\overline{p_i} + \epsilon)$, because were this not the case, an agent would have an envy-reducing play by bidding ϵ higher and thereby forcing the core solution. Letting $r_{pessimal} \leq r_c$ represent the revenue-pessimal core solution, it follows that $r_{pessimal} - n \cdot \epsilon \leq \sum \overline{p_i} \leq r$. Therefore, our revenue is no smaller than that of the revenue-pessimal core outcome minus $n \cdot \epsilon$. ■

Robustness of the approach

As we argued in the introduction, core-selecting combinatorial auctions are not incentive compatible. Abandoning incentive compatibility comes with a host of strategic concerns. How can we say how agents will play if they do not play truthfully? In this section, we explore the robustness of our mechanisms: what happens when agents either bid too little or too much, or attempt to otherwise manipulate the mechanism in ways that could be to their benefit. That is, what happens when agents fail—either due to willful manipulation or incompetence—to decrease their envy of envy quotes?

The optimal manipulation, if agents had perfect information and side payments were allowed, would be for each to shave her bid in a specific manner in order to achieve the core solution that minimizes the sum of the payments by the agents. This solution will coincide with the VCG solution if the VCG solution is in the core; this is the solution concept featured by Ausubel and Milgrom [2002] and Day and Raghavan [2007]. But in our setting, agents do not have perfect information, and the setting is iterated.

An agent can, of course, shave her bid too much. This is the great fear of running a non-incentive compatible mechanism—that agents, recognizing that they should shave their bids, will bid very little and the end result will be low revenue. As we show, however, we do not necessarily need to rely on agents being motivated only by envy to achieve solutions in the core.

Proposition 3. *If revenue is less than the revenue-minimal point in the (true) core, then some agent can make a new bid on a bundle she is losing that will myopically increase (or maintain) her utility.*

Proof. If revenue is less than the revenue-minimal point in the core, then some agent has an envy-reducing bid on an envy quote she is receiving on a losing bundle. Moreover, since envy quotes may coincide with a quote of the actual value she would need to pay to win the bundle, bidding in response to such a report can increase myopic utility. Specifically, if she captures the bundle at the higher price she will have higher utility, and if she fails to capture the bundle she will be no worse off. ■

Agent rationality provides a strong argument for our mechanism not returning a low revenue solution. On the other hand, if an agent shaves her bid too little, there might not be any straightforward way to achieve a better outcome for her. There may be a multitude of core outcomes with higher revenue, in which agents bid more than in the bidder-optimal solution. If an agent insufficiently shaves her bid, the mechanism might arrive at such a state. Since that outcome is in the core, only a global effort by a grand coalition of agents can force an outcome where the agents pay less.

In summary, our mechanism handles agents' mistakes in a revenue-optimizing way. If agents bid too low, self-interest will compel them to correct their bids. If agents bid too high, the structure of the core can lock agents into a high-revenue core solution that no agent can escape.

Another concern if agents do not play optimally is that, in an iterated setting, agents will move around in the state space of possible allocations, attempting to find advantageous outcomes in which other agents make errors that are beneficial to the agent that is causing the moving. One way of dealing with this possibility is to ignore it; as we have shown, only efficient outcomes can emerge as the fixed points of our iterated mechanisms. Therefore, the only way an agent will be able to take advantage of an inefficient outcome that yields low revenue for the auctioneer is for some other agent to not make an envy-reducing—and utility non-reducing—move that is made apparent to her by her envy quotes. Another possibility is to add an ascending clock to the auction, such that bids can only increase. This does not impact our convergence result, which relies on agents increasing their bids on bundles they are not winning.

Comparison with the ascending proxy auction

In this section, we discuss the differences between our approach and the iterative core-selecting combinatorial auction of Ausubel and Milgrom [2002], the *ascending proxy auction (APA)*.

The most important difference involves bidder strategies. The APA mandates that bidders behave in a specific way to achieve a solution in the core, namely, that agents raise their losing bids by small ϵ in each iterative round. Since that elicitation process is slow and there is no guarantee agents will behave this way in practice, Ausubel and Milgrom [2002] suggest that agents surrender their valuations to a proxy agent that bids in this manner on their behalf. Such a scheme loses out on perhaps the most important part of having an iterative auction in the first place: reducing the bidders' valuation deliberation efforts [Sandholm and Boutilier, 2006] (and potentially yielding higher revenue by making bidders feel more secure in their valuations [Cramton, 1998]). Furthermore, even with proxy agents, the APA is a very slow way to calculate a core solution from a set of bids [Day and Raghavan, 2007].

Essentially, what envy quotes provide is a more efficient way of conveying price information to losing agents. As we discuss in the next section, we can formalize this argument. If an *optimal* scheme for generating envy quotes with a particular core-selecting mechanism is used, price information is conveyed to losing bidders as efficiently as possible. In contrast, the APA uses the *least* efficient scheme for conveying prices to losing bidders.

Additionally, our approach works with any core-selecting combinatorial auction, while the APA was designed only for the bidder-optimal core-selecting combinatorial auction. Furthermore, the APA relies on an ascending global clock for setting prices while our approach assumes neither a global clock nor ascending prices.

Computing envy quotes

In this section we introduce the concept of optimal envy quotes and show how to compute optimal envy quotes for different core-selecting mechanisms.

Trivial envy quotes for any core-selecting combinatorial auction

One simple but valid scheme for producing envy quotes is to give an agent that is losing bundle S with a bid of b the envy quote of b. This trivially satisfies the envy quote definition. It is also fundamentally equivalent to the APA, because using this scheme the only feedback a bidder receives on a bundle she is losing is that she is losing the bundle at her current bid, and that a bid of ϵ higher may or may not win the bundle.

Recalling that envy quotes are bounded from above by the core support $c(S)$ (Definition 9), we have the following definition of what kinds of envy quotes accelerate the auction the most:

Definition 12. *A method for generating envy quotes is* optimal *for a core-selecting combinatorial auction if it always generates the largest possible envy quotes.*

It is never possible for an agent following an optimal envy quote to make an uncompetitive bid that does not change the outcome. Therefore, any such action causes a state change in the system. This is generally not the case with suboptimal envy quotes and thus the process with such quotes can take more steps to terminate. Conveying the best possible information about current prices to bidders allows them to make the best decisions about what bundles to bid on and how much to bid on them.

Optimal envy quotes for any core-selecting combinatorial auction

Intrinsically, the use of a core-selecting auction implies that, for each losing bid, there exists some threshold value (the optimal envy quote) above which the allocation and/or payments change and below which they do not. Because of this property, we can solve for the threshold value by treating the core-selecting process as a black box and using binary search to find that threshold. Letting v^* represent the sum total of the accepted bids of the current solution, we begin on the search interval $[0, v^* + \epsilon]$, and query as to whether the midpoint of the interval changes the current solution. If so, the midpoint becomes the new upper bound, and if not, the midpoint becomes the new lower bound. Each time this process is run, we produce an additional bit of accuracy with respect to finding the optimal envy quote. Because this approach treats the core-selecting mechanism as a black box, this process works with any core-selecting combinatorial auction.

However, this approach is likely to be slow because it must be run multiple times for each losing bid. To address this shortcoming, in the following sections we develop techniques to solve for optimal envy quotes in a single step for the two common kinds of core-selecting combinatorial auction.

Optimal envy quotes for bidder-optimal core-selecting combinatorial auctions

The most common core-selecting combinatorial auctions select a bidder-optimal (i.e., revenue-pessimal) point in the core (as defined by reported, rather than actual, valuations). Those methods vary depending on which bidder-optimal core solution is chosen [Day and Raghavan, 2007, Day and Cramton, 2008, Erdil and Klemperer, 2009].

We now describe a method for generating optimal envy quotes for any bidder-optimal core-selecting combinatorial auction. To construct an optimal envy quote for a bundle, we use a mixed integer program (MIP) to find the smallest price for that bundle that could be used to construct a blocking coalition. To compute an envy quote for agent i for bundle S, we do the following:

1. *Add each winners' accepted bid constrained by the price paid by the agent into the MIP*. For example, if an agent wins the bundle ABC at a price of 20, we add the constraint $\pi_A + \pi_B + \pi_C = 20$.
2. *For each bidder $j \neq i$, calculate j's surplus (i.e., reported valuation minus price), subtract it from j's losing bids, and add those revised losing bids as constraints into the MIP using the sum of item prices*. To illustrate this, imagine a winning agent with surplus 3 that has a losing bid for the bundle AB for a price of 10. This would be added as $\pi_A + \pi_B \geq 10 - 3$.
3. *Add all of i's losing bids as constraints, without subtracting i's surplus*. For example, if i has a losing bid on AB for a price of 5, then we add the constraint $\pi_A + \pi_B \geq 5$.
4. *For each agent, use binary variables and a constraint on them to ensure that at most one of the agent's losing bids is selected.*[2] This ensures that any blocking coalition of bids involves at most one bid from each agent. Thus none of the constraints from Step 2 or 3 cut into the core.
5. *The objective of the MIP is to calculate the lowest bundle price (as sum of item prices) for S*. For instance, consider bundle AB. Our objective function is $\min \pi_A + \pi_B$. This is equivalent to designating the envy quote for a bundle along a hyperplane normal to the items in the bundle, and is therefore more flexible than imposing only item prices.
6. *The solution to the objective, the minimum bundle price, is given to agent i as an envy quote for bundle S.*

Proposition 4. *The above method generates optimal envy quotes for any bidder-optimal core-selecting combinatorial auction.*

Proof. Suppose our method produces an envy quote for a bundle S of p. We will show that a bid of $p + \epsilon$ will change the mechanism's allocation-payment pair.

The constraints in the MIP that we use to generate the envy quote define the set of all possible blocking coalitions involving an agent's bids exactly. Since the combinatorial auction selects a bidder-optimal point in the core, this implies that the addition of a hyperplane corresponding to the items of S, $\sum_{x_i \in S} \pi_i = p + \epsilon$ cuts into the core, which implies the existence of a blocking coalition at the current set of prices based on the bid $p + \epsilon$. Because the mechanism selects a bidder-optimal point in the core, either the bid $p + \epsilon$ would be part of a new winning coalition or it would change the prices. ∎

One can make this MIP into an anytime algorithm by first adding the agent's losing bids as constraints, and then adding the constraints from the other agents into the MIP incrementally. After the optimization of every such addition, we have a valid envy quote, and those quotes increase as additional constraints are added. However, terminating the optimization prematurely will obviously not yield optimal envy quotes in every case.

Our method of generating envy quotes works with any bidder-optimal core-selecting combinatorial auction. Different allocations and different individual payments may be the outcome according to the different bidder-optimal core-selecting combinatorial auctions. This is captured in the above envy quoting method as different ways of calculating the surpluses in Step 2.

According to our Proposition 2, our process converges to some core solution of the true valuations. Furthermore, if no agent bids more than her true valuation, some *bidder-optimal* point in the true core is reached. This can be seen by analyzing two cases. First, if the current solution has revenue less than the bidder-optimal revenue in the true core, some agent has an envy-reducing move on an envy quote according to our Proposition 1, and thus the auction has not terminated. Second, if the current solution has revenue greater than the bidder-optimal revenue in the true core, some bidder is bidding more than her true valuation. This is because if agents bid truthfully, a bidder-optimal point in the true core is reached and revenue is nondecreasing in reported bids (given that a solution in the core of the reported valuations is chosen).

Optimal envy quotes for the first-price combinatorial auction

We can also develop optimal envy quoting methods for other core-selecting mechanisms. As an extreme (bidder-pessimal) example, consider a first-price mechanism, which charges winning bidders their bids, and selects the allocation that will maximize revenue. It is trivially core-selecting as long as no bidder bids above her true valuation.

For this setting, the optimal envy quote scheme is to produce the quote for each losing bid, that is, the threshold value at which the agent would go from losing to winning the bundle in question. Any lower envy quote is not optimal, because an agent could bid higher without altering the allocation-payment pair, and no envy quote can ever be higher than a quote. Solving for these quotes is NP-hard [Sandholm, 2002], but can usually be done fast in practice.

Conclusions

Core-selecting combinatorial auctions have recently emerged as an important direction both in theory and in practice. Core solutions avoid many pathologies of

[2] This can be implemented relatively efficiently in a MIP using the SOS1 construct.

VCG, including poor revenue, revenue non-monotonicity, and vulnerability to bidding with multiple pseudonyms. Prior models of core-selecting combinatorial auctions have either been too slow to yield truly iterative auctions and/or assumed agents know each others' valuations.

In this paper, we expanded upon previous advances by designing a truly iterative core-selecting combinatorial auction that does not make any assumptions about what agents know about each others' valuations. We showed that when agents as a group follow envy-reducing strategies, the resulting fixed point is in the core of true values. However, we demonstrated that agents acting individually to reduce their envy would be insufficient to arrive at a fixed point in the core, because of the combinatorial nature of the auction.

To remedy this problem, we developed a system of *envy quotes*, where agents are given estimates on their losing bids of the prices at which those bundles are being won. We proved that if agents act to reduce their envy of their envy quotes, an outcome in the core of true values is reached. Moreover, we discussed some of the safeguards our mechanism has for when agents do not behave this way: our method may enhance revenue via agent mistakes and the mechanism cannot reach an allocation-payment pair where agents pay too little without some agent having a clear potentially utility-increasing (and not merely envy-reducing) bid.

Finally, we developed various techniques for generating envy quotes. The trivial method of returning a losing agent's bid is valid for *any* core-selecting combinatorial auction, but not optimal in the sense that it will generate unnecessary iterations. This slowly-converging method is equivalent to (the truly iterative version of) the ascending proxy auction of Ausubel and Milgrom [2002]. We point out how binary search can be used to produce optimal envy quotes in general. However, that requires multiple rounds of optimization for each quote. To address this shortcoming, we developed ways in which an optimal envy quote can be generated in a single optimization. Specifically, we developed an optimal method for any bidder-optimal core-selecting combinatorial auction, and another optimal method for the bidder-pessimal core-selecting (first-price) combinatorial auction.

References

L. Ausubel and P. Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1(1):1–42, 2002.

L. M. Ausubel and P. Milgrom. The lovely but lonely Vickrey auction. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*. MIT Press, 2006.

L. M. Ausubel, P. Cramton, and P. Milgrom. The Clock-Proxy Auction: A Practical Combinatorial Auction Design. In *Combinatorial Auctions*, chapter 5. MIT Press, 2006.

V. Conitzer and T. Sandholm. Self-interested automated mechanism design and implications for optimal combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 132–141, New York, NY, 2004.

P. Cramton. A Review of the 10-40 GHz Auction. Technical report, Office of Communications, United Kingdom, 2008a.

P. Cramton. A Review of the L-Band Auction. Technical report, Office of Communications, United Kingdom, 2008b.

P. Cramton. Ascending auctions. *European Economic Review*, 42 (3-5):745–756, 1998.

R. Day and P. Cramton. The Quadratic Core-Selecting Payment Rule for Combinatorial Auctions. Technical report, University of Maryland, 2008.

R. Day and P. Milgrom. Core-selecting package auctions. *International Journal of Game Theory*, 36(3):393–407, 2008.

R. Day and S. Raghavan. Fair Payments for Efficient Allocations in Public Sector Combinatorial Auctions. *Management Science*, 53(9):1389–1406, 2007.

A. Erdil and P. Klemperer. A new payment rule for core-selecting package auctions. Technical report, Oxford University, 2009.

J. K. Goeree and Y. Lien. On the impossibility of core-selecting auctions. Technical report, Institute for Empirical Research in Economics, 2009.

K. Hoffman, D. Menon, S. van der Heever, and T. Wilson. Observations and Near-Direct Implementations of the Ascending Proxy Auction. In *Combinatorial Auctions*, chapter 17. MIT Press, 2006.

V. Krishna and M. Perry. Efficient mechanism design. Technical report, Hebrew Univ. of Jerusalem, March 1997.

H. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *The Journal of Political Economy*, pages 461–479, 1983.

A. Likhodedov and T. Sandholm. Methods for boosting revenue in combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 232–237, San Jose, CA, 2004.

A. Likhodedov and T. Sandholm. Approximating revenue-maximizing combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Pittsburgh, PA, 2005.

P. Milgrom. Incentives in core-selecting auctions. Technical report, Stanford Department of Economics, Oct. 2006.

R. Myerson. Optimal auction design. *Mathematics of Operation Research*, 6:58–73, 1981.

D. Parkes. Notes on indirect and direct implementations of core outcomes in combinatorial auctions. Technical report, Harvard University, 2002.

D. Parkes. iBundle: An efficient ascending price bundle auction. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 148–157, Denver, CO, Nov. 1999.

A. Roth and J. Vande Vate. Random Paths to Stability in Two-Sided Matching. *Econometrica*, 58(6):1475–1480, 1990.

M. Rothkopf. Thirteen reasons why the Vickrey-Clarke-Groves process is not practical. *Operations Research*, 55:191–197, 2007.

M. Rothkopf, T. Teisberg, and E. Kahn. Why are Vickrey auctions rare? *Journal of Political Economy*, 98(1):94–109, 1990.

T. Sandholm. Issues in computational Vickrey auctions. *International Journal of Electronic Commerce*, 4(3):107–129, 2000. Early version in ICMAS-96.

T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135:1–54, Jan. 2002.

T. Sandholm and C. Boutilier. Preference elicitation in combinatorial auctions. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, pages 233–263. MIT Press, 2006. Chapter 10.

L. Shapley and M. Shubik. The assignment game I: The core. *International Journal of Game Theory*, 1(1):111–130, 1971.

P. R. Wurman, J. Zhong, and G. Cai. Computing price trajectories in combinatorial auctions with proxy bidding. *Electronic Commerce Research and Applications*, 3(4):329 – 340, 2004.

M. Yokoo. Pseudonymous Bidding in Combinatorial Auctions. In *Combinatorial Auctions*, chapter 7. MIT Press, 2006.