

## Operations Research

Publication details, including instructions for authors and subscription information: http:// pubsonline.informs.org

## Automated Design of Revenue-Maximizing Combinatorial Auctions

Tuomas Sandholm, Anton Likhodedov

## To cite this article:

Tuomas Sandholm, Anton Likhodedov (2015) Automated Design of Revenue-Maximizing Combinatorial Auctions. Operations Research 63(5):1000-1025. http:// dx. doi.org/ 10.1287/ opre.2015.1398

## Full terms and conditions of use: http://pubsonline.informs.org/page/terms-and-conditions

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2015, INFORMS
Please scroll down for article-it is on subsequent pages

## informs

INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.
For more information on INFORMS, its publications, membership, or meetings visit http:// www. informs.org

# Automated Design of Revenue-Maximizing Combinatorial Auctions 

Tuomas Sandholm<br>Computer Science Department, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, sandholm@cs.cmu.edu<br>Anton Likhodedov<br>Deutsche Bank, EC2N 2DB London, United Kingdom, anton.likhodedov@gmail.com


#### Abstract

Designing optimal-that is, revenue-maximizing-combinatorial auctions (CAs) is an important elusive problem. It is unsolved even for two bidders and two items for sale. Rather than pursuing the manual approach of attempting to characterize the optimal CA, we introduce a family of CAs and then seek a high-revenue auction within that family. The family is based on bidder weighting and allocation boosting; we coin such CAs virtual valuations combinatorial auctions (VVCAs). VVCAs are the Vickrey-Clarke-Groves (VCG) mechanism executed on virtual valuations that are affine transformations of the bidders' valuations. The auction family is parameterized by the coefficients in the transformations. The problem of designing a CA is thereby reduced to search in the parameter space of VVCA-or the more general space of affine maximizer auctions.

We first construct VVCAs with logarithmic approximation guarantees in canonical special settings: (1) limited supply with additive valuations and (2) unlimited supply.

In the main part of the paper, we develop algorithms that design high-revenue CAs for general valuations using samples from the prior distribution over bidders' valuations. (Priors turn out to be necessary for achieving high revenue.) We prove properties of the problem that guide our design of algorithms. We then introduce a series of algorithms that use economic insights to guide the search and thus reduce the computational complexity. Experiments show that our algorithms create mechanisms that yield significantly higher revenue than the VCG and scale dramatically better than prior automated mechanism design algorithms. The algorithms yielded deterministic mechanisms with the highest known revenues for the settings tested, including the canonical setting with two bidders, two items, and uniform additive valuations. ${ }^{1}$


Keywords: combinatorial auction; optimal auction; revenue maximization; automated mechanism design (AMD); parametric mechanism design.
Subject classifications: games/group decisions: bidding/auctions; marketing: pricing; computer science: artificial intelligence.
Area of review: Financial engineering
History: Received September 2008; revision received November 2009; accepted January 2011. (This paper was accepted in January 2011 but due to unforeseen delays by the authors it did not enter production until March 2015. The references have been updated to reflect papers that were accepted after January 2011 but before the publication date.)

## 1. Introduction

Combinatorial auctions (CAs), where bidders can bid on bundles of items, are nowadays popular mechanisms for allocating items (goods, tasks, resources, services, etc.). They are desirable in settings where the bidders' valuations exhibit complementarity and/or substitutability in the items. In such settings CAs have numerous advantages over traditional auctions. For example, they tend to yield better allocations and remove the bidders' exposure problems. The interested reader can find out more about CAs in a modern textbook on the topic (Cramton et al. 2006).

Perhaps the most important and elusive open problem in CAs (and the whole field of mechanism design) is designing optimal auctions, that is, auctions that maximize the seller's expected revenue (Vohra 2001). Astonishingly, this problem is unsolved even for auctions with just two distinct
items on sale (e.g., a TV and a VCR) and two bidders. A historical advance on the problem was the design of the optimal one-item auction (Myerson 1981). In that auction, instead of determining the winner and the payment based on bids, they are determined based on virtual valuations, which are transformations of the bids in a way that makes weak bidders (i.e., bidders who are likely to have low valuations) artificially more competitive. Myerson's auction was later extended to the case of selling multiple copies of the same item (Maskin and Riley 1989). Hartline and Karlin (2007) provide an overview of revenue-maximizing mechanism design. However, the characterization of revenuemaximizing multi-item auctions has been obtained only for special cases of the two-item two-bidder setting (bidders drawing valuations for the items from the same binary distribution) (Avery and Hendershott 2000, Armstrong 2000).

While it might be surprising that the revenue-maximizing CA is unknown, we observe that this is actually what one should expect once one views the problem through a computational lens. Conitzer and Sandholm (2004) proved that the problem of finding a revenue-maximizing CA (among all deterministic CAs with discrete types) is NP-complete. Therefore, it is unlikely that a concise characterization of revenue-maximizing (deterministic) CAs can even exist.

Consequently, we deviate from the classical manual mechanism design approach of looking for a characterization and instead proceed down a different avenue. We introduce a broad parameterized family of CAs-virtual valuations combinatorial auctions (VVCAs)—and develop algorithms that search for a high-revenue CA within that family. Even though it is well known that randomization can increase revenue in CAs, we focus on deterministic CAs, because in many applications, randomization is not palatable and very few, if any, randomized CAs are used in practice.

The bundling literature is also closely related to optimal CA design. That literature is concerned with bundling decisions by the seller, and the effect they have on revenue, in the context of catalog offers by the seller to the buyers. For example, Palfrey (1983) proved that in certain models, the seller obtains higher expected revenue by bundling the items together when the number of bidders is small, and he should auction the items separately when the number of bidders increases.

In designing algorithms for generating high-revenue CAs, we will use ideas from the optimal auction literature and the bundling literature. Specifically, we adapt the idea of artificially making weak bidders more competitive using virtual valuations (from Myerson's one-item auction, except that we use a different transformation to get from bids to virtual valuations) and ideas on tweaking the allocation rule (from bundling research). These techniques beget our parametric family of CAs, which we call VVCAs.

A classic CA design—and a benchmark we will useis the Vickrey-Clarke-Groves (VCG) mechanism (Vickrey 1961, Clarke 1971, Groves 1973) (aka Generalized Vickrey Auction), which in the one-item case is the second-price sealed-bid auction (aka Vickrey auction). The VCG allocates the items in a way that maximizes the social welfare (SW) of the bidders (sum of their valuations for the allocated items), and each winning bidder pays the minimal amount that she would have had to bid on the bundle she won in order to win it (not considering competition from any other bundles that she herself may have bid on). In this mechanism, each bidder's (weakly) dominant strategy is to bid their true valuations.

The rest of the paper is organized as follows. In the first part of $\S 2$, we review the CA setting, introduce notation, review the needed basics of mechanism design, discuss the VCG mechanism along with reasons for its revenue deficiency, and draw ideas from Myerson's single-item auction.

In $\S 2.5$, we discuss how revenue can be increased over the VCG by weighting bidders and boosting allocations. We introduce our VVCA family of auctions. A VVCA is a VCG run on virtual valuations that are affine transformations of the bidders' actual valuations. The coefficients of these transformations parameterize the family of VVCAs; we prove that incentive compatibility precludes more general transformations of the valuations (except potentially a greater weight for some bidders on the grand bundle). We also review the more general family of affine maximizer auctions (AMAs).

In §2.6, we design particular randomized VVCAs that yield a logarithmic worst-case approximation and deterministic VVCAs that yield a logarithmic average-case approximation to the optimal auction, for the canonical settings of (1) items in limited supply and additive valuations (no complementary or substitutable items) and (2) items in unlimited supply and general valuations. These results suggest that VVCAs are a reasonably powerful class of CAs for revenue maximization, and thus provide justification for our use of VVCAs as one of the mechanism classes we will use in automated mechanism design (AMD) later in the paper. These results may also be of independent interest, and generalize earlier results on prior-free mechanism design.

In the main part of the paper, $\S 3$, we pursue the approach of designing high-revenue auctions automatically-for general valuations. Our algorithms always design for the specific setting at hand-specifically, seller's prior over the bidders' valuations. (It turns out that one cannot obtain high revenue in a prior-free way.) Our approach is a form of AMD (Conitzer and Sandholm 2002). In prior AMD work, the types of bidders had to be discretized and an optimal mechanism was searched for using general-purpose integer programming algorithms. Thus that AMD work only scales to tiny CAs (Conitzer and Sandholm 2003). Our approach turns out to be significantly more scalable for at least two reasons. First, it does not assume that a complete prior is given as input. Rather, only samples from the prior are used, which enables sparse sampling. Second, economic insights and results are used to guide the search for highrevenue CAs. Although our approach may not construct an optimal auction, it yields significant revenue improvement over the VCG and provides high-revenue mechanisms for settings for which none were known.

We prove several properties of the problem to help guide our design of appropriate algorithms. We then present a sequence of increasingly sophisticated algorithms for searching for parameters within the VVCA family and the AMA family. Some of the algorithms use economic insights to navigate the search space efficiently in order to enhance computational speed. The experiments (§3.4) show that our algorithms create mechanisms that yield significantly higher revenue than the VCG, that our algorithms scale significantly better than the prior AMD algorithms, and that the more sophisticated ones of our algorithms tend
to outperform the more obvious ones in absolute run-time and anytime performance.

The generated mechanisms are, to our knowledge, the highest-revenue mechanisms known to date for their respective settings. For example, for the well-studied canonical setting of two bidders, two items, and uniformly drawn additive valuations, we generated the highest revenue mechanism to date. Experiments suggest that it is an optimal AMA within the precision used.

Section 4 overviews additional related research. Section 5 presents conclusions, and $\S 6$ lays out future research directions.

## 2. Framework and Contributions on the Way to the Main Results

We study a setting with one seller (index 0 refers to the seller), a set $N$ of $n$ bidders, and a set $G=\left(g_{1}, \ldots, g_{m}\right)$ of heterogeneous indivisible items for sale.

In an auction, the bidders submit bids for the bundles of items and the auction rules determine the allocation $a$ and the payments $t$, where $a_{i}$ is the bundle of goods that Bidder $i$ receives and $t_{i}$ is the payment by Bidder $i$.

### 2.1. Utilities and Valuations

We make the standard assumption that each Bidder $i$ has a quasi-linear utility function $u_{i}=v_{i}(a)-t_{i}$, where $v_{i}$ is the valuation of Bidder $i$ for allocation $a$. Each bidder's true valuations are private information. We also make the following standard assumptions (Lavi et al. 2003):

1. no externalities: the valuation of any Bidder $i$ for each allocation $a$ depends only on the bundle $a_{i}$ that Bidder $i$ receives, not on how the items that $i$ does not receive get allocated,
2. free disposal: the value of a subset of a bundle is less than or equal to the value of a bundle $\left(\forall b^{\prime} \subset b, v_{i}\left(b^{\prime}\right) \leqslant\right.$ $\left.v_{i}(b)\right)$, and
3. each bidder's valuation for the empty bundle is zero, i.e., $v_{i}(\varnothing)=0$.

In a CA, the valuation function for Bidder $i$ is given by the vector $\left(v_{i}\left(b_{1}\right), \ldots, v_{i}\left(b_{2^{m}}\right)\right)$, where each number specifies the value that Bidder $i$ attaches to a certain bundle of items from $G$ (there are $2^{m}$ bundles). Let $V_{i}$ denote the set of all possible valuation functions for Bidder $i . V$ denotes $\times_{i=1, \ldots, n} V_{i}$. Unless explicitly stated, we make the following standard assumptions throughout the paper:

1. $V_{i}$ is a convex compact subset of $\mathfrak{R}^{\left|2^{m^{m}}\right|}$.
2. Each valuation function $v_{i}$ is generated from a continuous density $f_{i}$, and $f_{i}$ is positive on all $V_{i}$.
3. The valuations of different bidders are drawn independently of each other.

In terms of the distributions from which valuations are drawn, two classes of models are typically considered in literature. In the symmetric case, $f_{i}=f_{j}$ for all Bidders $i$ and $j$. In the asymmetric case, valuations of different bidders are drawn from different $f_{i}$. We will consider both cases.

### 2.2. Mechanism Design Principles

Each bidder's valuation function is private informationalthough the auctioneer and other bidders may know the distribution from which it is drawn. Thus a concern is that a bidder might not reveal her true valuation function when bidding-she might be able to obtain higher utility by submitting a different valuation function. A key goal in mechanism design is to incentivize the bidders to tell the truth (by the revelation principle (for a review, see Krishna 2002) this is without loss of generality: any social choice function that can be implemented using an arbitrary mechanism can also be implemented using a truth-promoting mechanism). As is common in much of mechanism design, especially within computer science, we focus on ex post incentive compatible mechanisms, that is, mechanisms where each bidder maximizes her utility by bidding truthfully, regardless of what valuations the other bidders reveal. Such mechanisms are also called dominant strategy mechanisms. They are robust in the sense that the bidders cannot benefit from counterspeculating each others' valuations and rationality: each bidder has an optimal strategy regardless of other bidders' strategies. This also means that we do not need any assumptions about the bidders' knowledge of each others' valuation functions. In particular, we do not need the unrealistic assumption of common knowledge of priors, which underlies work on Bayes-Nash implementation.

As usual, we also require that the mechanism be ex post individually rational: each bidder is no worse off by participating than not participating, for all possible valuation revelations of the other bidders.

### 2.3. VCG Mechanism and Reasons Why It Does Not Maximize Revenue

A classic example that satisfies the above conditions is the following mechanism (Vickrey 1961, Clarke 1971, Groves 1973).

Definition 2.1 (VCG Mechanism). Each Bidder $i$ submits a valuation function $v_{i}$. The allocation, $a$, is computed to maximize SW
$\mathrm{SW}(v)=\sum_{i=1}^{n} v_{i}(a)$
The payment by Bidder $i$ is
$t_{i}=\left[\sum_{j \neq i} v_{j}\left(a_{-i}\right)-\sum_{j \neq i} v_{j}(a)\right]$,
where
$a_{-i}=\underset{a}{\arg \max } \sum_{j \neq i} v_{j}(a)$
is the allocation that is optimal among the allocations where Bidder $i$ does not receive any items. The SW of allocation $a_{-i}$ is
$\left[\mathrm{SW}_{-i}\right](v)=\sum_{j \neq i} v_{j}\left(a_{-i}\right)$.

One can also interpret $t_{i}$ as the minimum valuation for $a_{i}$ (the bundle won by $i$ ), which $i$ would have had to bid in order to win $a_{i}$.

The VCG maximizes the welfare of the bidders. However, it can yield arbitrarily poor revenue to the seller compared to a revenue-maximizing mechanism (Conitzer and Sandholm 2006). There are several different reasons why the VCG can yield poor revenue:

1. Bundling effect. The following well-known simple example shows that bundling decisions of the seller may affect the revenue (even when there is no complementarity or substitutability).

Example 2.1. Consider an auction with $k$ items for sale $\left(g_{1}, \ldots g_{k}\right)$, and two bidders. Say that the bidders' valuations for the individual items are drawn independently and uniformly from some interval, and a bidder's valuation for a bundle is the sum of her valuations for the individual items in the bundle. The VCG would sell each item separately to the higher bidder, collecting payment equal to the valuation of the lower bidder for each item. Therefore the revenue is $\sum_{j=1}^{m} \min _{i \in\{1,2\}} v_{i}\left(g_{j}\right)$. However, should the seller decide to bundle all the items together and sell them as a whole via a Vickrey (second-price) auction, she would receive revenue $\min _{i \in\{1,2\}}\left[\sum_{j=1}^{m} v_{i}\left(g_{j}\right)\right]$, which is greater.
2. Asymmetry of valuation distributions. In the asymmetric case, it may happen that the distribution of valuations of Bidder $i$ for some bundle $b$ is concentrated around higher values than (or even stochastically dominates) the distributions of other bidders (such bidders are called "strong" and "weak," respectively). Higher revenue can be obtained by favoring weak bidders, as does the Myerson auction discussed shortly. As an extreme example, consider a one-item auction where the valuation of Bidder 1 is surely higher than the valuations of the other bidders. The VCG would charge the second highest bid price, whereas it would be easy to improve revenue beyond that by charging (at least) the lowest possible valuation of the strong bidder.
3. No reserve prices. Even if the valuation distributions are symmetric (or even if there is only one bidder), expected revenue can be improved by setting reserve prices. They force the bidders to bid at least that much in order to win.

In this paper, we will design mechanisms that yield significantly greater revenue than the VCG mechanism in CAs. But before that, we will review one more classic mechanism, namely, Myerson's optimal one-item auction. We do this because we will adapt some ideas from it to CAs.

### 2.4. Ideas from Myerson's One-Item Auction

For increasing revenue in CAs, we will draw some ideas from the optimal one-item auction (Myerson 1981).

Definition 2.2 (Myerson's one-Item Auction). Each Bidder $i$ submits her valuation $v_{i}$ for the item. The mechanism computes virtual valuations for the bidders:
$\tilde{v}_{i}\left(v_{i}\right)=v_{i}-\frac{1-F_{i}\left(v_{i}\right)}{f_{i}\left(v_{i}\right)}$.
( $F_{i}$ is the cumulative distribution function corresponding to $f_{i}$.) The allocation is computed to maximize the following objective:
$\operatorname{SW}(v)=\sum_{i=1}^{n} \tilde{v}_{i}$.
Thus the item is given to the bidder with highest virtual valuation. ${ }^{2}$ Therefore the allocation rule is the same as in the VCG (Definition 2.1), except that virtual valuations are used in place of real valuations. The payment by the winning bidder is equal to the minimum bid that she would have had to make in order to win (that is, $\tilde{v}_{i}^{-1}\left(v_{j}\right)$, where $v_{j}$ is the second-highest bid). The item is sold only if the virtual valuation of the winning bidder is above 0 (in other words, the mechanism also, in effect, uses reserve prices; these can differ between bidders in the asymmetric setting). Losing bidders pay nothing.

The intuition behind the mechanism is that it is biased in favor of weak bidders, thus creating an artificial competition between weak and strong bidders, and extracting more revenue from strong bidders. It is easy to check that the transformation (3) brings down the valuations of strong bidders more than those of weak bidders. This mechanism, in effect, enables the auctioneer to set a high sell price for a strong bidder while motivating the bidder to stay truthful (even if the bidder knows that her valuation is greater than the valuations of all other bidders).

Drawing from this intuition, we propose increasing revenue in CAs by designing certain virtual valuations, and then running VCG on those valuations. We argue that virtual valuations are capable of improving the VCG with respect to bundling aspects, asymmetry handling, and reserve pricing. In the next section, we discuss the forms of virtual valuations that we use.

### 2.5. Techniques for Increasing Revenue in CAs

In this section, we introduce two families of CAs that employ virtual valuations, and a hybrid family that combines the two. We then analyze the restrictions that individual rationality and incentive compatibility impose on virtual valuations.

### 2.5.1. Bidder Weighting Technique.

Definition 2.3 (Bidder-Weighted Auction). Each Bidder $i$ submits a valuation function $v_{i}$. An allocation, $a$, is chosen that maximizes
$\mathrm{SW}^{\mu}(v)=\sum_{i=1}^{n} \mu_{i} v_{i}(a)$.

The parameters, $\mu$, of the mechanism are positive real numbers. The payments by the bidders are
$t_{i}=\frac{1}{\mu_{i}}\left[\sum_{j \neq i} \mu_{j} v_{j}\left(a_{-i}\right)-\sum_{j \neq i} \mu_{j} v_{j}(a)\right]$.
The mechanism effectively replaces the valuation function $v_{i}$ of Bidder $i$ with $\mu_{i} v_{i}$. This is useful in asymmetric cases when valuations of some bidders are concentrated around higher values than those of other bidders. The proof of incentive compatibility of this mechanism follows that of the VCG.

In many cases, the same bidder can be strong with respect to some bundle and weak with respect to another bundle. It would thus seem to be helpful to allow the mechanism to assign a bidder different weights for different bundles. However, it is easy to show that such a mechanism would not be incentive compatible.

Theorem 2.1. No mechanism that chooses an allocation, $a$, that maximizes $\mathrm{SW}^{\mu(a)}(v)=\sum_{i=1}^{n} \mu_{i}\left(a_{i}\right) v_{i}(a)$ is incentive compatible for all possible valuations in the CA domain unless $\mu_{i}\left(a_{i}\right)$ is constant over $a_{i}$. The only exception is that the grand bundle, $G$, can potentially have a larger multiplier, $\mu_{i}(G)$ (and the empty bundle can have any multiplier since $\mu_{i}(\varnothing)$ does not matter because $\left.v_{i}(\varnothing)=0\right)$.

For readability of the body of this paper, all proofs are presented in the appendix.
2.5.2. Allocation Boosting Technique. Theorem 2.1 shows that there does not exist a general bundle- and bidder-specific multiplicative weighting mechanism. However, it turns out that it is possible to give a bidder-bundlespecific advantage in an incentive compatible way by using additive terms. Let $\lambda_{i}(a)=c_{\{i, b\}}$ for all allocations $a$ that give Bidder $i$ exactly bundle $b$ (where $b$ can be any bundle of items, including the empty bundle). Here, the $c_{\{i, b\}}$ values (which we will sometimes write as $\lambda_{i}(b), \lambda_{i}\left(a_{i}\right)$, or $\left.\lambda_{\{i, b\}}\right)$ are real numbers that the auction designer sets. We call this the allocation boosting technique.

One way to think about this technique is to consider it to be an artificial bidder who has preferences over allocations but does not actually take any items herself. This technique addresses bundling and reserve pricing in full generality.
2.5.3. Bidder Weighting and Allocation Boosting. Now, we define a mechanism that uses allocation boosting and bidder weighting can be defined as follows.

Definition 2.4 (Virtual Valuations CA (VVCA)). Each Bidder $i$ submits a valuation function $v_{i}$. The mechanism computes an allocation $a$ that maximizes the weighted SW
$\mathrm{SW}_{\lambda}^{\mu}(a)=\sum_{i=1}^{n}\left[\mu_{i} v_{i}(a)+\lambda_{i}(a)\right]$.
Here, $\mu_{i}$ are positive and $\lambda_{i}(a)=c_{\{i, b\}}$ for all allocations that give Bidder $i$ exactly bundle $b$ (where $b$ can be
any bundle of items, including the empty bundle). The $\mu_{i}$ and $c_{\{i, b\}}$ are parameters chosen by the auction designer. The payment rule is

$$
\begin{align*}
t_{i}=\frac{1}{\mu_{i}}[ & \sum_{j \neq i}\left[\mu_{j} v_{j}\left(a_{-i}\right)+\lambda_{j}\left(a_{-i}\right)\right] \\
& \left.\quad-\sum_{j \neq i}\left[\mu_{j} v_{j}(a)+\lambda_{j}(a)\right]-\lambda_{i}(a)\right] \tag{7}
\end{align*}
$$

VVCAs are a family of mechanisms, parameterized by the vectors $\mu$ and $\lambda$. VCG is the special case where for all Bidders $i, \mu_{i}=1$ and $\lambda_{i}(\cdot)=0$. A VVCA can be thought of as the VCG mechanism run on bidders' virtual valuations (by analogy to Myerson's single-item auction) rather than their actual valuations: the mechanism replaces the valuation of Bidder $i, v_{i}(a)$, with the virtual valuation $\mu_{i} v_{i}(a)+\lambda_{i}(a)$. In other words, $\mu_{i} v_{i}(a)+\lambda_{i}(a)$ can be viewed as a virtual valuation, $\left[\tilde{v}_{i}\right]_{\lambda}^{\mu}$, of Bidder $i$ for allocation $a$. This technique allows one to apply qualitative aspects of the ideas of Myerson's revenue-maximizing single-item auction to CAs: the revenue can be increased by setting reserve prices and boosting the valuations of weak bidders (that is, bidders who are likely to have low valuations). Furthermore, the VVCA parameters enable the designer to favor any desired form of bundling (via favoring the allocations that exhibit such bundling), thereby covering also the third of the reasons listed above for why the VCG does not maximize revenue. All of these levers can increase competition in the auction and can increase the seller's expected revenue.

The VVCA mechanism adds $c_{\{i, b\}}$ to the value of the objective on allocations where Bidder $i$ gets bundle $b .^{3}$ Obviously, the probability of Bidder $i$ winning $b$ is increasing in $c_{\{i, b\}}$. The proof of incentive compatibility of VVCAs follows that of the VCG.

Impossibility of nonlinear virtual valuations. In the Myerson auction, a bidder's virtual valuation can be a nonlinear function of her valuation. However, we defined VVCAs in a way where a bidder's virtual valuation is an affine transformation of the bidder's valuation. Does this unnecessarily restrict the space of mechanisms that we should consider? The answer is no, as will now be shown.
Incentive compatibility imposes limitations on the virtual valuations that can be used in a mechanism. In one-item auctions, it is sufficient for the virtual valuations $\tilde{v}_{i}$ to be increasing in $v_{i}$ (Myerson 1981). However, this is not sufficient in CAs. Lavi et al. (2003) showed that under certain natural assumptions, every incentive compatible CA is almost ${ }^{4}$ an affine maximizer. Affine maximizers were introduced by Roberts (1979). He proved that they are the only ex post incentive compatible mechanisms over unrestricted domains of valuations. The valuations in CAs are not unrestricted because there are no externalities, there is
free disposal, and $v_{i}(\varnothing)=0$. Therefore the seminal results of Roberts do not apply here.

Definition 2.5 (Affine Maximizer Auction (AMA)). Each Bidder $i$ submits a valuation function $v_{i}$. The allocation, $a$, is computed to maximize ${ }^{5}$
$\mathrm{SW}_{\lambda}^{\mu}(a)=\sum_{i=1}^{n} \mu_{i} v_{i}(a)+\lambda(a)$.
Here, $\mu_{i}$ are positive numbers and $\lambda$ is an arbitrary function of allocation. The payments are
$t_{i}=\frac{1}{\mu_{i}}\left[\sum_{j \neq i} \mu_{j} v_{j}\left(a_{-i}\right)+\lambda\left(a_{-i}\right)-\sum_{j \neq i} \mu_{j} v_{j}(a)-\lambda(a)\right]$,
where
$a_{-i}=\underset{a}{\arg \max }\left[\sum_{j \neq i} \mu_{j} v_{j}\left(a_{-i}\right)+\lambda(a).\right]$
It is easy to see that VVCA mechanisms are a strict subset of AMAs: a VVCA is an AMA with the restriction
$\lambda(a)=\sum_{i} \lambda_{i}\left(a_{i}\right)$.
The results of Lavi et al. (2003) imply that every "reasonable" incentive compatible and individually rational general CA mechanism is an AMA. (Non-AMAs might be incentive compatible for some specific distributions of valuations, but only AMAs are incentive compatible for all CA settings.) Since AMAs transform allocation values affinely, among mechanisms that are based on virtual valuations (transformations of valuations rather than transformations of allocation values) only those that use affine virtual valuations ( $\tilde{v}_{i}$ affine in $v_{i}$ ) are incentive compatible for all CA settings. The VVCA family captures all such virtual valuations, and is thus the most general class of incentive compatible CA mechanisms that use virtual valuations.

Bidder-specific reserve prices. Despite the simple form of VVCA, controlling the parameters $(\mu, \lambda)$ is a powerful tool. For instance, it allows the auction designer to enforce or prevent any bidder from receiving a certain bundle. Another important property of the VVCA is that it allows for bidder-specific reserve prices. (Bidderspecific reserve prices are, in effect, also used in Myerson's revenue-optimal one-item auction. Recall that the item is sold only if the virtual valuation of the winning Bidder $i$ is above 0 , which sets a reserve price for this bidder to $\left.\tilde{v}_{i}^{-1}(0)\right)$. However, the reserve price mechanism does not generalize straightforwardly to arbitrary CAs. The standard way to set reserve prices in CAs is to submit fake bids by the seller. That approach does not support bidder-specific reserve prices. In contrast, the VVCA does support bidderspecific reserve valuations.

### 2.6. Logarithmic Approximations of the Optimal CA for Restricted Valuations

The problem of designing high-revenue CAs can be analyzed in two different frameworks:

1. Average case analysis is the standard approach in designing high-revenue auctions, in economics and computer science. In this setup, we assume that the valuations of the bidders are drawn from some underlying probability distributions (not necessarily the same for different bidders), and the auction designer knows the distributions (but not the exact draws, i.e., valuations, of the bidders). We do not assume that the bidders know each others' distributions. In this framework, the goal is to construct an auction, which yields high revenue on average with respect to the distributions.
2. Worst-case analysis of the problem has sometimes been used in computer science: in that framework, the objective is to construct an auction with worst-case performance guarantees (Goldberg et al. 2001, Guruswami et al. 2005). The advantage is that the design typically does not require complete knowledge of the underlying distributions, although the mechanisms are not completely prior free. A disadvantage is lower expected revenue. An essential feature of auctions with worst-case performance guarantees is randomization: in many cases, deterministic auctions perform far worse than randomized ones with respect to the worst-case performance objective (Goldberg et al. 2001). In this paper, we mainly take the more standard approach of average-case analysis. However, in this section, we present results for the average and worst cases.

Specifically, in this section, we study two important subclasses of CA setting: additive valuations and unlimited supply. For each of the two settings, we derive VVCAs that guarantee average-case and worst-case revenue that are provably within a bound of optimal.

These results suggest that VVCAs are a reasonably powerful class of CAs for revenue maximization, and as such these results serve to further motivate the study of VVCAs as one of the mechanism classes we will study for AMD later in the paper. These results may also be of independent interest, and generalize prior worst-case results.
2.6.1. Additive Valuations. In this section we study the special case where valuations are additive $(\forall i \in N$, $\left.\forall b \in 2^{G}, v_{i}(b)=\sum_{g \in b} v_{i}(\{g\})\right)$. In addition, we make the following mild natural assumptions about the priors:

1. Define $l$ to be the lowest possible valuation of any bidder for any individual item. We assume $l>0$, and that the auction designer knows $l$.
2. Define $h$ to be the highest possible valuation of any bidder for any individual item. We assume that the auction designer knows $h$.

We now construct a CA, which guarantees a fraction $1 /(2+2\lfloor\log (h / l)\rfloor)$ of the revenue of the optimal CA, even in the worst case. This generalizes the result of Guruswami et al. (2005), which was for bidders who only demand one item each.

Theorem 2.2. Consider $V V C A^{k}$, which is a VVCA, where the seller submits a bid of $l \cdot 2^{k}$ for every item. Formally, the $V V C A^{k}$ parameters are:

1. $\mu_{i}=1$ for all Bidders $i$ and
2. $\lambda_{i}(b)=-|b| \cdot l \cdot 2^{k}$ for all bundles $b$ (including singleitem bundles), where $|b|$ is the number of items in $b$, for all Bidders $i$.

Consider mechanism $M$ that selects $k$ uniformly at random from $\{0,1, \ldots,\lfloor\log (h / l)\rfloor\}$ and runs $V V C A^{k}$. $M$ is ex post incentive compatible, ex post individually rational, and for every given set of valuations $v$ yields expected revenue at least
$\frac{R_{\mathrm{OPT}}}{2+2\lfloor\log (h / l)\rfloor}$,
where $R_{\mathrm{OPT}}$ is the revenue of the optimal CA.
The same bound can also be made to hold in the averagecase framework with a deterministic CA.
Corollary 2.1. There exists $k$ such that VVCA ${ }^{k}$ yields a fraction $1 /(2+2\lfloor\log (h / l)\rfloor)$ of the revenue of the optimal auction on an expected revenue basis.

The right value of $k$ can easily be found by enumeration of all $V V C A^{k}$ and evaluating their expected revenues.

We obtained the logarithmic bounds in Theorem 2.2 and Corollary 2.1 by comparing the revenue of our auctions to the welfare of an efficient allocation, $\mathrm{SW}\left(a^{\mathrm{EFF}}\right)$, which obviously bounds the revenue of any individually rational auction. This proof technique cannot get us past the logarithmic approximation.

Theorem 2.3. For some additive-valuations auctions,
$\frac{E_{v}\left[R_{\mathrm{OPT}}\right]}{E_{v}\left[\mathrm{SW}\left(a^{\mathrm{EFF}}\right)\right]} \leqslant\left(1-\frac{1}{h / l}\right) \frac{1}{\ln (h / l)}$.
The above theorem is also of independent interest. It shows how much of the surplus the designer is unable to capture due to incomplete information about the bidders' valuations, even if he had a prior over them (as OPT does in the theorem).
2.6.2. Unlimited Supply. Another special case of the optimal CA design problem is the case when items are available in unlimited supply: the auctioneer is still selling items $g_{1}, \ldots, g_{m}$, but each item is now available in an infinite number of copies. One classic example of unlimited supply is the sale of digital goods: music files, video files, electronic books, etc. There are other examples as well, such as the nonexclusive licensing of patents.

In this setting, we assume that each bidder is interested in at most one copy of every item. This is not a restrictive assumption, since the preferences of a bidder who wants several copies of the same item can be expressed by adding these copies to the set of items $G$. As in §2.6.1, we assume that the lowest and highest possible valuation (for any bidder for any bundle), $l$ and $h$, are known by the auction designer. We do not assume that valuations are additive.

Since items are available in unlimited supply, there is no competition among the bidders: under the efficient allocation every bidder is allocated her most-wanted bid. Due to free disposal, the allocation that allocates the grand bundle $G$ (i.e., the bundle that includes a copy of each item) to each bidder is also efficient. This observation means that the design is all about reserve pricing, which, in turn, enables us to prove the following.
Theorem 2.4. Let $V V C A^{\prime k}$ be the VVCA with

1. $\mu_{i}=1$ for all Bidders $i$,
2. $\lambda_{i}(G)=-l \cdot 2^{k}$ for all Bidders $i$, and
3. $\lambda_{i}(b)=-\infty$ for all $b \subset G$ for all Bidders $i$.

Consider mechanism $M^{\prime}$ that uniformly randomly selects $k$ from $\{0, \ldots,\lfloor\log (h / l)\rfloor\}$ and runs $V V C A^{\prime k} . M^{\prime}$ is ex post incentive compatible, ex post individually rational, and for every given set of valuations, $v$ yields expected revenue at least
$\frac{R_{\mathrm{OPT}}}{2+2\lfloor\log (h / l)\rfloor}$,
where $R_{\mathrm{OPT}}$ is the revenue of the optimal auction.
Again, the same bound can be obtained with a deterministic mechanism in the average-case model:
Corollary 2.2. There exists $k$ such that $V V C A^{\prime k}$ yields fraction $1 /(2+2\lfloor\log (h / l)\rfloor)$ of the revenue of the optimal auction on an expected revenue basis.

The right value of $k$ can easily be found by enumerating all $V V C A^{\prime k}$ and evaluating their expected revenues.

## 3. Automated Design of High-Revenue CAs for General Valuations

In §2.6, we designed auctions with logarithmic competitive revenue ratios for settings, where the valuations had particular special structure. On the positive side, the randomized mechanisms in that section-unlike the Myerson auction-require no knowledge of the priors, $f_{i}$, on bidders' valuations except a lower bound $l$ and an upper bound $h$ on the support. On the negative side, logarithmic guarantees on revenue are quite weak from a practical revenue-maximization perspective, and superior mechanisms can be constructed if the designer has knowledge about the priors-using the techniques we will develop in this section.
In this section, we will study the design of high revenue deterministic CAs for general valuations. We first prove that any completely prior-free mechanism-which does not even know $l$ and $h$-can do arbitrarily poorly in terms of revenue. ${ }^{6}$

Theorem 3.1. For every completely prior free, incentive compatible, individually rational deterministic CA mechanism $M$, and every $\epsilon>0$, there exist distributions of valuation functions $V$ such that
$\frac{E_{v}\left(R_{M}(v)\right)}{E_{v}(\operatorname{OPT}(v))}<\epsilon$.

Here, $\mathrm{OPT}(v)$ denotes the revenue-optimal mechanism and $E_{v}$ denotes the expectation over $V$. This holds even for auctions where there is only one item for sale.

Theorem 3.1 shows that in order to construct a high revenue mechanism, we need to use some knowledge about the priors over bidders' valuations. Therefore we will focus on mechanisms that are designed using information about the prior.

In this section we suggest several automated approaches for constructing such mechanisms. We focus on averagecase analysis.

We do not assume that the priors are given explicitly. This is crucial in the CA setting because writing down the complete prior explicitly is infeasible: even in the discrete case the support of any bidder's valuation distribution is doubly exponential $\left(z^{2^{m}}\right.$ where $m$ is the number of items and $z$ is the number of possible values that the bidder might have for any given bundle). Our algorithms only use samples from the prior, which enables sparse sampling. This is in contrast to prior experiments on AMD for CAs by Conitzer and Sandholm $(2004,2003)$ that assumed an explicit representation of the prior and thus only scaled to a couple of items and bidders, and most restrictively, only a couple of possible types (valuation functions) per bidder. Another point of deviation from that earlier work is that we do not assume that the valuation space is discretized. Furthermore, we present custom algorithms for the problem while their setting, with the explicit prior, was amenable to solving-in the small-by standard mixed-integer programming packages.

### 3.1. Our Approach: Searching for Good Auction Parameters

The main idea in our paper is that we search computationally for good parameters of the auction within some class of auctions, where every auction is incentive compatible and individually rational. Thereby we simplify mechanism design down to the task of finding good parameters.

VVCA and AMA define families of mechanisms, parameterized by $(\lambda, \mu)$. Depending on the value of the parameters, the seller's expected revenue may be greater or less than in the VCG. The seller's revenue (given the valuations $v$ ) in VVCA is

$$
\begin{align*}
R(\mu, \lambda, v)= & \sum_{i=1}^{n} t_{i}(\mu, \lambda, v)=-\left(\sum_{i=1}^{n} \frac{1}{\mu_{i}}\right) \mathrm{SW}_{\lambda}^{\mu}(v) \\
& +\sum_{i=1}^{n} v_{i}(a)+\sum_{i=1}^{n} \frac{1}{\mu_{i}}\left[\mathrm{SW}_{-i}\right]_{\lambda}^{\mu}(v) \tag{10}
\end{align*}
$$

In this section, we will discuss the problem of finding parameters that yield high revenue in expectation.

To illustrate this idea, let us first consider general AMAs. The expected revenue is a function of the AMA parameters. Thus the problem of designing a high-revenue auction is
reduced to a search for the maximum of expected revenue in the AMA parameter space.

In the experiments, we evaluate parameter vectors by sampling valuations from the prior distributions (every sample is one complete valuation function for each bidder). The expected revenue of the AMA with a given set of parameters is estimated by running that AMA on each sample and averaging. Our approach can also be used in settings where the designer may not know-or may be unable to fully communicate-the actual distributions, but can provide samples.

### 3.2. Theory of Searching for Good Auction Parameters

The following theorem states that $E_{v}[R(\mu, \lambda, v)]$ is a "wellbehaved" function of $(\lambda, \mu)$.
Theorem 3.2. The expected revenue of the AMA (and consequently $V V C A$ ) is continuous and almost everywhere differentiable in $\mu$ and $\lambda$.

This suggests the use of numerical methods such as hill climbing for estimating locally optimal values of those parameters. That requires evaluating $E_{v}[R(\mu, \lambda, v)]$ for given $(\lambda, \mu)$, which can be estimated by sampling valuations from the distributions $f_{i}$.

Finding $R(\mu, \lambda, v)$ for a given set of valuations $v$ requires determining the AMA (6). Any optimal CA winner determination subroutine can be used here: the affine maximization problem can be converted into the standard CA winner determination problem by preprocessing the bids with the multiplicative and additive terms. We will discuss choices of winner determination subroutines in the context of specific auction design algorithms in §3.3.7

The main problem in optimization is that the number of parameters in $(\lambda, \mu)$ is exponential in the number of items for sale: $\mu$ is just a vector of size $n$, but the length of $\lambda$ is
$n 2^{m}$
in VVCA (for every bidder there is one parameter for every bundle, including the empty bundle) and
$(n+1)^{m}$
in AMA (one parameter for every allocation).
It would be helpful if we could discard some choices of $\lambda_{\{i, b\}}$ (by which we mean the variable $\lambda_{i}(b)$ corresponding to a specific $b$ in VVCA) from the search space in advance, thereby simplifying the optimization process. Unfortunately, the theorem below shows that there cannot exist a polynomial-time algorithm capable of always determining the optimal value for any $\lambda_{\{i, b\}}$, even if the valuations of the bidders are given. Moreover, no polynomial-time algorithm can always determine whether the mechanism with $\lambda_{\{i, b\}}$ set to some particular value $\lambda_{1}$ yields higher revenue than the mechanism with $\lambda_{\{i, b\}}$ set to $\lambda_{2}$.

Figure 1. Three-dimensional projection of the expected revenue surface in $(\mu, \lambda)$ space.


Note. The details of the setup behind this figure are explained later in the experimental section, §3.4.

Theorem 3.3. For any parameter $\lambda_{\{i, b\}}$ in VVCA and any pair of values of this parameter $\left(\lambda_{1}\right.$ and $\left.\lambda_{2}\right)$, there does not exist an algorithm that determines whether $R\left(\mu,\left(\lambda_{-\{i, b\}}, \lambda_{1}\right)\right)>R\left(\mu,\left(\lambda_{-\{i, b\}}, \lambda_{2}\right)\right)$ in polynomial time, even if the valuations $v$ of the bidders are given, unless $P=N P$. (Here, $\lambda_{-\{i, b\}}$ denotes the set of all $\lambda$ parameters except for $\lambda_{\{i, b\}}$.) The same is true for any parameter in AMA.

Theorem 3.3 shows that there is no easy general method to decide whether one set of parameters is better than another. Therefore, there is no easy way to fix some of the parameters up front without compromising optimality. Any search algorithm that guarantees the optimum for every distribution of valuations must run optimization over all parameters. ${ }^{8}$

A related problem that also makes optimization complicated is that the surface of $E_{v}[R(\mu, \lambda, v)]$ is nonconvex even in simple cases and can have ridges, see Figure 1. Therefore, local search algorithms can get stuck in local optima, which will be borne out in the experiments.

### 3.3. Viable Algorithms for Searching for Good Auction Parameters

The theory above curtails the space of viable algorithms for our task. In this section, we will introduce such algorithms.

If the number of items and bidders is small, we can run grid enumeration over all parameters and find an auction that is optimal (modulo grid discretization), as we will do in the small-scale experiments later in the paper. In contrast, for larger problems, we only search for a local optimum (because we are forced to search over the entire set of parameters by Theorem 3.3). We conduct this search using forms of hill climbing, suggested by Theorem 3.2.

We also tackle the complexity by introducing subfamilies of AMAs that have fewer parameters to optimize over, but
which still produce high revenue as our large-scale experiments will show.

We are now ready to present our first algorithm.
Algorithm 1 (BLAMA: Basic local optimization of AMA)

1. Sample the valuations from the prior distributions.
2. Start at some known AMA (typically the VCG or one of the AMAs with average-case performance guarantees from §2.6). Evaluate the mechanism at the sample points.
3. Run (Fletcher-Reeves conjugate) gradient ascent [Stoer and Bulirsch 1980] in the AMA parameter space from the starting point.

Algorithm 1 is still susceptible to the problem that we may have a prohibitive number of optimization parameters. For one, in order to compute the gradient for choosing the direction of the climb at every step, the algorithm must consider an exponential number of parameters.

To address this problem, we introduce new algorithms that guess the climbing direction based on insights drawn from the fact that we are in a CA domain. The idea of the first of these algorithms is from Equation (9), that is, the payment rule of AMA. If the payment, $t_{i}$, of Bidder $i$ in allocation $a$ is much lower than her valuation for $a$, one should expect that the her payment could have been increased. ${ }^{9}$ The payment can be increased directly only by (1) decreasing $\lambda(a)$, (2) increasing $\lambda\left(a_{-i}\right)$, or (3) modifying the $\mu$ parameters.

Algorithm 2 (ABAMA: Allocation boosting of AMA)

1. Sample the valuations from the prior distributions.
2. Start at some known AMA (typically the VCG or one of the auctions from §2.6).
3. For every sample point, compute the revenue loss on the winning allocation $a$ (we call this variant of the algorithm $A B A M A a$ ) or the second-best allocation (we call this variant of the algorithm $A B A M A b$ ). (The revenue loss from a bidder is the difference between the bidder's valuation and her payment. The revenue loss is the sum of the bidders' revenue losses.) Note that each allocation may be associated with multiple samples. Let the revenue loss of an allocation be the sum of the revenue losses of the samples associated with the allocation. Make a list of allocations in decreasing order of revenue loss.
4. Choose the first allocation, $a$, from the list. If the list is empty, exit.
5. Run (Fletcher-Reeves conjugate) gradient ascent in the $\{\mu, \lambda(a)\}$ subspace of the AMA parameter space; leave the other parameters unchanged. If the values of $\{\mu, \lambda(a)\}$ did not change (i.e., we cannot further improve the revenue by modifying $\{\mu, \lambda(a)\})$, remove $a$ from the list and go to Step 4. Otherwise, go to Step 3.

The only parameters considered by Algorithm 2 at each step are the $\mu$ and $\lambda$ corresponding to the winning or second-best allocations. In practice, the number of those allocations is small, which dramatically decreases the number of parameters in consideration.

A computational issue in algorithms that optimize over the entire AMA parameter space is that in the input to the winner determination, the parameter $\lambda$ can be different for every possible allocation, necessitating the explicit enumeration of all $(n+1)^{m}$ allocations in the winner determination. This further hinders the scalability.

To mitigate this problem, and to search in a smaller number of parameters than the number of parameters that AMAs have, we can focus on VVCAs instead (a VVCA has "only" $n 2^{m}$ parameters, one for every bidder-bundle pair). The parameters of VVCA are valuation (and not allocation) specific, so they can simply be preprocessed in before winner determination. Thus, any standard winner determination subroutine can be used: there is no need to explicitly enumerate all allocations, so each iteration of the auction design algorithm will run faster. For example, an integer programming package such as CPLEX can be used. Alternatively, one could use the dynamic program of Rothkopf et al. (1998), which runs in time $O\left(n^{1.59}\right)$ time if the input is represented in the flat way (one value for each bidderbundle pair) (Sandholm 2002).

Other than those differences, the design algorithm for VVCAs is similar to Algorithm 2 for AMAs.

Algorithm 3 (BBBVVCA: Bidder-bundle boosting of VVCA)

1. Sample the valuations from the prior distributions.
2. Start at some known VVCA (typically the VCG or one of the auctions from §2.6).
3. For every sample point, compute the payments of winning bidders. For every Bidder $i$ winning bundle $b$ and paying $t_{i}$, compute $v_{i}(b)-t_{i}$, i.e., the revenue loss for that bidder-bundle pair. Sum up the revenue losses over the sample and make a list of bidder-bundle pairs in decreasing order of the revenue loss.
4. Choose the first bidder-bundle pair, $\{i, b\}$, from the list. If the list is empty, exit.
5. Run (Fletcher-Reeves conjugate) gradient ascent in the $\left\{\mu, c_{\{i, b\}}\right\}$ subspace of the VVCA parameter space ( $\{i, b\}$ is the bidder-bundle pair, which incurs the highest revenue loss). Leave the values of all the other parameters unchanged. If the new values of $\left\{\mu, c_{\{i, b\}}\right\}$ do not change (i.e., we cannot improve the revenue further by modifying $\left.\left\{\mu, c_{\{i, b\}}\right\}\right)$, remove $\{i, b\}$ from the list and go to Step 4. Otherwise, go to Step 3.

### 3.4. Experiments

We conducted computational experiments with our CA design algorithms. In §3.4.1, we present experiments that compare the revenue of the techniques to each other, to grid search techniques, and to VCG. In §3.4.2, we present highaccuracy experiments to design the best mechanism for the canonical setting of two bidders, two items, and additive uniformly drawn valuations. In §3.4.3, we study the
scalability-in terms of computation time and revenueof the techniques to larger numbers of items and bidders. Finally, §3.4.4 studies the anytime performance of the design algorithms.
3.4.1. Revenue Compared to VCG and Iterated Grid Search-Based Parameter Optimization. We compared the four local search algorithms described in this paperBLAMA, allocation boosting AMA (ABAMA variants a and b), and BBBVVCA-against the VCG and against iterated grid search algorithms for optimizing the mechanism parameters.

The first iterated grid search algorithm, AMA*, optimizes AMA parameters as follows. In the first iteration, each dimension of the parameter space is discretized into $k$ values. The algorithm loops over all the grid points and evaluates the auction at each point. Then, the grid search is repeated but within a hyperrectangle that is centered around the optimal point from the previous iteration, has one $k$ th the size in each dimension, and still has $k$ discrete values in each dimension. This process is repeated a given number of times to narrow down on a good AMA parameter vector.

The second iterated grid search algorithm, $V V C A^{*}$, works analogously but optimizes VVCA parameters instead.

In each iteration of a grid search, a training set of valuation vectors is drawn (each vector includes a valuation function for each bidder), and each grid point is evaluated on that training set. Once all the iterations of the increasingly focused grid search have been completed, we evaluate the final auction parameter vector on a separate larger test set of valuation vectors.

One issue that these experiments uncovered is that the mechanism design algorithm—be it grid search or one of the local search methods-will eventually overfit the auction parameters to the training set. In fact, in the experiments in this section, already around the third iteration of grid search, the revenue of the auction stopped increasing (and often decreased somewhat) when evaluated on the test set. Increasing the size of the training set made this problem significantly less pronounced, as one would expect. In contrast, the time it takes to evaluate a grid point grows roughly linearly with the size of the training set. So, there is a trade-off.

We ran the local search methods to local optimum on the training set. A single training set throughout each run of each local search was used, but different training sets across different runs. Finally, we evaluated the mechanism parameter vectors on the test set.

The experiments in this section have two items, $g_{1}$ and $g_{2}$, and two bidders with valuation functions, $v_{1}$ and $v_{2}$, respectively. Valuations $v_{1}\left(g_{1}\right)$ and $v_{1}\left(g_{2}\right)$ are independently drawn from distribution $F_{1}$. Valuations $v_{2}\left(g_{1}\right)$ and $v_{2}\left(g_{2}\right)$ are independently drawn from distribution $F_{2}$. The valuation of Bidder 1 for the bundle of two items is given by $v_{1}\left(g_{12}\right)=$ $v_{1}\left(g_{1}\right)+v_{1}\left(g_{2}\right)+c_{1}$, where $c_{1}$ is a complementarity parameter drawn from distribution $C$. Analogously, $v_{2}\left(g_{12}\right)=$
$v_{2}\left(g_{1}\right)+v_{2}\left(g_{2}\right)+c_{2}$, where $c_{2}$ is also independently drawn from $C$. The results for various distributions $F_{1}, F_{2}$, and $C$ are given in Table 1. The columns correspond to the three different settings studied in this section. The first three rows specify the setting (distributions $F_{1}, F_{2}$, and $C$ ) and the last eight rows give the estimates of the expected revenue of the mechanisms found by the different algorithms.

All of the algorithms find mechanisms that generate significantly more revenue than the VCG-even in symmetric settings. All four local search techniques get stuck in local optima. Nevertheless, they close more than half of the revenue gap between VCG and the best AMA. The four techniques yield very similar revenue here, and have selective superiority across the settings.

Reducing the dimensionality of the search problem by fixing one $\mu$ parameter and one $\lambda$ parameter-something that can be done without loss of generality as was discussed in §3.2-did not affect revenue much for the local search methods or for the iterated grid search methods.

To provide visual insight into the problem, Figure 1 illustrates the AMA revenue surface as a function of the AMA parameters in Setting I. To visualize the revenue surface that is a function of two $\mu$ parameters and nine $\lambda$ parameters in just three dimensions, we fixed both $\mu$ 's and all $\lambda$ 's except for the following. The parameter $\lambda_{00}$ favors allocations where both items are kept by the seller, and $\lambda_{10}$ favors allocations where item 1 is allocated to Bidder 1 and the other item is kept. The analogous parameters $\lambda_{01}$ (referring to Bidder 1 and item 2), $\lambda_{20}$ (Bidder 2 and item 1), and $\lambda_{02}$ (Bidder 2 and item 2) are set equal to $\lambda_{10}$. The figure shows revenue on the training set.

We also studied the design of a bidder-symmetric AMA, that is, one where we force the additional equalities $\lambda_{01}=$ $\lambda_{02}, \lambda_{10}=\lambda_{20}, \lambda_{12}=\lambda_{21}$, and $\lambda_{11}=\lambda_{22}$. We call the algorithm that searches through the remaining space of $\mu$ and $\lambda$ parameters-again using iterated grid search-AMA ${ }_{b s y m}^{*}$. Here, fixing one of the remaining open lambda parameters
would not be without loss of generality because the symmetry constraints already reduce the degrees of freedom. Therefore, in the variant that tries to fix some parameters to reduce the dimensionality of the problem (italicized results in Table 1 and parameter vectors in Table 2), we only fix $\mu_{1}=1$.

Table 2 contains the AMA parameters discovered by the algorithms for each of the three settings. The parameters found by the different algorithms differ widely. This suggests that there are a wide variety of mechanisms that are almost optimal.

### 3.4.2. High-Accuracy Experiment to Design the Best

 Mechanism for the Canonical Setting (Setting I). There has been significant interest in revenue maximization in Setting I, a canonical setting. In this section, we will therefore study that setting in more detail.Since Setting I has additive valuations (complementarity parameter $C=0$ ), we can compare the revenue to running a separate Myerson auction for each item. ${ }^{10}$ That mechanism generates revenue $\frac{5}{6} \approx 0.833$ (Tang and Sandholm 2012). $A M A^{*}, V V C A^{*}$, and $A M A_{\text {bsym }}^{*}$ generate higher revenue than that (Table 1). This proves that our mechanisms outperform separate Myerson auctions even though there is no complementarity or substitutability. While this may seem surprising, it is what one should expect in light of known results in bundle pricing in catalog offers (Adams and Yellen 1976, McAfee et al. 1989, Bakos and Brynjolfsson 1999). Bakos and Brynjolfsson (1999) proceed to analyze this phenomenon in a setting with a large number of items for sale, and point out that "the law of large numbers makes it much easier to predict consumers' valuations for a bundle of goods than their valuations for the individual goods," and thus the seller can capture more of the surplus by bundling. In Setting I, bundling the two goods and running a Myerson auction on the bundle yields revenue 0.839 (Tang and Sandholm 2012). Our best mechanisms

Table 1. Revenue performance of the algorithms.

|  | Setting I |  |  |  | Setting II |  |  |  | Setting III |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $F_{1}$ | $U[0,1]$ |  |  |  | $U[1,2]$ |  |  |  | $U[1,2]$ |  |  |  |
| $F_{2}$ | $U[0,1]$ |  |  |  | $U[1,2]$ |  |  |  | $U[1,5]$ |  |  |  |
| C | 0 |  |  |  | $U[-1,1]$ |  |  |  | $U[-1,1]$ |  |  |  |
| VCG | $2 / 3 \approx 0.6667$ |  |  |  | 2.405 |  |  |  | 2.847 |  |  |  |
| AMA* | 0.860 | (0.847) | 0.868 | (0.864) | 2.76 | (2.73) | 2.77 | (2.76) | 4.24 | (4.17) | 4.16 | (4.14) |
| $A M A_{\text {bsym }}^{*}$ | 0.872 | (0.863) | 0.862 | (0.858) | 2.78 | (2.75) | 2.78 | (2.71) | 3.74 | (3.69) | 3.75 | (3.72) |
| $V V C A^{*}$ | 0.866 | (0.860) | 0.869 | (0.865) | 2.77 | (2.76) | 2.75 | (2.74) | 4.24 | (4.20) | 4.21 | (4.20) |
| BLAMA | 0.786 | (0.767) | 0.786 | (0.780) | 2.63 | (2.61) | 2.63 | (2.59) | 4.08 | (3.83) | 4.03 | (3.97) |
| $A B A M A ~_{a}$ | 0.786 | (0.784) | 0.786 | (0.784) | 2.63 | (2.62) | 2.63 | (2.62) | 4.01 | (3.83) | 4.00 | (3.79) |
| $A B A M A ~_{\text {b }}$ | 0.787 | (0.780) | 0.786 | (0.783) | 2.63 | (2.63) | 2.63 | (2.63) | 4.02 | (3.89) | 3.99 | (3.95) |
| BBBVVCA | 0.776 | (0.774) | 0.775 | (0.773) | 2.62 | (2.61) | 2.61 | (2.61) | 4.01 | (3.99) | 4.05 | (4.01) |

[^0]Table 2. AMA parameters computed by the various algorithms.

| Setting | Algorithm | $\mu_{1}$ | $\mu_{2}$ | $\lambda_{00}$ | $\lambda_{01}$ | $\lambda_{02}$ | $\lambda_{10}$ | $\lambda_{20}$ | $\lambda_{12}$ | $\lambda_{21}$ | $\lambda_{11}$ | $\lambda_{22}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| I | AMA* | 1 | 1.28 | 2 | 0 | 1.16 | 0 | 1.19 | 1.72 | 0.69 | 1.19 | 0.81 |
|  | $A M A_{\text {bsym }}^{*}$ | 1 | 1 | 1.15 | 0 | 0 | 0 | 0 | 0.05 | 0.05 | 0.33 | 0.33 |
|  | $V V C A * *$ | 1 | 1.09 | 1.28 | 0.66 | 0.66 | 0.72 | 0.63 | 0.09 | 0 | 0.39 | 0.30 |
|  | BLAMA | 1 | 1.00 | 0.16 | 0.16 | 0.16 | 0.16 | 0.16 | 0 | 0.01 | 0.31 | 0.32 |
|  | $A B A M A_{a}$ | 1 | 0.97 | 0.05 | 0.05 | 0.05 | 0.05 | 0.05 | 0 | 0.05 | 0.33 | 0.37 |
|  | $A B A M A_{b}{ }_{\text {a }}$ | 1 | 0.99 | 0.35 | 0.35 | 0.35 | 0.35 | 0.35 | 0 | 0.03 | 0.35 | 0.36 |
|  | BBBVVCA | 1 | 1.00 | 0.62 | 0.44 | 0 | 0.62 | 0.62 | 0 | 0.44 | 0.70 | 0.67 |
| II | AMA* | 1 | 1 | 1.25 | 1.25 | 1.25 | 1.25 | 1.25 | 0 | 0 | 0 | 0 |
|  | $A M A_{b s y m}^{*}$ | 1 | 1 | 2.95 | 1.53 | 1.53 | 1.485 | 1.48 | 0.27 | 0.27 | 0.44 | 0.44 |
|  | VVCA* ${ }^{\text {a }}$ | 1 | 1.43 | 3.50 | 2.25 | 1.88 | 2.25 | 1.88 | 0.63 | 0.63 | 1.52 | 0 |
|  | BLAMA | 1 | 0.99 | 0.28 | 0.28 | 0.28 | 0.28 | 0.28 | 0.02 | 0 | 0.56 | 0.57 |
|  | $A B A M A ~_{a}$ | 1 | 1.03 | 0.19 | 0.19 | 0.19 | 0.19 | 0.19 | 0.08 | 0 | 0.67 | 0.57 |
|  | $A B A M A_{b}$ | 1 | 1.00 | 0.60 | 0.60 | 0.60 | 0.60 | 0.60 | 0 | 0.02 | 0.60 | 0.60 |
|  | BBBVVCA | 1 | 0.97 | 0.92 | 0.92 | 0 | 0.92 | 0.56 | 0 | 0.56 | 0.92 | 1.06 |
| III | AMA* | 1 | 0.72 | 1.97 | 2.16 | 1.34 | 2.09 | 1.47 | 0 | 0.02 | 1.09 | 0.34 |
|  | $A M A_{\text {bsym }}^{*}$ | 1 | 1 | 5.14 | 1.70 | 1.70 | 1.70 | 1.70 | 0.60 | 0.60 | 0.03 | 0.03 |
|  | $V V C A^{*}$ | 1 | 0.88 | 3.88 | 2.63 | 1.38 | 2.75 | 1.25 | 0.25 | 0 | 1.88 | 0.06 |
|  | BLAMA | 1 | 0.77 | 0.30 | 0.30 | 0.30 | 0.30 | 0.30 | 0.06 | 0.05 | 1.08 | 0 |
|  |  | 1 | 0.62 | 1.24 | 1.24 | 1.24 | 1.24 | 1.24 |  | 0.81 | 1.24 | 0.97 |
|  | ABAMA $^{\text {b }}$ | 1 | 0.57 | 0.44 | 0.44 | 0.44 | 0.44 | 0.44 | 0 | 0.14 | 0.54 | 0.44 |
|  | BBBVVCA | 1 | 0.69 | 14.76 | 14.76 | 13.46 | 14.76 | 0 | 13.46 | 0 | 14.76 | 14.27 |

Notes. The parameters computed by the variant that fixed $\mu_{1}=1$ and the first lambda to zero are shown, that is, the solution corresponding to the italicized revenue numbers from Table 1. Many parameter vectors had negative lambdas; for readability, we scaled the lambdas additively within each parameter vector so the smallest lambda equals zero; this is without loss of generality as discussed earlier in the paper. (For $A M A_{b s y m}^{*}$ that would not be without loss of generality, so there we did not fix any lambda. To make the $A M A_{b s y m}^{*}$ parameter vector readable and comparable to the others, we simply constrained its lambdas to be nonnegative during the iterated grid search.)
yield higher revenue than that, so they do more than pure bundling with an optimal reserve price.

In a mixed bundling auction (Jehiel et al. 2007), there is a fake bidder who is only interested in the allocations where the entire set of items is allocated to any one bidder. The resulting auction works as if it gave a discount for a bidder who wins the grand bundle. The highest-revenue mixed bundling auction for Setting I yields revenue 0.786. Mixed bundling auctions can also be complemented with reserve prices, which are defined by further including the seller into consideration. This class is called a mixed-bundling auction with reserve prices (MBARP) (Tang and Sandholm 2012). In an MBARP, the seller submits a phantom bid of price $a$ on item 1 , a phantom bid of price $b$ on item 2 , and a phantom bid of price $a+b$ on the bundle; furthermore, there is
an extraneous preference of value $c$ to allocate both items together to either one of the bidders. MBARPs are a subset of $\lambda$-auctions, which are a subset of VVCAs, which are a subset of AMAs. Tang and Sandholm (2012) derived an analytical formula for MBARP revenue as a function of $a$, $b$, and $c$, and found the optimal MBARP for Setting I. It is shown in the first row in Table 3.

We conducted a high accuracy experiment with a training set of 40,000 and a test set of $40,000,000$ in order to find a high-revenue AMA for Setting I. We ran each of our local search algorithms starting from the optimal MBARP parameter vector. ABAMAb generated the highest-revenue parameters. We then manually enforced certain equalities (item and bidder symmetry) and did some rounding to make the AMA "nicer," as detailed in Table 3. As expected,

Table 3. AMA parameters and expected revenue in our high accuracy experiment for Setting I.

|  | $\mu_{1}$ | $\mu_{2}$ | $\lambda_{00}$ | $\lambda_{01}$ | $\lambda_{02}$ | $\lambda_{10}$ | $\lambda_{20}$ | $\lambda_{12}$ | $\lambda_{21}$ | $\lambda_{11}$ | $\lambda_{22}$ | Expected <br> revenue |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Optimal MBARP | 1 | 1 | 1.154 | 0.577 | 0.577 | 0.577 | 0.577 | 0 | 0 | 0.265 | 0.265 | 0.8705 |
| $\downarrow$ <br> ABAMA <br> $\downarrow$ <br> Enforcing symmetries <br> and rounding <br> $\downarrow$ <br>  | 1 | 0.98843 | 1.154 | 0.50890 | 0.51606 | 0.51280 | 0.52273 | 0.01729 | 0.01842 | 0.24118 | 0.25692 | 0.8744 |
| Iterated grid search | 1 | 1 | 1.15 | 0.5 | 0.5 | 0.5 | 0.5 | 0 | 0 | 0.25 | 0.25 | 0.8741 |

Notes. For displaying in this table, we round double-precision floating-point variables to five decimals. Given the test set size, there may be some amount of inaccuracy in the last reported digit of the expected revenue.
this caused the expected revenue to drop slightly. Finally, we conducted three rounds of iterated grid search, varying the three parameters $\lambda_{00}, \lambda_{01}=\lambda_{02}=\lambda_{10}=\lambda_{20}$, and $\lambda_{11}=\lambda_{22}$ in each iteration (we kept $\lambda_{12}=\lambda_{21}=0$ without loss of generality). This brought the revenue back up to what it was before we enforced the "niceness" (within the precision measurable on the test set).

We then ran each of our four local search-based algorithms on the final AMA parameters-without enforcing any symmetry or niceness constraints. Some of them did not change the parameter vector at all. Others improved the parameter vector slightly with respect to the training set, but decreased revenue on the test set. This suggests that we have reached a local optimum (within precision used in our algorithms), and the experiments from the previous section suggest that there is no better solution in the entire AMA search space. (Naturally, it may be possible to refine the AMA parameters further by searching additional digits of precision while using an even larger training and test set. We ran the final grid search on the aforementioned three lambda parameters down to changes of 0.005 .)

We also ran an iterated grid search to find the highestrevenue item-symmetric bidder-symmetric VVCA for Setting I using the high-accuracy training and test sets. In that experiment, $\mu_{1}=\mu_{2}=1, \lambda_{1}(\varnothing)=\lambda_{2}(\varnothing)=0$, and we varied the two parameters $\lambda_{1}(\{1\})=\lambda_{2}(\{1\})=\lambda_{1}(\{2\})=$ $\lambda_{2}(\{2\})$ and $\lambda_{1}(\{1,2\})=\lambda_{2}(\{1,2\})$. In each run, we used five consecutively finer grids, and in each grid, each dimension had nine grid points. Over 10 runs on different training sets, the highest revenue on the test set was 0.8703 (the associated AMA parameters are shown in Table 4, and are very close to those of the optimal MBARP in Table 3) and the average was 0.8702 . These are slightly below the best achieved with the optimal MBARP, but within numerical tolerance given the sizes of the training and test sets. At the same time, these are significantly lower than the AMA revenue in Table 3, 0.8744 . Interestingly, this suggest that in Setting I, generalizing the mechanism design from MBARPs to VVCAs does not yield additional revenue, but generalizing further to AMAs does.

It is well known that in some CA settings, randomized mechanism yield higher revenue than deterministic ones. The best achievable revenue in Setting I is unknown. It is bounded above by the revenue that the seller could extract if she knew the bidders' valuations exactly (that is, SW), which is $\frac{4}{3} \approx 1.33$. In this paper, we focus on deterministic AMAs, and the experiments suggest that we have succeeded in generating an optimal one within the numeric accuracy used.

Table 4. AMA parameters for the highest-revenue itemsymmetric bidder-symmetric VVCA found.

| $\mu_{1}$ | $\mu_{2}$ | $\lambda_{00}$ | $\lambda_{01}$ | $\lambda_{02}$ | $\lambda_{10}$ | $\lambda_{20}$ | $\lambda_{12}$ | $\lambda_{21}$ | $\lambda_{11}$ | $\lambda_{22}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 1 | 1.149 | 0.574 | 0.574 | 0.574 | 0.574 | 0 | 0 | 0.264 | 0.264 |

If our AMA is indeed optimal, that means that bidder and item symmetry $\left(\mu_{1}=\mu_{2}, \lambda_{01}=\lambda_{02}=\lambda_{10}=\lambda_{20}, \lambda_{11}=\right.$ $\lambda_{22}$, and $\lambda_{12}=\lambda_{21}$ ) can be simultaneously imposed on the AMA design here without compromising revenue.
3.4.3. Scalability Experiments. We also tested the scalability of the algorithms to larger problems. In the rest of the experiments, we will focus solely on the local searchbased algorithms because the grid search-based algorithms do not scale due to a combinatorial explosion. ${ }^{11}$

We generated the problem instances for the scalability experiments as follows. Each single item $g_{j}$ is assigned a value $v_{i}\left(g_{j}\right)$ drawn from $U[0,1]$. Then, each bundle $B$ containing two or more items is assigned the value
$v(B)=c_{B}+\sum_{j \in B} v_{i}\left(g_{j}\right)$,
where $c_{B}$ is drawn from $U[-|B| / m,|B| / m]$. Recall that $m$ is the total number of items. In addition, we check that the valuations satisfy monotonicity (free disposal). If monotonicity does not hold for a bundle, we add random noise to the valuation until monotonicity does hold. We do this by drawing a random number uniformly from [ $0, \frac{1}{2}$ ], and repeating as many times as necessary.

As described above, the valuations are symmetric in that the same procedure is used to generate the valuations $v_{i}$ for each Bidder $i$. We also consider a simple asymmetric version of the distribution in which the valuations of each Bidder $i$ are multiplied by $i$.

In each of the experiments below, the algorithms sampled 10 valuation functions for each bidder from the prior. The fact that the distributions are sampled, rather than modeled precisely, causes some strictly increasing curves (such as SW when the number of bidders increases) to appear nonmonotonic.

Scalability with respect to the number of bidders. Our first scalability experiment tests the algorithms' ability to scale as the number of bidders grows. We fixed the number of items to three and varied the number of bidders from 2 to 30 . Figure 2 Left shows how the algorithms scale on the symmetric instances and Figure 2 Right is for the asymmetric instances. The run-times are not affected by the symmetry of the distributions. In both settings, BBBVVCA is by far the fastest algorithm. This is partially due to the fact that it has the fewest parameters. Also, BBBVVCA uses CPLEX (or could use dynamic programming) as the winner determination subroutine while the other algorithms need to use allocation enumeration as explained above. At 30 bidders, BBBVVCA is about an order of magnitude faster than ABAMAa and ABAMAb, and about two orders of magnitude faster than BLAMA.

We are also interested in the algorithms' economic performance. Figures 3 and 4 show the expected revenue that the auctions designed by the algorithms generate for symmetric and asymmetric distributions, respectively. Each of the two figures contains four plots. The first row reports

Figure 2. (Color online) Scalability as the number of bidders increases.

absolute revenue and the second row reports revenue as a multiple VCG revenue. The second column contains a "zoomed-in" version of the plot to the left for readability. In these graphs, we also plot the expected revenue generated by the VCG, and the SW (which is an upper bound that is not actually achievable by any mechanism because

the designer has incomplete information about the bidders' valuations).

All of the algorithms improve revenue over the VCG. The relative improvement is greatest when the number of bidders is small and when the valuation distributions are asymmetric. Intuitively, as the number of bidders increases-

Figure 3. (Color online) Revenue of the mechanisms found by the algorithms on the symmetric distribution.


Note. Social welfare (SW) and VCG revenue are also shown. The graphs on the right are "zoomed-in" versions of the graphs on the left.

Figure 4. (Color online) Revenue of the mechanisms found by the algorithms on the asymmetric distribution.

especially in the symmetric setting-there is more competition, and the competition drives the VCG (or any other reasonable auction for that matter) to achieve high revenue. In fact, Monderer and Tennenholtz (2005) have proven that in the symmetric case, the VCG revenue is asymptotically optimal as the number of bidders goes to infinity. Our experiments also exhibit that phenomenon.

The fastest algorithm, BBBVVCA, does well in terms of revenue compared to the other algorithms when the number

of bidders is small, especially in the symmetric setting. However, its relative revenue performance becomes poor as the number of bidders increases. ABAMAa and ABAMAb yield comparable revenue to BLAMA and they are significantly faster than BLAMA. Perhaps surprisingly, BLAMA yields relatively poor revenue in the asymmetric case when the number of bidders is very small.

Scalability with respect to the number of items. Figure 5 shows how the algorithms scale with the number of items.

Figure 5. (Color online) Scalability as the number of items increases.


Figure 6. (Color online) Revenue of the mechanisms found by the algorithms on the symmetric distribution.


We fixed the number of bidders to three and varied the number of items from 2 to 10 . Again, we see that BBBVVCA is by far the fastest. It may seem that the run-time increases rapidly, but recall that the input size increases exponentially with the number of items.

Figures 6 and 7 show the expected revenue that the auctions designed by the algorithms generate for symmetric and asymmetric distributions, respectively. The fastest algorithm, BBBVVCA, does well in terms of revenue only in the asymmetric case when the number of items is less than four; as the number of items increases, BBBVVCA yields hardly any additional revenue over the VCG. ABAMAa and ABAMAb yield comparable-or higher-revenue than BLAMA and they are significantly faster than BLAMA.
3.4.4. Anytime Performance. The results above suggest that some of the algorithms are faster and yield lower revenue while others are slower and yield higher revenue. One can potentially achieve better trade-offs between runtime and revenue by not running an algorithm to completion, but by using an interim solution. All of the algorithms that we designed have the anytime property: they have a solution available at any time, and the more time is allocated to the algorithm, the better the solution it finds (until it terminates).


Figure 8 demonstrates the anytime performance of the algorithms. Two test instances were used. Both have seven items and seven bidders. The first has valuations drawn from the symmetric distribution (Figure 8 Left). The second has valuations drawn from the asymmetric distribution (Figure 8 Right). For each algorithm, the expected revenue of the best mechanism it has found so far is plotted against time. The expected revenue generated by the VCG mechanism is represented as a flat line. In the symmetric setting, all of the algorithms have selective superiority: which one should be selected depends on how much time is available. In the asymmetric setting, BLAMA is dominated but the other algorithms have selective superiority.

## 4. Additional Related Research

We discussed the most closely related research in the body of the paper. Here, we will discuss some additional interesting related work.

Levin (1997) also studies the problem of designing revenue-maximizing CAs, but he limits his analysis to the case where the buyers' preferences are strictly complementary. Moreover, he makes the highly simplifying assumption that the buyers' valuations are parameterized by a single type, in addition to satisfying a number of technical conditions on the valuation distributions. In contrast,

Figure 7. (Color online) Revenue of the mechanisms found by the algorithms on the asymmetric distribution.


Figure 8. (Color online) The algorithms' anytime performance.

the approaches presented in our paper use much milder assumptions on the valuations.

Bartal et al. (2003) develop incentive compatible auctions that are not affine maximizers, focusing on multi-unit CAs, where bidders desire a small number of units of each item. They work in the worst-case framework and their objective is efficiency.

In parallel with our work and independently, Hitt and Chen (2005) studied mechanism design in a combinatorial

Anytime performance ( 7 bidders, 7 items, asymmetric dist.)

market. Their setting differs from ours in a number of ways. First, they assume that the seller knows the valuations of all of the types of buyers (as well as the market proportion of each type), and is interested in finding prices, one for each type of buyer, to optimize revenue. A second difference is that their paper restricts attention to mechanisms where the buyers are only allowed to choose a fixed number of the items for a single fixed price.

There has also been interesting related work subsequent to the conference versions of our work (Likhodedov and Sandholm 2004, 2005), but prior to the submission of our journal version in 2008. We will discuss that research in what follows.

Ledyard (2007) finds the revenue-maximizing CA in the restricted setting of known single-minded bidders, that is, setting in which each bidder only has value for one bundle and that bundle is known to the seller.

Balcan et al. (2008b) study item pricing (bundle pricing not allowed) in a setting where buyers may have combinatorial valuations. In the unlimited supply setting, they get a revenue approximation of $O(\log n+\log m)$, where $n$ is the number of items and $m$ is the number of bidders. They do not place any assumptions on the valuations, including not even requiring monotonicity (that is, free disposal). In our paper, we get a $(2+2\lfloor\log (h / l)\rfloor)$ approximation, where $h$ and $l$ are the highest and the lowest possible valuation, respectively. This result requires monotonicity, and our mechanism is more general than item pricing.

Virtual valuations have recently been used in mechanism design in novel ways. Hartline and Roughgarden (2008) use them in settings where monetary transfers are not possible. Chawla et al. (2007) study multi-item pricing in a setting where the seller has priors, buyers have unit demand, and the buyers' valuations for the different items are independent random variables. They present a constant approximation algorithm for this problem that makes use of a connection between this problem and virtual valuations.

Balcan et al. (2008a) use techniques from sample complexity in machine learning to reduce problems of incentive compatible mechanism design to standard algorithmic questions, for a class of revenue-maximization problems. For those problems, their technique enables one to convert an optimal (or $\beta$-approximation) algorithm for an algorithmic pricing problem into a $(1+\epsilon)$-approximation (or $\beta(1+\epsilon)$-approximation) for the incentive compatible mechanism design problem, as long as the number of bidders is at least $O\left(\beta / \epsilon^{2}\right)$ times a measure of the complexity of the class of allowable pricings. They apply the results to the problem of auctioning a digital good, to the attribute auction problem, which includes a variety of discriminatory pricing problems, and to the problem of item pricing in unlimited-supply CAs. The work is in the prior-free setting. The idea of using samples of some bidders to price on other bidders was, to our knowledge, first applied to the multi-item prior-free setting by Goldberg et al. (2001). That work is different from ours in many ways. For one, they study unlimited supply. An overview of those and other revenue-maximization techniques is given by Hartline and Karlin (2007).

Recently, Wu et al. (2008) studied nonlinear mixedinteger programming to find bundle prices. That work falls within the AMD framework and is thus related to ours, but the setting is posted pricing while ours is an auction. Their setting is also more restricted than ours. For example, the
consumers are only allowed to choose up to $N$ goods out of a larger pool of $J$ goods.

## 5. Conclusions

The design of optimal-that is, revenue-maximizing-CAs (CAs) is an important recognized elusive research problem. The characterization is open even for two items and two bidders. Recent results show that the problem of finding an optimal CA (among all deterministic CAs in the setting with discrete types of agents) is NP-complete. That casts doubts on whether the manual approach of characterizing the optimal CA can succeed because a concise characterization is unlikely to exist.

Therefore, in this paper we developed a new approach to the problem creating high-revenue CAs. Instead of attempting a full characterization, we developed methods for modifying rich parametric auction mechanisms to obtain high-revenue CAs.

We introduced a general family of auctions, based on the techniques of bidder weighting and allocation boosting, which we call virtual valuations combinatorial auctions (VVCAs). All auctions in the family are based on the VCG mechanism, executed on virtual valuations that are affine transformations of the bidders' valuations. The VVCA family is parameterized by the multipliers and constants in the transformations. The VVCA family is a subset of a more general parametric family called affine maximizer auctions (AMAs). Each auction in the family is ex post individually rational and dominant strategy incentive compatible. (Recent theory suggests that AMAs are almost the broadest family with this required property.) Therefore the problem of designing high-revenue CAs reduces to search in the parameter space of AMA or VVCA. AMA has many more parameters; our VVCA family is the restriction where the auction has to be based on virtual valuations-an idea motivated by Myerson's optimal single-item auction. We proved that VVCAs are the most general family of virtual valuations-based CAs (except that the grand bundle could have higher special multipliers for some bidders).

We first studied the design of VVCAs in a setting where the designer does not know the bidders' valuation distributions (aka priors), but only knows an upper and lower bound on valuations. We designed randomized VVCAs that yield a logarithmic worst-case approximation and deterministic VVCAs that yield a logarithmic average case bound from optimal revenue, for the canonical settings of (1) items in limited supply and additive valuations (this generalizes earlier results by others that were for singleunit demand) and (2) unlimited supply. We also proved that in the former setting, there is a logarithmic gap between the revenue of the optimal unrestricted mechanism and SW: the seller fails to capture that much of the surplus due to incomplete information about the bidders' valuations (even if he knows the priors).

In the main part of the paper, we developed algorithms for automatically designing high-revenue CAs for unrestricted
valuations. The algorithms work by searching in the parameter space of VVCA or AMA. They use knowledge about the priors. This is in fact necessary in order to achieve high revenue: we proved that any (deterministic) mechanism that is completely prior free can yield an arbitrarily small fraction of the revenue that an optimal mechanism yields. More recently, others have shown that with a known upper and lower bound on valuations, a logarithmic approximation can be achieved but no better. A logarithmic bound is not meaningful from the perspective of practical revenue maximization, so we wanted to achieve revenues drastically higher than that. Therefore we use priors. Our algorithms do not require explicit representation of the priors as input. Rather, only samples from the priors are used, which enables sparse sampling. This is a significant deviation point from prior research on AMD and is a key to scalability, because in CAs, each prior is doubly exponential even in the discrete case.

Our design of the algorithms was guided by properties of the problem that we proved. First, the revenue of VVCA and AMA is a continuous and almost everywhere differentiable function, suggesting the use of hill-climbing methods for parameter search. On the negative side, we proved that the revenue surface is nonconvex and showed that it has ridges. We also proved that it is NP-hard to determine whether one setting of any of the additive parameters is better than another. It follows that one has to search over all those parameters if one is interested in finding a provably optimal solution.

Despite those negative results, near-optimal parameters for VVCA and AMA can be found in settings with few items and bidders. Experiments on small auctions showed that they yield a drastic increase in revenue over the VCG.

The generated mechanisms are, to our knowledge, the highest-revenue mechanisms known to date for their respective settings. For example, for the canonical setting of two bidders, two items, and uniformly drawn additive valuations, we generated the highest-revenue mechanism to date. Experiments suggest that it is an optimal AMA within the precision used. Furthermore, the mechanism achieves that while honoring bidder and item symmetry.

With larger numbers of bidders and items, locally optimal (or merely suboptimal) parameters can be used within VVCA and AMA. This can be done in a way that guarantees that revenue is no worse than that of the VCG, by starting the hill climbing in parameter space from the VCG. The best of our hill-climbing algorithms use economic insights to navigate the search space efficiently to enhance computational speed and revenue lift. Experiments showed that they yield significantly higher revenue than the VCG, that they scale much better than prior AMD algorithms (which assumed discretized type spaces and an explicitly represented prior), and that the more sophisticated methods indeed tend to outperform the more obvious ones in absolute run-time and anytime performance.

## 6. Future Research

A wealth of interesting future research questions remain. For example, how much-and exactly how-should one sample from the actual valuation distribution to construct the sample distribution used by our AMD algorithms? The run-time of the algorithms increases roughly linearly with the number of samples. In contrast, if too small a sample is used, overfitting occurs. The algorithms could use, for example, cross-validation (see, e.g., Mitchell 1997) to detect and deal with such overfitting.

Another direction involves developing sparse representations of priors that are nevertheless expressive enough to capture the essence of the particular auction setting to which AMD is being applied. Sparse representations are important for scalability because, as we discussed, the flat representation of the prior is doubly exponential even in the discrete case. Furthermore, it would be desirable to also have sparse representations for the samples (bidders' valuations) that the algorithms draw from the prior. Our algorithms use winner determination as a subroutine, and modern scalable winner determination algorithms based on tree search/integer programming take advantage of such sparseness. ${ }^{12}$ They can optimally solve very large instances in practice (for a review of optimal winner determination algorithms, see Sandholm 2006, and for a perspective on state-of-the-art scalability, see Sandholm 2007, 2013).

There is also interesting work left on the organization of the search of the parameter space. Although we proved that any algorithm that guarantees optimality must run optimization in an exponential number of parameters in general, for special applications or special priors, there might be more efficient algorithms for optimizing the parameters-or there might even be good simpler mechanism families to optimize over. It also remains interesting to try to design better algorithms for the general case.

It would be interesting to take this paradigm to realworld applications: estimating priors, designing mechanisms, and comparing the different design algorithms on real data.

From a theoretical and practical perspective, it would be interesting to see whether one can draw additional insights or learn some principles for enhancing revenue by studying the mechanisms that have been automatically designed. This may require tools for visualizing those mechanisms in order to better understand them.

It would also be desirable to extend the AMD approach to other settings. For example, it would be desirable to automatically create auction mechanisms on the fly for selling tail keywords (that have little or no competition) in sponsored search. There is also recent work on optimizing bundling and bundle discounts in catalog sales (Benisch and Sandholm 2011) and work on bundling items when the bundles will be sold in a VCG (Kroer and Sandholm 2015). As a different type of example, in combinatorial exchanges with multiple buyers and multiple sellers-such as spectrum exchanges and airport landing slot exchanges-the
two-sided asymmetric information about valuations hurts efficiency (Myerson and Satterthwaite 1983), and we believe that automated design algorithms that use priors to minimize the inefficiency would be desirable.

## Acknowledgments

The authors thank Andrew Gilpin and Pingzhong Tang for their significant help in programming and conducting some of the experiments, and for commenting on drafts of the journal version of this paper. This material is based on work supported by the National Science Foundation [Grants IRI-9703122, IIS-9800994, IIS-0081246, IIS-0121678, IIS-0427858, IIS-1320620, and IIS1546752], as well as a Sloan Fellowship.

## Appendix A. Proofs

## Proof of Theorem 2.1.

Step I. Suppose there exists a Bidder $i$ and nonempty bundles $a_{i}^{1}$ and $a_{i}^{2}$ such that $\mu_{i}\left(a_{i}^{1}\right) \neq \mu_{i}\left(a_{i}^{2}\right)$. Without loss of generality assume $\mu_{i}\left(a_{i}^{1}\right)<\mu_{i}\left(a_{i}^{2}\right)$. Here, in Step I, we analyze the Case $a_{i}^{1} \not \subset a_{i}^{2}$. We will construct a valuation space where the allocation rule that maximizes $\mathrm{SW}^{\mu(a)}(v)$ violates weak monotonicity (WMON) (which is a necessary condition for incentive compatibility, proved by Lavi et al. 2003).
Definition A. 1 (W-MON). Suppose the allocation chosen under valuations $\left(v_{i}, v_{-i}\right)$ of Bidder $i$ and other bidders is $a$, and the allocation chosen under valuations $\left(u_{i}, v_{-i}\right)$ is $\tilde{a}$. A mechanism satisfies W-MON if
$u_{i}(\tilde{a})-v_{i}(\tilde{a}) \geqslant u_{i}(a)-v_{i}(a)$
holds for all $i, u_{i}, v_{i}$, and $v_{-i}$.
Denote $\mu_{i}\left(a_{i}^{1}\right)$ by $\mu_{1}$ and $\mu_{i}\left(a_{i}^{2}\right)$ by $\mu_{2}$. Set all the valuations of all bidders except Bidder $i$ to some small $\epsilon>0$. For Bidder $i$, set the valuations for $a_{i}^{1}$ and all supersets of $a_{i}^{1}$ to
$v_{i}\left(a_{i}^{1}\right)=\frac{2}{\mu_{1}}$.
Also, set the valuations for $a_{i}^{2}$ and all supersets of $a_{i}^{2}$, that are not supersets of $a_{i}^{1}$, to
$v_{i}\left(a_{i}^{2}\right)=\frac{1}{\mu_{2}}$.
These valuations honor free disposal, no externalities, and $v_{i}(\varnothing)=0$.

Take
$\epsilon<\frac{1}{n \cdot \max _{\{i, a\}} \mu_{i}\left(a_{i}\right)}$.
This choice of valuations ensures that Bidder $i$ gets bundle $a_{i}^{1}$. Also, the difference in the value of
$\mathrm{SW}^{\mu(a)}(v)=\sum_{i=1}^{n} \mu_{i}\left(a_{i}\right) v_{i}(a)$
for $a$ (the best allocation where $i$ gets $a_{i}^{1}$ ) and $\tilde{a}$ (the best allocation where $i$ gets $a_{i}^{2}$ ) is less than 2 .

Suppose now the valuation of Bidder $i$ for bundle $a_{i}^{2}$ increases by
$\Delta_{2}=\frac{2+\mu_{1} \epsilon}{\mu_{2}-\mu_{1}}$
and the valuation for bundle $a_{i}^{1}$ increases by
$\Delta_{1}=\Delta_{2}+\epsilon$.
The value of weighted SW on $a$ increases by
$\Delta\left(\mathrm{SW}_{1}^{\mu(a)}\right)=\mu_{1}\left(\frac{2+\mu_{1} \epsilon}{\mu_{2}-\mu_{1}}+\epsilon\right)$,
while the value of the same objective on $\tilde{a}$ increases by

$$
\begin{aligned}
\Delta\left(\mathrm{SW}_{2}^{\mu(a)}\right) & =\mu_{2}\left(\frac{2+\mu_{1} \epsilon}{\mu_{2}-\mu_{1}}\right)=\left[\mu_{1}+\left(\mu_{2}-\mu_{1}\right)\right]\left(\frac{2+\mu_{1} \epsilon}{\mu_{2}-\mu_{1}}\right) \\
& =\mu_{1}\left(\frac{2+\mu_{1} \epsilon}{\mu_{2}-\mu_{1}}\right)+2+\mu_{1} \epsilon=\Delta\left(\mathrm{SW}_{1}^{\mu(a)}\right)+2
\end{aligned}
$$

Therefore, under the new valuations, the allocation rule will choose $\tilde{a}$ despite the fact that the valuation of Bidder $i$ for $a$ increased more than her valuation for $\tilde{a}$. That contradicts W-MON. Therefore the mechanism is not incentive compatible under the assumption used in Step I, i.e., $a_{i}^{1} \not \subset a_{i}^{2}$. We have thus shown that the mechanism is not incentive compatible if (a) neither $a_{i}^{1}$ nor $a_{i}^{2}$ is a subset of the other or (b) $a_{i}^{2} \subset a_{i}^{1}$.

Step II. By (a), we have that $\mu_{i}(\{j\})=\mu_{i}(\{k\})$ for all items $j$ and $k$, and that $\mu_{i}(S)=\mu_{i}(\{j\})$ for all nonempty bundles $S \subseteq$ $G-\{j\}$ and all $j$. It follows that $\mu_{i}(S)$ has to be constant for all bundles $S$, except perhaps the grand bundle $G$ (and the empty bundle whose $\mu_{i}(\varnothing)$ does not matter because $\left.v_{i}(\varnothing)=0\right)$. From (b), it follows that $\mu_{i}(G) \geqslant \mu_{i}(S)$.
Proof of Theorem 2.2. Before giving the proof, we need to introduce the following notation. Let $a^{\mathrm{EFF}}$ be an efficient allocation, $a^{k}$ be the winning allocation of $V V C A^{k}$, and $a_{-i}^{k}$ be the allocation that would have won had Bidder $i$ not submitted any bids. Let $v_{N}\left(g_{j}\right)$ be the highest bid for item $g_{j}: v_{N}\left(g_{j}\right)=\max _{i^{\prime} \in N} v_{i^{\prime}}\left(g_{j}\right)$. Also, let $v_{N \cup\{0\}}^{k}\left(g_{j}\right)$ be the highest bid for item $g_{j}$, including the bid of the seller
$v_{N \cup\{0\}}^{k}\left(g_{j}\right)=\max \left\{v_{N}\left(g_{j}\right), l \cdot 2^{k}\right\}$.
Finally, let $v_{N \cup\{0\} \backslash\{i\}}^{k}\left(g_{j}\right)$ be the highest bid for item $g_{j}$, including the bid of the seller, but excluding the bid of Bidder $i$
$v_{N \cup\{0 \backslash \backslash i\}}^{k}\left(g_{j}\right)=\max _{i^{\prime} \in\{1, \ldots, n\} \backslash\{i\}}\left\{v_{i^{\prime}}\left(g_{j}\right), l \cdot 2^{k}\right\}$.
Because the valuations are additive, $a^{\mathrm{EFF}}$ allocates every item $g_{j}$ according to $v_{N}\left(g_{j}\right)$, that is, to Bidder $i^{\prime}\left(1 \leqslant i^{\prime} \leqslant n\right)$ that submitted $v_{N}\left(g_{j}\right)$. Since the seller's bids are also additive, $a^{k}$ allocates every item $g_{j}$ according to $v_{N \cup\{0\}}^{k}\left(g_{j}\right)$ and $a_{-i}^{k}$ allocates every item $g_{j}$ according to $v_{N \cup\{0 \backslash \backslash\{i,}^{k}\left(g_{j}\right)$.

We will use the following lemma in the proof.
Lemma A.1. Consider a set of bidders' valuations v. If Bidder $i$ wins bundle $b$ in $V V C A^{k}$, she pays at least $|b| \cdot l \cdot 2^{k}$.

Proof of Lemma. By Equation (7), the payment of Bidder $i$ is

$$
\begin{aligned}
t_{i}= & \mathrm{SW}_{\lambda}^{\mu}\left(a_{-i}^{k}\right)-\mathrm{SW}_{\lambda}^{\mu}\left(a^{k}\right)+v_{i}(b) \\
= & \left(\sum_{g_{j} \notin b} v_{N \cup\{0 \backslash \backslash i\}}^{k}\left(g_{j}\right)+\sum_{g_{j} \in b} v_{N \cup\{0 \backslash \backslash i\}}^{k}\left(g_{j}\right)\right) \\
& -\sum_{j=1}^{m} v_{N \cup\{0\}}^{k}\left(g_{j}\right)+v_{i}(b) .
\end{aligned}
$$

Obviously, $v_{i}(b)=\sum_{g_{j} \in b} v_{N \cup\{0\}}^{k}\left(g_{j}\right)$. Thus the last two terms simplify to
$-\sum_{j=1}^{m} v_{N \cup\{0\}}^{k}\left(g_{j}\right)+v_{i}(b)=-\sum_{g_{j} \notin b} v_{N \cup\{0\}}^{k}\left(g_{j}\right)$.
For the items that are not allocated to Bidder $i$, we have
$\sum_{g_{j} \notin b} v_{N \cup\{0 \backslash \backslash i\}}^{k}\left(g_{j}\right)=\sum_{g_{j} \notin b} v_{N \cup\{0\}}^{k}\left(g_{j}\right)$.
Therefore

$$
t_{i}=\sum_{g_{j} \in b} v_{N \cup\{0 \backslash \backslash i\}}^{k}\left(g_{j}\right),
$$

which by definition of $v_{N \cup\{0\} \backslash\{i\}}^{k}$ is no less than $|b| \cdot l \cdot 2^{k}$.
Proof of Theorem 2.2 Continued. Since every $V V C A^{k}$ is ex post incentive compatible and ex post individually rational and $M$ is a randomization over $V V C A^{k}, M$ is also ex post incentive compatible and ex post individually rational.

We now prove the revenue bound. By Lemma A.1, any bidder who wins a bundle, $b$, in $V V C A^{k}$, pays at least $|b| \cdot l \cdot 2^{k}$. Because valuations are additive, $a^{k}$ allocates every item $g_{j}$ to the same bidder as $a^{\mathrm{EFF}}$ if $v_{N}\left(g_{j}\right) \geqslant l \cdot 2^{k}$, and leaves the item for the seller otherwise. Therefore the revenue in $V V C A^{k}$ is at least $n_{k} l \cdot 2^{k}$, where $n_{k}$ is the number of such $g_{j}$ that $v_{N}\left(g_{j}\right) \geqslant l \cdot 2^{k}$ :
$n_{k}=\sum_{j=1}^{m} I\left[v_{N}\left(g_{j}\right) \geqslant l \cdot 2^{k}\right]$,
where $I$ is an indicator function, which equals 1 if its argument is true and 0 otherwise.

So, when the valuations of bidders are given by $v$, the expected revenue of mechanism $M, E_{k}\left[R_{M}(v)\right]$, is at least

$$
\begin{align*}
& \frac{1}{1+\lfloor\log (h / l)\rfloor} \sum_{k=0}^{\lfloor\log (h / l)\rfloor} l \cdot 2^{k} \sum_{j=1}^{m} I\left[v_{N}\left(g_{j}\right) \geqslant l \cdot 2^{k}\right] \\
& \quad=\frac{1}{1+\lfloor\log (h / l)\rfloor} \sum_{j=1}^{m} \sum_{k=0}^{\lfloor\log (h / l)\rfloor} I\left[v_{N}\left(g_{j}\right) \geqslant l \cdot 2^{k}\right] l \cdot 2^{k} . \tag{11}
\end{align*}
$$

The sum on the right of (11) can be bounded as follows:

$$
\begin{align*}
v_{N}\left(g_{j}\right) & \leqslant l+\sum_{k=0}^{\lfloor\log (h / l)\rfloor} I\left[v_{N}\left(g_{j}\right) \geqslant l \cdot 2^{k}\right] l \cdot 2^{k} \\
& \leqslant 2 \cdot \sum_{k=0}^{\lfloor\log (h / l)\rfloor} I\left[v_{N}\left(g_{j}\right) \geqslant l \cdot 2^{k}\right] l \cdot 2^{k} . \tag{12}
\end{align*}
$$

Substituting (12) into (11), we obtain
$E_{k}\left[R_{M}(v)\right] \geqslant \frac{1}{2+2\lfloor\log (h / l)\rfloor} \sum_{j=1}^{m} v_{N}\left(g_{j}\right)$.
Here, $\sum_{j=1}^{m} v_{N}\left(g_{j}\right)$ is the welfare of the efficient allocation. No individually rational auction can yield more revenue than that. Therefore the revenue of the optimal auction is bounded from above by $\sum_{j=1}^{m} v_{N}\left(g_{j}\right)$. It follows that
$E_{k}\left[R_{M}(v)\right] \geqslant \frac{R_{\mathrm{OPT}}(v)}{2+2\lfloor\log (h / l)\rfloor}$.
Proof of Corollary 2.1. By construction of $M$ in Theorem 2.2, we have
$E_{k}\left[R_{M}(v)\right]=\frac{1}{1+\lfloor\log (h / l)\rfloor} \sum_{k=0}^{\lfloor\log (h / l)\rfloor} R_{V V C A^{k}}(v)$.
Substituting $E_{k}\left[R_{M}(v)\right]$ into (14) and taking expectations over $v$, we obtain

$$
\frac{\sum_{k=0}^{\lfloor\log (h / l)\rfloor} E_{v}\left[R_{V V C A^{k}}(v)\right]}{1+\lfloor\log (h / l)\rfloor} \geqslant \frac{E_{v}\left[R_{\mathrm{OPT}}(v)\right]}{2+2\lfloor\log (h / l)\rfloor} .
$$

Since the sum of $V V C A^{k}$ contains exactly $1+\lfloor\log (h / l)\rfloor$ terms, there must exist a $k^{*}$ such that
$E_{v}\left[R_{V V C A^{*}}(v)\right] \geqslant \frac{E_{v}\left[R_{\mathrm{OPT}}(v)\right]}{2+2\lfloor\log (h / l)\rfloor}$.
Note that $k^{*}$ can be found by enumerating all $V V C A^{k}$ and evaluating their expected revenues.

Proof of Theorem 2.3. We first scale the problem down by a factor $l$, so the support of each bidder's valuation distribution is $[1,(h / l)]$. In this proof, we will use the shorthand notation $H=h / l$.

Consider an $n$-item auction with $n$ bidders. Assume the valuation of Bidder $i$ for item $i$ is drawn from distribution $F_{i}$ with density
$f_{i}\left(v_{i}\right)= \begin{cases}\frac{H}{(H-1) v_{i}^{2}}, & \text { for } v_{i} \in[1, H], \\ 0 & \text { otherwise. }\end{cases}$
The valuations of other bidders for item $i$ are 0 . The valuations for the bundles are additive.

Here, the expected welfare of an efficient allocation is

$$
\begin{aligned}
E_{v}\left[\mathrm{SW}\left(a^{\mathrm{EFF}}\right)\right] & =E_{v}\left[\sum_{i=1}^{n} v_{i}\right]=\sum_{i=1}^{n} E\left[v_{i}\right] \\
& =n \int_{1}^{H} v_{i} f_{i}\left(v_{i}\right) d v_{i}=(\ln H) \frac{n H}{H-1} .
\end{aligned}
$$

Our example is a special case of single-minded bidders, that is, settings where each bidder has a nonzero valuation for one bundle and all the supersets of this bundle, and 0 valuations for all other bundles. We will use a characterization result for single-minded bidders introduced by Mu'alem and Nisan (2002). We start by reviewing two definitions.

Definition A.2. An auction mechanism $M$ is monotone if for any Bidder $i$ and any bids of the other bidders $v_{-i}$, we have that if $v_{i}$ is a bid that wins bundle $a_{i}$, then any higher bid $v_{i}^{\prime}\left(a_{i}\right)>v_{i}\left(a_{i}\right)$ also wins bundle $a_{i}$.
Definition A.3. A payment scheme (associated with monotone auction mechanism $M$ ) is based on a critical value $p$, if it has the following form: payment $t_{i}=p\left(a_{i}, v_{-i}\right)$ if Bidder $i$ wins $a_{i}$, and $t_{i}=0$ otherwise.

Mu'alem and Nisan (2002) showed that when bidders are single minded, an auction is incentive compatible if and only if it is monotone and its payment scheme is based on a critical value.

Let $M$ be an optimal individually rational incentive compatible mechanism. Consider a mechanism $M^{\prime}$, which is like $M$ except that it holds back the goods allocated to bidders with 0 valuations for them:

1. Whenever $M$ allocates a bundle $b$ such that item $i \in b$ to Bidder $i, M^{\prime}$ allocates item $i$ to Bidder $i$ and holds back the rest of the bundle.
2. Whenever $M$ allocates a bundle $b$ such that item $i \notin b$ to Bidder $i, M^{\prime}$ holds back the entire bundle $b$.
3. The payment scheme is the same as in $M$.

Denote the allocation under mechanism $M$ by $a$, and the allocation under mechanism $M^{\prime}$ by $a^{\prime}$. Then, for any Bidder $i, v_{i}(a)=$ $v_{i}\left(a^{\prime}\right)$. This makes $M^{\prime}$ incentive compatible. Obviously, the revenue of $M^{\prime}$ is the same as that of $M$; therefore $M^{\prime}$ is also optimal.

We now obtain a bound on the expected payment that mechanism $M^{\prime}$ extracts from Bidder $i$ :

$$
E_{v}\left[t_{i}^{M^{\prime}}(v)\right]=E_{v_{-i}}\left[E_{v_{i}}\left[p_{i}^{M^{\prime}}\left(v_{-i}, v_{i}\right)\right]\right]
$$

Since $M^{\prime}$ must be based on a critical value, the latter equals
$E_{v_{-i}}\left[p_{i}^{M^{\prime}}\left(v_{-i}\right) \cdot \operatorname{Prob}\left(\right.\right.$ Bidder $i$ wins $\left.\left.g_{i} \mid v_{-i}\right)\right]$.
Since $M^{\prime}$ is an optimal incentive compatible individually rational auction, it must allocate item $i$ to Bidder $i$ whenever $v_{i}$ is greater than $p_{i}^{M^{\prime}}\left(v_{-i}\right)$ : by construction $M^{\prime}$ either keeps item $i$ or allocates it to Bidder $i$. Also, $M^{\prime}$ cannot sell item $i$ when $v_{i}$ is less than $p_{i}^{M^{\prime}}\left(v_{-i}\right)$; otherwise, the auction will not be individually rational. Finally, if $M^{\prime}$ does not allocate item $i$ to Bidder $i$ when $v_{i}$ is larger than $p_{i}^{M^{\prime}}\left(v_{-i}\right)$, then an auction that differs from $M^{\prime}$ by allocating item $i$ to Bidder $i$ would yield higher revenue while still remaining incentive compatible (since it is still monotone and based on a critical value).

Therefore

$$
\begin{array}{r}
E_{v_{-i}}\left[p_{i}^{M^{\prime}}\left(v_{-i}\right) \cdot \operatorname{Prob}\left(\operatorname{Bidder} i \text { wins } g_{i} \mid v_{-i}\right)\right] \\
=E_{v_{-i}}\left[p_{i}^{M^{\prime}}\left(v_{-i}\right) \cdot \operatorname{Prob}\left(v_{i}>p_{i}^{M^{\prime}}\left(v_{-i}\right)\right)\right]
\end{array}
$$

In this example, the valuations of bidders other than $i$ do not provide any information about $v_{i}$.

Since $M^{\prime}$ is an optimal auction, we have

$$
p_{i}^{M^{\prime}}=\underset{p}{\arg \max }\left[p \cdot\left(1-F_{i}(p)\right)\right]
$$

By substituting $F_{i}$, it is easy to show that $p_{i}^{M^{\prime}}=1$ for all $i$. Therefore the expected revenue of $M^{\prime}$ is $n$.

Since $R_{M}=R_{M^{\prime}}$, we obtain
$\frac{E_{v}\left[R_{M}\right]}{E_{v}\left[\mathrm{SW}\left(a^{\mathrm{EFF}}\right)\right]} \leqslant\left(1-\frac{1}{H}\right) \frac{1}{\ln H}$.
Proof of Theorem 2.4. $M^{\prime}$ is incentive compatible and individually rational because it is a randomization over incentive compatible and individually rational auctions. Let $a^{k}$ be the winning allocation in $V V C A^{\prime k}$ and $a_{-i}^{k}$ be the allocation that would have been optimal had Bidder $i$ not submitted any bids. Since there is no competition, $a_{-i}^{k}$ and $a^{k}$ are the same for all bidders except for Bidder $i$. By construction of $V V C A^{\prime k}$, Bidder $i$ wins the grand bundle $G$ iff $v_{i}(G) \geqslant-l \cdot 2^{k}$ and wins nothing otherwise.

Since $a_{-i}^{k}$ and $a^{k}$ are equivalent for bidders other than $i$, the payment of Bidder $i$ for bundle $G$ is

$$
\begin{aligned}
t_{i} & =\left(\mathrm{SW}_{\lambda}^{\mu}\left(a_{-i}^{k}\right)-\mathrm{SW}_{\lambda}^{\mu}\left(a^{k}\right)\right)+v_{i}(G) \\
& =\left(-v_{i}(G)+l \cdot 2^{k}\right)+v_{i}(G)=l \cdot 2^{k}
\end{aligned}
$$

Denote by $a^{\text {EFF }}$ the (efficient) allocation that allocates a copy of every item to every bidder. Using the notation of Theorem 2.2, the expected revenue of mechanism $M^{\prime}, E_{k}\left[R_{M}(v)\right]$, can be written as

$$
\begin{aligned}
& \frac{1}{1+\lfloor\log (h / l)\rfloor} \sum_{k=0}^{\lfloor\log (h / l)\rfloor} l \cdot 2^{k} \sum_{i=1}^{n} I\left[v_{i}(G) \geqslant l \cdot 2^{k}\right] \\
& \geqslant \frac{\sum_{i=1}^{n} v_{i}(G)}{2+2\lfloor\log (h / l)\rfloor}=\frac{\mathrm{SW}\left(a^{\mathrm{EFF}}\right)}{2+2\lfloor\log (h / l)\rfloor} \\
& \geqslant \frac{R_{\mathrm{OPT}}(v)}{2+2\lfloor\log (h / l)\rfloor} \cdot
\end{aligned}
$$

Proof of Corollary 2.2. Analogous to the proof of Corollary 2.1.
Proof of Theorem 3.1. We will prove the claim for a one-item auction. This is without loss of generality because the argument applies directly to CAs where the bundle of all items is the only bundle that any agent is interested in (i.e., values for all other bundles are zero). Let $M$ be a completely prior-free mechanism. We will construct a distribution $V$ of valuation functions where the ratio of the expected revenues of $M$ and some (suboptimal) incentive compatible individually rational mechanism is less than $\epsilon$.

In order to be incentive compatible for all distributions of valuations, $M$ must have the following property. Assume valuations of Bidders 2 to $n$ are given by a vector $v_{-1}=\left\{v_{2}, \ldots, v_{n}\right\}$. If Bidder 1 with valuation $v_{1}$ wins, the price paid, $p_{1}$, cannot depend on $v_{1}$. Otherwise Bidder 1 with sufficiently high $v_{1}$ would submit the valuation $\arg \min _{\hat{v}_{1} \in V_{1}} p_{1}\left(\hat{v}_{1}, v_{-1}\right)$ rather than her true valuation. Since $M$ is a deterministic prior-free mechanism, $p_{1}$ is a function of $v_{-1}$ only. (We allow $p_{1}\left(v_{-1}\right)$ to be infinite if Bidder 1 never wins for some $v_{-1}$.)

We now fix the distributions of valuations of all bidders except Bidder 1. We let their valuation distributions be bounded; we denote by $h$ the highest valuation that any of those other bidders might have.

Consider now the following set of segments:
$\left\{T_{l}\right\}_{0}^{\infty}=\left[h, \frac{h}{\delta}\right),\left[\frac{h}{\delta}, \frac{h}{\delta^{2}}\right), \ldots,\left[\frac{h}{\delta^{n-1}}, \frac{h}{\delta^{n}}\right), \ldots$
Take
$\delta<\frac{\epsilon}{2-\epsilon}$.

Consider $p_{1}\left(v_{-1}\right)$. Obviously,
$\sum_{l=0}^{\infty} \operatorname{Pr}\left\{p_{1}\left(v_{-1}\right) \in T_{l}\right\} \leqslant 1$.
Therefore, for all $\gamma>0$, there exists some $T_{l(\gamma)}$ such that
$\operatorname{Pr}_{v_{-1}}\left\{p_{1}\left(v_{-1}\right) \in T_{l(\gamma)}\right\} \leqslant \gamma$.

Now, take
$\gamma \leqslant \frac{\epsilon(\delta+1)-4 \delta}{4-4 \delta}$.
It is easy to check that $\gamma>0$ because of our choice of $\delta$. Let $v_{1}$ be uniformly distributed on $T_{l(\gamma)}$. For convenience, denote the left end point of $T_{l(\gamma)}$ by $x_{\gamma}$ :
$T_{l(\gamma)}=\left[x_{\gamma}, \frac{x_{\gamma}}{\delta}\right)$
The expected revenue of OPT is at least the revenue obtainable by making a take-it-or-leave-it offer to Bidder 1 with a price set at the middle of the interval $T_{l(\gamma)}$ :
$\hat{p}=x_{\gamma} \frac{\delta+1}{2 \delta}$.
Since $v_{1}$ is uniformly distributed, the revenue of OPT is at least
$\operatorname{Pr}\left\{v_{1}>\hat{p}\right\} \cdot \hat{p}=\frac{\hat{p}}{2}=x_{\gamma} \frac{\delta+1}{4 \delta}$.
The revenue of $M$ can be calculated as follows: if $M$ does not allocate the item to Bidder 1, by construction of the valuation set, the revenue of the mechanism is at most $x_{\gamma}$.

Since
$\underset{v}{\operatorname{Pr}}\left\{p_{1} \in T_{l(\gamma)}\right\}=\operatorname{Pr}_{v_{-1}}\left\{p_{1}\left(v_{-1}\right) \in T_{l(\gamma)}\right\} \leqslant \gamma$
the probability that Bidder 1 wins is at most $\gamma$. By individual rationality of $M$, whenever Bidder 1 wins, the revenue of $M$ is at most $x_{\gamma} / \delta$.

The expected revenue of $M$ can therefore be bounded as follows:
$E_{v}\left(R_{M}(v)\right) \leqslant \gamma \frac{x_{\gamma}}{\delta}+x_{\gamma}(1-\gamma)$.
Finally, for the chosen $\delta$ and $\gamma$, we obtain

$$
\begin{aligned}
\frac{E_{v}\left(R_{M}(v)\right)}{E_{v}(\mathrm{OPT}(v))} & \leqslant \frac{\gamma\left(x_{\gamma} / \delta\right)+x_{\gamma}}{\frac{1}{4} x_{\gamma}((\delta+1) /(2 \delta))} \leqslant 4 \frac{\delta+\gamma(1-\delta)}{\delta+1} \\
& \leqslant 4 \frac{\delta+\epsilon(\delta+1) / 4-\delta}{\delta+1}=\epsilon .
\end{aligned}
$$

Proof of Theorem 3.2. It is easy to show that the expected revenue of the VVCA and AMA is Lipschitz continuous in all parameters as follows.

Take any $\mu_{j}^{1}$ and $\mu_{j}^{2}$. We have

$$
\begin{aligned}
& \left|\operatorname{SW}_{\lambda}^{\left(\mu_{-j}, \mu_{j}^{1}\right)}(v)-\operatorname{SW}_{\lambda}^{\left(\mu_{-j}, \mu_{j}^{2}\right)}(v)\right| \\
& =\mid \sum_{i \neq j}\left[\mu_{i} v_{i}\left(a_{1}\right)+\lambda_{\left\{i, a_{1}(i)\right\}}\left(a_{1}\right)\right]+\mu_{j}^{1} v_{j}\left(a_{1}\right)+\lambda_{\left\{j, a_{1}(j)\right\}}\left(a_{1}\right) \\
& \quad-\sum_{i \neq j}\left[\mu_{i} v_{i}\left(a_{2}\right)+\lambda_{\left\{i, a_{2}(i)\right\}}\left(a_{2}\right)\right]-\mu_{j}^{2} v_{j}\left(a_{2}\right)-\lambda_{\left\{j, a_{2}(j)\right\}}\left(a_{2}\right) \mid \\
& \leqslant \max \left\{\left|\mu_{j}^{1}-\mu_{j}^{2}\right| v_{j}\left(a_{2}\right),\left|\mu_{j}^{1}-\mu_{j}^{2}\right| v_{j}\left(a_{1}\right)\right\} .
\end{aligned}
$$

The last inequality is by definition of $a_{1}$ and $a_{2}$. Since the valuations are bounded, $E_{v}\left[\mathrm{SW}_{\lambda}^{\mu}\right]$ is Lipschitz continuous (the same property for $\lambda$ can be verified in a similar manner). The same property is true for $E_{v}\left[\mathrm{SW}_{-i}\right]_{\lambda}^{\mu}$ for all $i$. Therefore the expected revenue, $E_{v}(R(\mu, \lambda, v))$, where $R$ is given by (10), is also Lipschitz continuous.

By the Rademacher theorem, this yields the statement of the theorem.

Alternatively, continuity can be shown using convexity of SW. By the definition of $\mathrm{SW}_{\lambda}^{\mu}$, it can be represented as
$\sup _{a \in A} f(\lambda, \mu, a)$,
where
$f(\lambda, \mu, a)=\sum_{j=1}^{n}\left(\mu_{j} v_{j}(a)+\lambda_{\left\{j, a_{j}\right\}}(a)\right)$
$f$ is linear in $(\lambda, \mu)$ and therefore convex in $(\lambda, \mu)$. Since pointwise maximum preserves convexity, $\mathrm{SW}_{\lambda}^{\mu}$ is also convex in $f(\lambda, \mu)$. Also, since integration preserves convexity, $E_{v}\left[\mathrm{SW}_{\lambda}^{\mu}(v)\right]$ is convex and thereby continuous.

Similarly, it can be shown that $E_{v}\left[\left(\mathrm{SW}_{-i}\right)_{\lambda}^{\mu}(v)\right]$ is continuous. $E_{v}[R(\mu, \lambda, v)]$ is then continuous as a linear combination of continuous functions.

Proof of Theorem 3.3. Without loss of generality, assume $\lambda_{1}<$ $\lambda_{2}$. Assume there exists a polynomial-time algorithm $A$ that for some Bidder $i$, some bundle $b$, and arbitrary set of valuations $v$, can tell whether
$R\left(\mu,\left(\lambda_{-\{i, b\}}, \lambda_{1}\right)\right)>R\left(\mu,\left(\lambda_{-\{i, b\}}, \lambda_{2}\right)\right)$.
Also, assume $\mu_{i}=1$.
We now show that $A$ can be used to find a solution to the independent set problem. The proof works as follows. We first show that $A$ can be used to compute the optimal allocation in AMA with single-minded bidders; as defined by Lehmann et al. (2002), this term refers to auctions where each bidder bids for only one bidder-specific bundle. We then show how to convert an instance of the independent set problem into an instance of optimal allocation in affine maximization.

Let $\hat{v}$ be a given set of valuations. We assume that every bidder bids just for one bundle. Moreover, assume that all the bundles have the same number of items, $g$. If this assumption does not hold, augment the smaller bundles with fake items, such that each item is a part of only one bundle-so that there is no competition for fake items.

Denote the set of parameters $\left(\lambda_{-\{i, b\}}, \lambda_{1}\right)$ by $\bar{\lambda}_{1}$ and $\left(\lambda_{-\{i, b\}}, \lambda_{2}\right)$ by $\bar{\lambda}_{2}$. Define $a\left(\lambda_{1}\right)$ to be an optimal allocation under $\bar{\lambda}_{1}$ :

$$
a\left(\lambda_{1}\right)=\underset{a \in A}{\arg \max } \operatorname{SW}_{\bar{\lambda}_{1}}^{\mu}(\hat{v}),
$$

and define $a\left(\lambda_{2}\right)$ to be an optimal allocation under $\bar{\lambda}_{2}$ :
$a\left(\lambda_{2}\right)=\underset{a \in A}{\arg \max } \mathrm{SW}_{\bar{\lambda}_{2}}^{\mu}(\hat{v})$.
We now show that $A$ can be used to compute the optimal (in terms of $\operatorname{SW}_{\lambda}^{\mu}(\hat{v})$ ) allocation $a\left(\lambda_{1}\right)$.

We start by proving the following proposition.
Proposition A.1. If A outputs
$R\left(\mu, \bar{\lambda}_{2}\right) \geqslant R\left(\mu, \bar{\lambda}_{1}\right)$,
then $b$ is not part of $a\left(\lambda_{1}\right)$. If A outputs
$R\left(\mu, \bar{\lambda}_{2}\right)<R\left(\mu, \bar{\lambda}_{1}\right)$,
then the weighted $S W$, corresponding to the allocation where $b$ is given to $i$ and other bundles are allocated optimally, is at most $\Delta \lambda=\lambda_{2}-\lambda_{1}$ suboptimal to the weighted SW of $a\left(\lambda_{1}\right)$.
Proof. Consider the following three cases:

1. $a\left(\lambda_{1}\right)$ and $a\left(\lambda_{2}\right)$ allocate $b$ to $i$.
2. Neither $a\left(\lambda_{1}\right)$ nor $a\left(\lambda_{2}\right)$ allocates $b$ to $i$.
3. $a\left(\lambda_{2}\right)$ allocates $b$ to $i$, but $a\left(\lambda_{1}\right)$ does not.

Clearly, since $\lambda_{2}>\lambda_{1}$, these are the only three cases possible (if $b$ is a part of an optimal allocation under $\lambda_{\{i, b\}}=\lambda_{1}$, it is also a part of optimal allocation under $\lambda_{\{i, b\}}=\lambda_{2}$ ).

Consider the revenue of the seller, given by (10). In Case 1, $b$ is part of $a\left(\lambda_{1}\right)$ and $a\left(\lambda_{2}\right)$, which means that $v_{i}(b)+\lambda_{1}$ is high enough to make $b$ a part of the optimal allocation (same goes for $v_{i}(b)+\lambda_{2}$ ). The optimal allocation of $G \backslash\{b\}$ is the same; therefore we have $a\left(\lambda_{1}\right)=a\left(\lambda_{2}\right)$ (the allocations are the same). Similarly, $a_{-j}\left(\lambda_{1}\right)=a_{-j}\left(\lambda_{2}\right)$.

Therefore
$\mathrm{SW}_{\bar{\lambda}_{2}}^{\mu}(\hat{v})=\operatorname{SW}_{\bar{\lambda}_{1}}^{\mu}(\hat{v})+\Delta \lambda$.
Also,
$\left[\mathrm{SW}_{-i}\right]_{\bar{\lambda}_{2}}^{\mu}(\hat{v})=\left[\mathrm{SW}_{-i} i_{\bar{\lambda}_{1}}^{\mu}(\hat{v})\right.$
and for $j \neq i$,
$\left[\mathrm{SW}_{-j}\right]_{\lambda_{2}}^{\mu}(\hat{v}) \leqslant\left[\mathrm{SW}_{-j}\right]_{\bar{\lambda}_{1}}^{\mu}(\hat{v})+\Delta \lambda$.
Substituting the above equalities in (10), we get that in Case 1, $R\left(\mu, \bar{\lambda}_{1}\right)-R\left(\mu, \bar{\lambda}_{2}\right) \geqslant \Delta \lambda>0$.

In Case 2, since $b$ is not part of $a$ under either values of parameters,
$\operatorname{SW}_{\bar{\lambda}_{2}}^{\mu}(\hat{v})=\operatorname{SW}_{\bar{\lambda}_{1}}^{\mu}(\hat{v})$
and $\sum_{i=1}^{n} \lambda_{\{i, a(i)\}}$ is the same in both cases. Also, by definition, $\mathrm{SW}_{\lambda}^{\mu}$ is increasing in any of the $\lambda$ parameters; therefore
$\sum_{j=1}^{n}\left[\mathrm{SW}_{-j}\right]_{\bar{\lambda}_{2}}^{\mu}(\hat{v}) \geqslant \sum_{j=1}^{n}\left[\mathrm{SW}_{-j}\right]_{\bar{\lambda}_{1}}^{\mu}(\hat{v})$
So in Case 2,
$R\left(\mu, \bar{\lambda}_{1}\right)-R\left(\mu, \bar{\lambda}_{2}\right) \leqslant 0$.
Therefore, depending on the output of $A$, we either conclude that $i$ does not get $b$ in $a\left(\lambda_{1}\right)$, or that $i$ does get $b$ in $a\left(\lambda_{2}\right)$. In the latter case, allocating $b$ to $i$ under $\left(\mu, \bar{\lambda}_{1}\right)$ decreases the optimal weighted SW by at most $\Delta \lambda$. This proves the proposition.

We can use the following algorithm to find whether $i$ should obtain $b$ in the optimal allocation.

Algorithm 4 (Determines whether $i$ should get $b$ )

1. Run $A$.
2. If $A$ outputs $R\left(\mu, \bar{\lambda}_{1}\right)>R\left(\mu, \bar{\lambda}_{2}\right)$, then allocate $b$ to $i$, otherwise do not allocate.
Proposition A. 1 shows that if this algorithm makes mistakes, the SW goes down by at most $\Delta \lambda$.

Algorithm $A$ can be used to determine whether any of the other Bidders $i^{\prime}$ gets her preferred bid (recall that we consider singleminded bidders). Suppose Bidder $i^{\prime}$ bids for bundle $b^{\prime}$. Since we transformed the problem into one where all the bidders bid for bundles of the same size, $|b|=\left|b^{\prime}\right|=g$. Rename the items so that $b^{\text {new }}=b^{\prime}$ and $\left[b^{\prime}\right]^{\text {new }}=b$ (it is always possible to do that). Now, change $\hat{v}_{i}(b)$ so that the new value of
$\left[\hat{v}_{i}^{\text {new }}\right]_{\lambda_{1}}^{\mu}\left(b^{\text {new }}\right)=\mu_{i} \hat{v}_{i}^{\text {new }}\left(b^{\text {new }}\right)+\lambda_{\{i, b\}}$
equals the old value for $\left[\hat{v}_{i^{\prime}}\right]_{\bar{\lambda}_{1}}^{\mu}\left(b^{\prime}\right)$ by setting
$\hat{v}_{i}^{\text {new }}=\frac{\mu_{i^{\prime}} \hat{v}_{i^{\prime}}\left(b^{\prime}\right)+\lambda_{\left\{i^{\prime}, b^{\prime}\right\}}-\lambda_{\{i, b\}}}{\mu_{i}}$.
Similarly, modify $\hat{v}_{i^{\prime}}\left(b^{\prime}\right)$ :
$\hat{v}_{i^{\prime}}^{\text {new }}=\frac{\mu_{i} \hat{v}_{i}(b)+\lambda_{\{i, b\}}-\lambda_{\left\{i^{\prime}, b^{\prime}\right\}}}{\mu_{i^{\prime}}}$.
This operation effectively interchanges $i$ and $i^{\prime}$ and now Algorithm $A$ can be used to find whether $i^{\prime}$ obtains $b^{\prime}$ in $a\left(\lambda_{1}\right)$. The only problem with such a transformation is that the new real (but not virtual) valuations might become negative (we explain how to deal with this later in the proof).

We now show that Algorithm $A$ can be used to compute $a\left(\lambda_{1}\right)$.

Proposition A.2. Algorithm A can be used to compute $a\left(\lambda_{1}\right)$ in polynomial time.
Proof. Since there are at most $n$ bids in the auction, we can run Algorithm 4 on each one of those bundles to determine which bundles are parts of optimal allocation $a\left(\lambda_{1}\right)$. The only problem here is that Algorithm 4 can make mistakes when telling us to allocate the bundle, reducing the value of the allocation by at most $\Delta \lambda$, and more importantly, possibly yielding an invalid allocation.

We can deal with this problem as follows. Define $\epsilon$ to be the minimal positive difference in weighted SW of any two allocations:

$$
\begin{aligned}
\epsilon= & \min \left\{\operatorname{SW}_{\bar{\lambda}_{1}}^{\mu}\left(a_{1}\right)-\operatorname{SW}_{\bar{\lambda}_{1}}^{\mu}\left(a_{2}\right)\right\}, \quad \text { where } a_{1}, a_{2} \in A, \\
& \text { s.t. } \operatorname{SW}_{\bar{\lambda}_{1}}^{\mu}\left(a_{1}\right)>\operatorname{SW}_{\bar{\lambda}_{1}}^{\mu}\left(a_{2}\right)
\end{aligned}
$$

Now, we make the following transformation of all real valuations of all bidders:
$\hat{v}_{j}^{\text {new }}=\frac{c \mu_{j} \hat{v}_{j}+(c-1) \lambda_{\left\{j, b_{j}\right\}}}{\mu_{j}}$.
Here, $b_{j}$ is the bundle that Bidder $j$ bid on, $\lambda_{\left\{j, b_{j}\right\}}$ is the corresponding parameter, and $c$ is some constant.

This transformation increases all virtual valuations by the same factor $c$. Trivially, this increases $\epsilon$ to $c \epsilon$. Choose $c$ to be larger than $\Delta \lambda / \epsilon$. After transformation (17), the weighted SW of all allocations differ by more than $\Delta \lambda$.

Note that since all the valuations increase by factor $c$, such an operation does not change the optimal allocation $a\left(\lambda_{1}\right)$. But we know that if Algorithm 4 makes a mistake and chooses the bundle which is not a part of $a\left(\lambda_{1}\right)$, the weighted SW decreases by at most $\Delta \lambda$, which means that on these modified valuations, Algorithm 4 cannot choose suboptimal allocations.

Therefore, by applying transformation (17) and running Algorithm 4, we can compute the optimal allocation $a(\lambda)$.

The "trick" with transformation (17) can be also applied to deal with negative real valuations in Proposition A.1.

Finally, Algorithm 4 can be used to solve the independent set problem as follows. We create an item for each edge of the graph. We create a bid for each vertex of the graph; the items included in the bid correspond to the edges that are connected to the vertex. We set the valuations for bundles so that virtual valuations are equal to some constant $c$ (since virtual valuations are increasing functions of real valuations, it is always possible to do that). The affine maximizing allocation yields the maximum independent set of the graph.

## Endnotes

1. This paper combines and extends upon two short conference papers (Likhodedov and Sandholm 2004, 2005).
2. For the mechanism to be incentive compatible, $\tilde{v}_{i}$ should be increasing in $v_{i}$. If (3) does not satisfy this condition, an "ironing" technique is used to make $\tilde{v}_{i}$ nondecreasing (Myerson 1981).
3. The basic idea of adding an allocation-specific constant to the objective was also discussed by Jehiel et al. (2007) in parallel with our work (Likhodedov and Sandholm 2004, 2005), for the purpose of tuning the bundling policy. Their $\lambda$-auction is a special case of the VVCA where (1) no weights are used ( $\mu_{i}=1$ for all $i$ ) and (2) the same constant is added to the objective whenever all items are sold to the same bidder (i.e., $\lambda_{i}(G)=c$ for the grand bundle $G$, and $\lambda_{i}(b)=0$ for all other bundles $b$ ). In a symmetric additive valuations model, they show that the $\lambda$-auction can increase the revenue over both pure bundling auctions (which always sells all the items together) and separate auctioning of individual items (this is what the VCG does in the additive valuations model). Our work differs from theirs in several ways: (1) we study a significantly larger family of mechanisms that subsumes the family they study; (2) unlike us, they assume that each bidder's valuation is additive, that is, no bidder considers any items as complements or substitutes; (3) we ask different questions; and (4) our main thrust is on automated design rather than manual analysis.
4. A mechanism is an almost affine maximizer if it is an affine maximizer for sufficiently high valuations. Lavi et al. (2003) conjecture that the "almost" qualifier is merely technical, and can be removed in future research.
5. Throughout this paper, ties in allocation rules can be broken arbitrarily.
6. In a recent personal email communication (8/30/2008), Jason Hartline points out that if one knows the highest possible valuation, $h$, and the lowest possible valuation, $l$, then one can achieve, in expectation over random bits, a $\log (h / l)$-approximation to revenue by running the VCG first, and then uniformly increasing every winner's price by $2^{k}$ (where $k$ is randomly chosen in $[l, h]$ ), and then giving each winner the option to not buy. A recent technical report shows that even with knowledge of those two
bounds, one cannot achieve a better than logarithmic approximation (Micali and Valiant 2008). Both of those results are subsequent work to the conference versions of our paper (Likhodedov and Sandholm 2004, 2005).
7. In the additive valuations setting, the optimal allocation is trivial to find: every item is sold to the bidder with the highest virtual valuation $\left[\tilde{v}_{i}\right]_{\lambda}^{\mu}$ for that item. If that virtual valuation is negative, the seller keeps the item.
8. Naturally, one can fix one of the $\mu$ parameters in AMA (e.g., $\mu_{1}=1$ ) without loss of generality because the other $\mu$ parameters and the $\lambda$ parameters can be scaled multiplicatively accordingly. In addition, one can fix any one of the $\lambda$ parameters in AMA without loss of generality (for example, the one corresponding to the seller keeping all the items) because that simply amounts to adding a constant to the value of each allocation, which affects neither the allocation nor the payments. However, if one does this, one may have to allow for negative values for some of the other $\lambda$ parameters. Similarly, in VVCA, one can fix one $\mu$ parameter (e.g., $\mu_{1}=1$ ) and one $\lambda$ parameter (e.g., $\lambda_{1}(\varnothing)=0$ ) without loss of generality. If one imposes additional restrictions on the mechanism (for example, item symmetry), that fixes more of the variables and the search for a mechanism can take place in a lower-dimensional space.
9. All of the algorithms that we introduce in the rest of this section are motivated by the idea of minimizing this "revenue loss" (difference between the bidder's valuation and her payment). This economic motivation is intuitively desirable and the algorithms based on it perform quite well, as our experiments will show. However, such algorithms can get stuck in local optima because minimizing revenue loss is not the same as maximizing revenue. As an extreme example, the allocation where no items are allocated to any bidders has zero revenue loss, yet it has zero revenue. 10. We do not perform a similar comparison in Settings II and III since in those settings there is a nonzero complementarity parameter. In that case, it would no longer be a dominant strategy for the bidders to report their true valuations in the individual auctions since their valuation can depend on the outcome of the auction for the other item.
10. In these experiments, we do not fix any lambda parameter or any $\mu_{i}$ although one of each could be fixed without loss of generality, as was discussed in $\S 3.2$.
11. The affinely transformed winner determination problem here in the case of sparse input representation is NP-complete. It is obviously in NP, and it is NP-hard because the regular sparse CA winner determination problem, a special case, is NP-complete (Rothkopf et al. 1998).

## References

Adams WJ, Yellen JL (1976) Commodity bundling and the burden of monopoly. Quart. J. Econom. 90(3):475-498.
Armstrong M (2000) Optimal multi-object auctions. Rev. Econom. Stud. 67:455-481.
Avery C, Hendershott T (2000) Bundling and optimal auctions of multiple products. Rev. Econom. Stud. 67:483-497.
Bakos Y, Brynjolfsson E (1999) Bundling information goods: Pricing, profits, and efficiency. Management Sci. 45(12):1613-1630.
Balcan MF, Blum A, Mansour Y (2008b) Item pricing for revenue maximization. Proc. ACM Conf. Electronic Commerce (ACM-EC), Chicago, IL (ACM, New York), 50-59.
Balcan MF, Blum A, Hartline J, Mansour Y (2008a) Reducing mechanism design to algorithm design via machine learning. J. Comput. System Sci. 74:78-89.

Bartal Y, Gonen R, Nisan N (2003) Incentive compatible multi-unit combinatorial auctions. Theoretical Aspects of Rationality and Knowledge (TARK) (Bloomington, IN).
Benisch M, Sandholm T (2011) A framework for automated bundling and pricing using purchase data. Proc. Conf. Auctions, Market Mechanisms and Their Appl., AMMA '11 (Springer, Berlin), 40-52.
Chawla S, Hartline J, Kleinberg R (2007) Algorithmic pricing via virtual valuations. Proc. ACM Conf. Electronic Commerce, EC '07 (ACM, New York), 243-251.
Clarke EH (1971) Multipart pricing of public goods. Public Choice 11: 17-33.
Conitzer V, Sandholm T (2002) Complexity of mechanism design. Proc. 18th Annual Conf. Uncertainty in Artificial Intelligence, UAI '02 (Morgan Kaufmann, San Francisco), 103-110.
Conitzer V, Sandholm T (2003) Applications of automated mechanism design. UAI-03 Workshop on Bayesian Modeling Appl., Acapulco, Mexico.
Conitzer V, Sandholm T (2004) Self-interested automated mechanism design and implications for optimal combinatorial auctions. Proc. ACM Conf. Electronic Commerce, ACM-EC '04 (ACM, New York), 132-141.
Conitzer V, Sandholm T (2006) Failures of the VCG mechanism in combinatorial auctions and exchanges. Internat. Conf. Autonomous Agents and Multi-Agent Systems, AAMAS '06 (ACM, New York), 521-528.
Cramton P, Shoham Y, Steinberg R (2006) Combinatorial Auctions (MIT Press, Boston).
Goldberg A, Hartline J, Wright A (2001) Competitive auctions and digital goods. Annual ACM-SIAM Sympos. Discrete Algorithms, SODA '01 (ACM, New York), 735-744.
Groves T (1973) Incentives in teams. Econometrica 41:617-631.
Guruswami V, Hartline J, Karlin A, Kempe D, Kenyon C, McSherry F (2005) On profit-maximizing envy-free pricing. Annual ACMSIAM Sympos. Discrete Algorithms, SODA '05 (SIAM, Philadelphia), 1164-1173.
Hartline J, Karlin A (2007) Profit maximization in mechanism design. Algorithmic Game Theory, Chapter 13 (Cambridge University Press, Cambridge, UK), 331-362.
Hartline J, Roughgarden T (2008) Optimal mechanism design and money burning. Proc. Annual Sympos. Theory of Comput., STOC '08 (ACM, New York), 75-84.
Hitt LM, Chen P (2005) Bundling with customer self-selection: A simple approach to bundling low-marginal-cost goods. Management Sci. 51(10):1481-1493.
Jehiel P, Meyer-Ter-Vehn M, Moldovanu B (2007) Mixed bundling auctions. J. Econom. Theory 127(1):494-512.
Krishna V (2002) Auction Theory (Academic Press, San Diego).
Kroer C, Sandholm T (2015) Computational bundling for auctions. Internat. Conf. Autonomous Agents and Multi-Agent Systems, AAMAS ' 15 (IFAAMAS, Richland County, SC), 317-326. Extended version: Carnegie Mellon University, Computer Science Department technical report CMU-CS-13-111, 2013.
Lavi R, Mu'alem A, Nisan N (2003) Towards a characterization of truthful combinatorial auctions. Proc. Annual Sympos. Foundations of Comput. Sci., FOCS '03 (IEEE Computer Society, Washington, DC), 574-583.
Ledyard JO (2007) Optimal combinatoric auctions with single-minded bidders. Proc. ACM Conf. Electronic Commerce, ACM-EC '07 (ACM, New York), 237-242.
Lehmann D, O’Callaghan LI, Shoham Y (2002) Truth revelation in rapid, approximately efficient combinatorial auctions. J. ACM 49(5): 577-602.
Levin J (1997) An optimal auction for complements. Games Econom. Behav. 18:176-192.
Likhodedov A, Sandholm T (2004) Methods for boosting revenue in combinatorial auctions. Proc. National Conf. Artificial Intelligence, AAAI '04 (AAAI Press, Palo Alto, CA), 232-237.

Likhodedov A, Sandholm T (2005) Approximating revenue-maximizing combinatorial auctions. Proc. National Conf. Artificial Intelligence, AAAI '05 (AAAI Press, Palo Alto, CA), 267-274.
Maskin E, Riley J (1989) Optimal multi-unit auctions. Hahn F, ed. The Economics of Missing Markets, Information, and Games, Chapter 14 (Clarendon Press, Oxford, UK), 312-335.
McAfee RP, McMillan J, Whinston MD (1989) Multiproduct monopoly, commodity bundling, and correlation of values. Quart. J. Econom. 104(2):371-383.
Micali S, Valiant P (2008) Revenue in truly combinatorial auctions and adversarial mechanism design. Technical report, MIT, MIT-CSAIL-TR-2008-039, June 26.
Mitchell T (1997) Machine Learning (McGraw-Hill, New York).
Monderer D, Tennenholtz M (2005) Asymptotically optimal multi-object auctions. Mimeo.
Mu'alem A, Nisan N (2002) Truthful approximate mechanisms for restricted combinatorial auctions. Proc. National Conf. Artificial Intelligence, AAAI '02 (AAAI Press, Palo Alto, CA), 379-384.
Myerson R (1981) Optimal auction design. Math. Oper. Res. 6:58-73.
Myerson R, Satterthwaite M (1983) Efficient mechanisms for bilateral trading. J. Econom. Theory 28:265-281.
Palfrey T (1983) Bundling decisions by a multiproduct monopolist with incomplete information. Econometrica 51:463-484.
Roberts K (1979) The characterization of implementable social choice rules. Laffont J-J, ed. Aggregation and Revelation of Preferences (North-Holland, Amsterdam), 312-348.
Rothkopf M, Pekeč A, Harstad R (1998) Computationally manageable combinatorial auctions. Management Sci. 44(8):1131-1147.
Sandholm T (2002) Algorithm for optimal winner determination in combinatorial auctions. Artificial Intelligence 135:1-54.
Sandholm T (2006) Optimal winner determination algorithms. Cramton P, Shoham Y, Steinberg R, eds. Combinatorial Auctions, Chapter 14 (MIT Press, Cambridge, MA), 337-368.
Sandholm T (2007) Expressive commerce and its application to sourcing: How we conducted $\$ 35$ billion of generalized combinatorial auctions. AI Magazine 28(3):45-58.
Sandholm T (2013) Very-large-scale generalized combinatorial multiattribute auctions: Lessons from conducting $\$ 60$ billion of sourcing. Neeman Z, Roth A, Vulkan N, eds. Handbook of Market Design (Oxford University Press, Oxford, UK), 379-412.
Stoer J, Bulirsch R (1980) Introduction to Numerical Analysis (Springer, New York).
Tang P, Sandholm T (2012) Mixed-bundling auctions with reserve prices. Internat. Conf. Autonomous Agents and Multi-Agent Systems, AAMAS '12, (ACM, New York), 729-736.
Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. J. Finance 16:8-37.
Vohra R (2001) Research problems in combinatorial auctions. Mimeo, version Oct. 29.
Wu S, Hitt LM, Chen P, Anandalingam G (2008) Customized bundle pricing for information goods: A nonlinear mixed-integer programming approach. Management Sci. 54(3):608-622.

Tuomas Sandholm is a professor in the Computer Science Department at Carnegie Mellon University. He is also a serial entrepreneur. His research interests include market design, optimization, algorithms for solving games, advertising markets, kidney exchange, prediction markets, (automated) market making, voting, preference elicitation, coalition formation, and normative models of bounded rationality.

Anton Likhodedov is head of emerging markets derivatives and flow credit trading at Deutsche Bank.


[^0]:    Notes. The highest expected revenue over 10 runs is reported and the average expected revenue is reported in parenthesis. Training set size was 1,000 and test set size $10,000,000$. In each of the grid searches $\left(A M A^{*}, A M A_{b s y m}^{*}\right.$, and $\left.V V C A^{*}\right)$, five consecutively finer grids were used, each being a hyperrectangle around the optimal parameter vector from the previous grid. At each grid resolution, each search dimension used five grid points. The results in italics are for the variant where we fixed $\mu_{1}=1$ and the first $\lambda$ parameter to zero. (For the earlier, nonitalicized results, we used a slightly different setup. Training set size was 400 and test set size 250,000 . For $A M A^{*}$ and $V V C A^{*}$, we used five grid points on each dimension and four iterations. For $A M A_{b s y m}^{*}$, we used 10 grid points on each dimension and 3 iterations.)

