Truth and Complexity in Allocation Games

Dave Buchfuhrer

April 20, 2010

Outline

Outline

Outline

What is an auction?

An auction is a mechanism by which to divvy up a collection of items among 2 or more bidders. An auction problem can be parameterized by:

Auction Parameters

- The number of bidders n
- The number of items m
- The valuations $v_1, \ldots, v_n$ of the bidders for subsets of items
- The type of payments allowed.
- The goal of the mechanism
- The solution concept

Truth

Auction Parameters

- The number of bidders n
- The number of items m
- The valuations $v_1, \ldots, v_n$ of the bidders for subsets of items
- The type of payments allowed.
- The goal of the mechanism
- The solution concept

Truth

The solution concept we use is called truthfulness.

Truth

The solution concept we use is called truthfulness.

Definition (Direct Revelation)

A mechanism M is a direct revelation mechanism if it has the following form:

1. Bidders supply a description of their valuation functions v_i
2. The mechanism determines an allocation and payments

Truth

The solution concept we use is called truthfulness.

Definition (Direct Revelation)

A mechanism M is a direct revelation mechanism if it has the following form:

1. Bidders supply a description of their valuation functions v_i
2. The mechanism determines an allocation and payments

Definition (Truthful)

A direct revelation mechanism M is truthful if each bidder's utility (value minus payments) is maximized by reporting v_i truthfully for any fixed reporting of the other bidders' valuations.

Goals

Auction Parameters

- The number of bidders n
- The number of items m
- The valuations $v_1, \ldots, v_n$ of the bidders for subsets of items
- The type of payments allowed.
- The goal of the mechanism
- The solution concept

Goals

We are interested in two goals.

Definition (Social Welfare)

The social welfare of an allocation $S_1, \ldots, S_n$ is

$$\sum_{i \in [n]} v_i(S_i)$$

Definition (Revenue)

The revenue of an allocation with prices $p_1, \ldots, p_n$ is

$$\sum_{i \in [n]} p_i$$

Payments

Auction Parameters

- The number of bidders n
- The number of items m
- The valuations $v_1, \ldots, v_n$ of the bidders for subsets of items
- The type of payments allowed.
- The goal of the mechanism
- The solution concept

We allow any positive payments for this part of the talk.

Number of items

Auction Parameters

- The number of bidders n
- The number of items m
- The valuations $v_1, \ldots, v_n$ of the bidders for subsets of items
- The type of payments allowed.
- The goal of the mechanism
- The solution concept

Number of items

For a single item, the standard Vickrey auction is truthful:

Example (Vickrey Auction)

Allocation: The item goes to the highest bidder
Payments: The highest bidder pays the second-highest price

Number of items

For a single item, the standard Vickrey auction is truthful:

Example (Vickrey Auction)

Allocation: The item goes to the highest bidder
Payments: The highest bidder pays the second-highest price

The VCG mechanism is a more general way to get truthful auctions. For the auctions in this talk, all VCG-based mechanisms are maximal-in-range (S. Dobzinski and N. Nisan. Limitations of VCG-based Mechanisms. *STOC*. 2007).

Number of items

For a single item, the standard Vickrey auction is truthful:

Example (Vickrey Auction)

Allocation: The item goes to the highest bidder
Payments: The highest bidder pays the second-highest price

The VCG mechanism is a more general way to get truthful auctions. For the auctions in this talk, all VCG-based mechanisms are maximal-in-range (S. Dobzinski and N. Nisan. Limitations of VCG-based Mechanisms. *STOC*. 2007).

Definition (Maximal-in-Range)

Any mechanism M has a range of possible outcomes R. M is maximal-in-range if it maximizes the social welfare over R.

Valuation Functions

Auction Parameters

- The number of bidders n
- The number of items m
- The valuations $v_1, \ldots, v_n$ of the bidders for subsets of items
- The type of payments allowed.
- The goal of the mechanism
- The solution concept

Valuation Functions

For general $v_1 \ldots, v_n$, the problem is NP-hard to approximate to within a constant factor and requires exponential communication in the worst case.

Valuation Functions

For general $v_1 \ldots, v_n$, the problem is NP-hard to approximate to within a constant factor and requires exponential communication in the worst case.

We consider functions which

- are submodular (easy to approximate to within $1 - 1/e$)
- can be described succinctly (no communication issues)

Number of Bidders

Auction Parameters

- The number of bidders n
- The number of items m
- The valuations $v_1, \ldots, v_n$ of the bidders for subsets of items
- The type of payments allowed.
- The goal of the mechanism
- The solution concept

Number of Bidders

If the number of bidders is very high, say $n = m^m$, then we can solve the auction by brute force:

- Enumerate all partitions of items into several bundles
- For each set of bundles, assign them optimally via bipartite matching

So we require that $n \in O(poly(m))$.

Auction Setting

We considered auctions with the following parameters:

Our Parameters

- There are $n = n(m) \in O(poly(m))$ bidders
- Any number m of items
- $v_1, \ldots, v_n$ restricted to a subset of submodular valuations
- Any positive payments are allowed
- The goal is to maximize social welfare $(\sum_i v_i(S_i))$
- The mechanism must be truthful

In order to achieve positive payments with a truthful mechanism, we restrict our attention to maximal-in-range mechanisms.

Auction Setting

We considered auctions with the following parameters:

Our Parameters

- There are $n = n(m) \in O(poly(m))$ bidders
- Any number m of items
- $v_1, \ldots, v_n$ restricted to a subset of submodular valuations
- ~~Any positive payments are allowed~~
- The goal is to maximize social welfare ($\sum_i v_i(S_i)$)
- ~~The mechanism must be truthful~~

In order to achieve positive payments with a truthful mechanism, we restrict our attention to maximal-in-range mechanisms.

Auction Setting

We considered auctions with the following parameters:

Our Parameters

- There are $n = n(m) \in O(poly(m))$ bidders
- Any number m of items
- $v_1, \ldots, v_n$ restricted to a subset of submodular valuations
- The goal is to maximize social welfare $(\sum_i v_i(S_i))$
- The mechanism must be maximal-in-range

In order to achieve positive payments with a truthful mechanism, we restrict our attention to maximal-in-range mechanisms.

Result

Theorem (BDFKMPSSU 10)

An n bidder, m item maximal-in-range auction mechanism can't beat the approximation ratio $\min(n, \sqrt{m})$ *unless* $NP \subseteq P/poly$.

Result

Theorem (BDFKMPSSU 10)

An n bidder, m item maximal-in-range auction mechanism can't beat the approximation ratio $\min(n, \sqrt{m})$ *unless* $NP \subseteq P/poly$.

We based our work on (C. Papadimitriou, M. Schapira and Y. Singer. On the Hardness of Being Truthful. *FOCS*. 2008):

- Show that a better than $\sqrt{m}$ approximation implies a large range
- Sauer's lemma implies a large VC dimension
- A large VC dimension implies a subset of items allocated in every possible way
- An MIR mechanism solves exactly on these items

Result

Theorem (BDFKMPSSU 10)

An n bidder, m item maximal-in-range auction mechanism can't beat the approximation ratio $\min(n, \sqrt{m})$ *unless* $NP \subseteq P/poly$.

We based our work on (C. Papadimitriou, M. Schapira and Y. Singer. On the Hardness of Being Truthful. *FOCS*. 2008):

- Show that a better than $\sqrt{m}$ approximation implies a large range
- Sauer's lemma implies a large VC dimension
- A large VC dimension implies a subset of items allocated in every possible way
- An MIR mechanism solves exactly on these items

Sauer's lemma is useful when the allocation are in $\{0, 1\}^m$, but if items can be unallocated, the allocations could be anything in $\{0, 1, 2\}^m$.

Sketch of our proof

Theorem

An n bidder, m item maximal-in-range auction mechanism can't beat the approximation ratio $\min(n, \sqrt{m})$ *unless* $NP \subseteq P/poly$.

Sketch of Proof.

- Assume an approximation better than $\min(n, \sqrt{m})$

$\square$

Sketch of our proof

Theorem

An n bidder, m item maximal-in-range auction mechanism can't beat the approximation ratio $\min(n, \sqrt{m})$ *unless* $NP \subseteq P/poly$.

Sketch of Proof.

- Assume an approximation better than $\min(n, \sqrt{m})$
- Restrict the auction to a subset of items that are always allocated with a large range

□

Sketch of our proof

Theorem

An n bidder, m item maximal-in-range auction mechanism can't beat the approximation ratio $\min(n, \sqrt{m})$ *unless* $NP \subseteq P/poly$.

Sketch of Proof.

- Assume an approximation better than $\min(n, \sqrt{m})$
- Restrict the auction to a subset of items that are always allocated with a large range
- Combine all but one bidder into a single meta-bidder

□

Sketch of our proof

Theorem

An n bidder, m item maximal-in-range auction mechanism can't beat the approximation ratio $\min(n, \sqrt{m})$ *unless* $NP \subseteq P/poly$.

Sketch of Proof.

- Assume an approximation better than $\min(n, \sqrt{m})$
- Restrict the auction to a subset of items that are always allocated with a large range
- Combine all but one bidder into a single meta-bidder
- Make use of existing Sauer's lemma machinery □

Sketch of our proof

Theorem

An n bidder, m item maximal-in-range auction mechanism can't beat the approximation ratio $\min(n, \sqrt{m})$ *unless* $NP \subseteq P/poly$.

Sketch of Proof.

- Assume an approximation better than $\min(n, \sqrt{m})$
- Restrict the auction to a subset of items that are always allocated with a large range
- Combine all but one bidder into a single meta-bidder
- Make use of existing Sauer's lemma machinery □

D. Buchfuhrer, S. Dughmi, H. Fu, R. Kleinberg, E. Mossel, C. Papadimitriou, M. Schapira, Y. Singer and C. Umans. Inapproximability for VCG-Based Combinatorial Auctions. *SODA*. 2010

NP-Hard 2-Bidder Auction

We noted that a 2-bidder auction with 1 *additive* bidder and 1 *budget-additive* bidder is NP-hard.

NP-Hard 2-Bidder Auction

We noted that a 2-bidder auction with 1 *additive* bidder and 1 *budget-additive* bidder is NP-hard.

Definition (Budget-additive)

A budget-additive bidder has value v_i for item i and a budget B. His valuation function is

$$v(S) = \min\left(\sum_{i \in S} v_i, B\right)$$

Additive and Budget Additive

We simulate these 2 bidders in an n-bidder budget-additive auction

- One bidder i^* is chosen to represent the budget-additive bidder
- All other bidders have valuation equal to the additive bidder

Additive and Budget Additive

We simulate these 2 bidders in an n-bidder budget-additive auction

- One bidder i^* is chosen to represent the budget-additive bidder
- All other bidders have valuation equal to the additive bidder

We showed that we can restrict to a large subset of the items for which the mechanism range includes every full assignment.

Outline

Public Projects?

A public project is similar to an auction, but all items are shared, and we can only choose k items.

Public Projects?

A public project is similar to an auction, but all items are shared, and we can only choose k items.

Definition (Social Welfare for Public Projects)

In a public project where players have valuations $v_1, \ldots, v_n$, a set S has social welfare

$$\sum_{i \in [n]} v_i(S)$$

Public Projects?

A public project is similar to an auction, but all items are shared, and we can only choose k items.

Definition (Social Welfare for Public Projects)

In a public project where players have valuations $v_1, \ldots, v_n$, a set S has social welfare

$$\sum_{i \in [n]} v_i(S)$$

Example

Imagine a city block with k empty storefronts. m businesses would like to open shop on this block. You wish to make the n people living near this block as happy as possible with your choices.

History

Public projects were first introduced in this form 2 years ago [PSS 08]. This paper had 2 interesting results:

- Approximation ratios better than $\sqrt{m}$ require exponential communication
- Truthful submodular auctions are NP-hard to approximate to better than $\sqrt{m}$. The techniques in this proof were a precursor to our auction result.

History

Public projects were first introduced in this form 2 years ago [PSS 08]. This paper had 2 interesting results:

- Approximation ratios better than $\sqrt{m}$ require exponential communication
- Truthful submodular auctions are NP-hard to approximate to better than $\sqrt{m}$. The techniques in this proof were a precursor to our auction result.

A matching truthful $\sqrt{m}$ mechanism for arbitrary subadditive public projects was shown later that year (M. Schapira and Y. Singer. Inapproximability of Combinatorial Public Projects. *WINE*. 2008).

Our Work

As with auctions, we focus on succinctly describable subsets of submodular valuation functions. We begin with one of the simplest valuation functions.

Our Work

As with auctions, we focus on succinctly describable subsets of submodular valuation functions. We begin with one of the simplest valuation functions.

Definition (Unit-Demand)

A valuation v is unit-demand if

$$v(S) = \max_{i \in S} v(\{i\})$$

Sharing is hard

Definition (Unit-Demand)

A valuation v is unit-demand if

$$v(S) = \max_{i \in S} v(\{i\})$$

In auctions, unit-demand is easy. Not so in public projects:

- Restrict all valuations to be 0 or 1
- For each item i, let $C(i)$ be the set of players with value 1 for i
- A player gets value 1 for set S if he is in $\bigcup_{i \in S} C(i)$
- So this is essentially max-k-cover
- max-k-cover is hard to approximate better than $1 - 1/e$

Negative Results

We were able to show NP-hardness for all classes we considered:

- unit-demand with n agents
- OXS with ≥ 3 agents
- budget-additive with ≥ 2 agents
- XOS with ≥ 2 agents
- coverage with ≥ 1 agent

For each class, we showed that no poly-time MIR mechanism can beat a $\sqrt{m}$ ratio unless $NP \subseteq P/poly$. This matches the known $\sqrt{m}$ MIR approximation.

Positive Results

If we have n unit-demand players where

- All values are either 0 or 1
- At most 2 items have value 1

The problem is still NP-hard, and no MIR mechanism can beat $\sqrt{m}$

Positive Results

If we have n unit-demand players where

- All values are either 0 or 1
- At most 2 items have value 1

The problem is still NP-hard, and no MIR mechanism can beat $\sqrt{m}$

Theorem

The below mechanism is a truthful 2-approximation.

Mechanism

- Rank each item by how many players have value 1 for it
- Choose the top k items, breaking ties by numerical order

So what?

In public projects, truthfulness becomes an issue for much simpler valuations, where it is clear that MIR mechanisms are not the end of the story.

Outline

So what?

So far, we've seen that

- MIR subadditive auctions can't beat $\min(n, \sqrt{m})$
- The same is true for many subadditive public projects
- Some public projects have better truthful mechanisms

So what?

So far, we've seen that

- MIR subadditive auctions can't beat $\min(n, \sqrt{m})$
- The same is true for many subadditive public projects
- Some public projects have better truthful mechanisms

Can the MIR auction bounds be overcome?

So what?

So far, we've seen that

- MIR subadditive auctions can't beat $\min(n, \sqrt{m})$
- The same is true for many subadditive public projects
- Some public projects have better truthful mechanisms

Can the MIR auction bounds be overcome? We don't think so.

How could we prove it?

One way to show bounds on optimal truthful mechanisms:

1. Show that the best truthful mechanisms are MIR.
2. Show bounds on efficient MIR mechanisms.

How could we prove it?

One way to show bounds on optimal truthful mechanisms:

1. Show that the best truthful mechanisms are MIR.

✓ Show bounds on efficient MIR mechanisms.

So we need only show that truthful mechanisms achieving a better ratio than $\min(n, \sqrt{m})$ are MIR.

Proving a Roberts Theorem

Theorem (Roberts 1979)

Any truthful auction mechanism for general valuations is MIR.

- Multi-unit auctions with 2 bidders can't be truthfully approximated better than 2 if all items must be allocated (R. Lavi, A. Mu'alem and N. Nisan. Towards a characterization of truthful combinatorial auctions. *FOCS*. 2003)
- A recent paper simplified the proof of Roberts' theorem to make it easier to adapt (R. Lavi, A. Mu'alem and N. Nisan. Two simplified proofs for Roberts' theorem. *Social Choice and Welfare*. 2009).
- Our auctions work was able to transform a 2-bidder result into a stronger *n*-bidder result.
- Perhaps with techniques like ours and simpler proofs of Roberts' theorem, we can characterize combinatorial auctions.

Outline

What's wrong with Truth?

- In the public projects we studied, NP-hardness implies that MIR mechanisms have bad approximation ratios

What's wrong with Truth?

- In the public projects we studied, NP-hardness implies that MIR mechanisms have bad approximation ratios
- Surprisingly, this holds for a single coverage valuation agent.

What's wrong with Truth?

- In the public projects we studied, NP-hardness implies that MIR mechanisms have bad approximation ratios
- Surprisingly, this holds for a single coverage valuation agent.
- More surprisingly, truth implies MIR for a single agent.

What's wrong with Truth?

- In the public projects we studied, NP-hardness implies that MIR mechanisms have bad approximation ratios
- Surprisingly, this holds for a single coverage valuation agent.
- More surprisingly, truth implies MIR for a single agent.
  - Agent's goal: to maximize his utility
  - Mechanism's goal: to maximize the agent's welfare
  - Best poly-time approximation: $1 - 1/e$
  - Best truthful poly-time approximation: $1/\sqrt{m}$ (**D. Buchfuhrer**, M. Schapira and Y. Singer. Computation and Incentives in Combinatorial Public Projects. *EC*. 2010)

What's wrong with Truth?

- In the public projects we studied, NP-hardness implies that MIR mechanisms have bad approximation ratios
- Surprisingly, this holds for a single coverage valuation agent.
- More surprisingly, truth implies MIR for a single agent.
  - Agent's goal: to maximize his utility
  - Mechanism's goal: to maximize the agent's welfare
  - Best poly-time approximation: $1 - 1/e$
  - Best truthful poly-time approximation: $1/\sqrt{m}$ (D. Buchfuhrer, M. Schapira and Y. Singer. Computation and Incentives in Combinatorial Public Projects. *EC*. 2010)
- Truthfulness is causing the obviously correct thing to fail.

Solutions

- In our example, the incentivized lies will increase social welfare
- So should we ditch truth in favor of some equilibrium notion?

Solutions

- In our example, the incentivized lies will increase social welfare
- So should we ditch truth in favor of some equilibrium notion?
- Truth has the following nice properties we'd like to keep:
  - Simplified design space
  - The mechanism has access to the actual valuations
  - Easy to demonstrate

Solutions

- In our example, the incentivized lies will increase social welfare
- So should we ditch truth in favor of some equilibrium notion?
- Truth has the following nice properties we'd like to keep:
  - Simplified design space
  - The mechanism has access to the actual valuations
  - Easy to demonstrate
- These properties are nice for more complicated settings.

Outline

Why rationality leads to strange results

Consider the single-player public project where the greedy algorithm is optimal.

- A rational player would choose the best outcome in the range.
- But this is NP-hard because the range is every allocation.
- The problem is that we are constraining the computation of the mechanism, but not the players

First attempt

What if we limit the player to polynomial computation?

First attempt

What if we limit the player to polynomial computation?

There are 2 problems with this:

- A player can have a specialized algorithm that computes the best allocation for his valuation
- A player could use heuristics to improve some allocations

Assuming the problem away

In order to get something like rationality, we need per instance optimality.

Assuming the problem away

In order to get something like rationality, we need per instance optimality.

What if we assume that the player can't find a better allocation regarding the social welfare?

Universal VCG Mechanisms

If we assume that the players can't find an assignment with better social welfare, VCG payments make any algorithm truthful.

Universal VCG Mechanisms

If we assume that the players can't find an assignment with better social welfare, VCG payments make any algorithm truthful.

Possible conclusions:

- This idea is stupid because it assumes away the real problems of mechanism design

Universal VCG Mechanisms

If we assume that the players can't find an assignment with better social welfare, VCG payments make any algorithm truthful.

Possible conclusions:

- This idea is stupid because it assumes away the real problems of mechanism design
- This idea is stupid because it pays out $n - 1$ times the social welfare.

Universal VCG Mechanisms

If we assume that the players can't find an assignment with better social welfare, VCG payments make any algorithm truthful.

Possible conclusions:

- ~~This idea is stupid because it assumes away the real problems of mechanism design~~
- This idea is stupid because it pays out $n - 1$ times the social welfare.

Universal VCG Mechanisms

If we assume that the players can't find an assignment with better social welfare, VCG payments make any algorithm truthful.

Possible conclusions:

- This idea is stupid because it pays out $n - 1$ times the social welfare.

Universal VCG Mechanisms

If we assume that the players can't find an assignment with better social welfare, VCG payments make any algorithm truthful.

Possible conclusions:

- This idea is stupid because it pays out $n - 1$ times the social welfare.

Solution: Find a good pivot rule.

Pivot Rule

Clarke Pivot

Player i is charged $v_{-i}(A(v_{-i}, 0)) - v_{-i}(A(v_{-i}, v_i))$

The idea is to offset the large payment to the players by an amount that does not depend on the player's bid.

- Individual rationality (non-negative utility for each player)
- No payments (or only small payments) made to the players

Pivot Rule

Clarke Pivot

Player i is charged $v_{-i}(A(v_{-i}, 0)) - v_{-i}(A(v_{-i}, v_i))$

The idea is to offset the large payment to the players by an amount that does not depend on the player's bid.

- Individual rationality (non-negative utility for each player)
- No payments (or only small payments) made to the players

We can get both using the above Clarke pivot if

- $\min_{v_i} v(A(v_{-i}, v_i)) = v(A(v_{-i}, 0))$
- $\max_{v_i} v_{-i}(A(v_{-i}, v_i)) = v_{-i}(A(v_{-i}, 0))$

Previous Work

This payment scheme was studied previously (N. Nisan and A. Ronen. Computationally Feasible VCG Mechanisms. *Journal of Artificial Intelligence Research*. 2007), but they only focus on individual rationality and a "second-chance" idea:

Previous Work

This payment scheme was studied previously (N. Nisan and A. Ronen. Computationally Feasible VCG Mechanisms. *Journal of Artificial Intelligence Research*. 2007), but they only focus on individual rationality and a "second-chance" idea:

Mechanism (Second-Chance)

1. Take bids in the form of each players' valuation function

Previous Work

This payment scheme was studied previously (N. Nisan and A. Ronen. Computationally Feasible VCG Mechanisms. *Journal of Artificial Intelligence Research*. 2007), but they only focus on individual rationality and a “second-chance” idea:

Mechanism (Second-Chance)

1. Take bids in the form of each players' valuation function
2. Compute an allocation

Previous Work

This payment scheme was studied previously (N. Nisan and A. Ronen. Computationally Feasible VCG Mechanisms. *Journal of Artificial Intelligence Research*. 2007), but they only focus on individual rationality and a "second-chance" idea:

Mechanism (Second-Chance)

1. Take bids in the form of each players' valuation function
2. Compute an allocation
3. Take bids describing other possible allocations

Previous Work

This payment scheme was studied previously (N. Nisan and A. Ronen. Computationally Feasible VCG Mechanisms. *Journal of Artificial Intelligence Research*. 2007), but they only focus on individual rationality and a "second-chance" idea:

Mechanism (Second-Chance)

1. Take bids in the form of each players' valuation function
2. Compute an allocation
3. Take bids describing other possible allocations
4. Use the allocation maximizing social welfare

Previous Work

This payment scheme was studied previously (N. Nisan and A. Ronen. Computationally Feasible VCG Mechanisms. *Journal of Artificial Intelligence Research*. 2007), but they only focus on individual rationality and a "second-chance" idea:

Mechanism (Second-Chance)

1. Take bids in the form of each players' valuation function
2. Compute an allocation
3. Take bids describing other possible allocations
4. Use the allocation maximizing social welfare
5. Charge VCG payments

What will we add?

We intend to improve on the Nisan/Ronen work by

- Looking more at constraining the payments made by the mechanism
- Showing auctions in which these mechanisms help

What will we add?

We intend to improve on the Nisan/Ronen work by

- Looking more at constraining the payments made by the mechanism
- Showing auctions in which these mechanisms help

Example

Any allocation algorithm for two players can be made truthful in this model. Simply output the best allocation from the one output by the algorithm, and the 2 in which one player gets all the items.

So what?

This mechanism is great for 2-player budget-additive auctions. The best known truthful mechanism gets a ratio of $\min(n, \sqrt{m})$, but now we can use a known FPTAS.

Outline

Another Solution

Rather than assume limited players, what if we assume players have special information?

- Problems occur when it's hard to determine what the player wants
- If we assume the player knows better, we should take advantage

Another Solution

Rather than assume limited players, what if we assume players have special information?

- Problems occur when it's hard to determine what the player wants
- If we assume the player knows better, we should take advantage
- Idea from communication complexity: demand queries

Demand Queries

Demand Query

Input: A price p_i for each item i
Output: A set S maximizing $v(S) - \sum_i p_i$

Demand queries can be used to solve auctions via linear programming in some situations (S. Bikhchandani and J.W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*. 1997)

Demand Queries

Demand Query

Input: A price p_i for each item i
Output: A set S maximizing $v(S) - \sum_i p_i$

Demand queries can be used to solve auctions via linear programming in some situations (S. Bikhchandani and J.W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*. 1997)

k-Demand Queries

Input: A price p_i for each item i and a positive integer $k \leq m$.
Output: A set S, $|S| \leq k$ maximizing $v(S) - \sum_{i \in S} p_i$

Questions

While demand queries have been studied extensively from a communication complexity viewpoint, succinctly described valuations can just be directly revealed, so demand queries have not received much attention.

- Can demand queries solve our single-player public project?
- Can k-demand queries solve more complicated public projects?
- Can demand queries solve hard auctions?

Demand queries and single player public projects

Theorem

There exists a class of public projects with description length $O(m)$ which requires exponentially many demand queries to solve.

Demand queries and single player public projects

Theorem

There exists a class of public projects with description length $O(m)$ which requires exponentially many demand queries to solve.

Proof.

Let $|S^*| = k$. Define

$$v_{S^*}(S) = \begin{cases} 0 & |S| \leq k, S \neq S^* \\ 1 & S = S^* \\ 3 & |S| > k \end{cases}$$

Demand queries and single player public projects

Theorem

There exists a class of public projects with description length $O(m)$ which requires exponentially many demand queries to solve.

Proof.

Let $|S^*| = k$. Define

$$v_{S^*}(S) = \begin{cases} 0 & |S| \leq k, S \neq S^* \\ 1 & S = S^* \\ 3 & |S| > k \end{cases}$$

If the demand query with prices $p_1, \ldots, p_m$ returns S^*, $p_i \geq 2$ for $i \notin S^*$ and $p_i \leq 1$ for $i \in S^*$. □

Solving auctions with demand queries

Consider a 2-player auction with 1 additive player and 1 budget-additive player.

Additive Player

Values $u_1, \ldots, u_m$ and valuation function $U(S) = \sum_{i \in S} u_i$

Budget-Additive Player

Values $v_1, \ldots, v_m, B$ and valuation $V(S) = \min(\sum_{i \in S} v_i, B)$

Solving auctions with demand queries

Additive Player

Values $u_1, \ldots, u_m$ and valuation function $U(S) = \sum_{i \in S} u_i$

Budget-Additive Player

Values $v_1, \ldots, v_m, B$ and valuation $V(S) = \min(\sum_{i \in S} v_i, B)$

- Query the budget-additive player with $p_i = u_i$
- He returns a set S maximizing $V(S) - \sum_{i \in S} u_i$
- This also maximizes $V(S) + U(S^C)$
- We can similarly solve the same public project using negative prices

Going further

Can demand queries be used to solve

- auctions and public projects with 2 budget-additive players?
- the single player coverage valuation public project?
- other hard auctions and public projects?

Outline

Revenue Maximization

So far, we've worried about maximizing the *social welfare*.

Less altruistic auctioneers are more concerned about revenue.

Revenue Maximization

So far, we've worried about maximizing the *social welfare*.

Less altruistic auctioneers are more concerned about revenue.

- Consider a 1 item, 2-bidder auction
- Bidder 1 values the item at $1, bidder 2 at $1,000

Revenue Maximization

So far, we've worried about maximizing the *social welfare*.

Less altruistic auctioneers are more concerned about revenue.

- Consider a 1 item, 2-bidder auction
- Bidder 1 values the item at $1, bidder 2 at $1,000
- The standard Vickrey auction will sell the item for $1

Revenue Maximization

So far, we've worried about maximizing the *social welfare*.
Less altruistic auctioneers are more concerned about revenue.

- Consider a 1 item, 2-bidder auction
- Bidder 1 values the item at $1, bidder 2 at $1,000
- The standard Vickrey auction will sell the item for $1
- A smart auctioneer would put a reserve price close to $1,000

Revenue Maximization

So far, we've worried about maximizing the *social welfare*.
Less altruistic auctioneers are more concerned about revenue.

- Consider a 1 item, 2-bidder auction
- Bidder 1 values the item at $1, bidder 2 at $1,000
- The standard Vickrey auction will sell the item for $1
- A smart auctioneer would put a reserve price close to $1,000
- The auctioneer loses if the actual value is $900 or $10,000

Revenue Maximization

So far, we've worried about maximizing the *social welfare*.
Less altruistic auctioneers are more concerned about revenue.

- Consider a 1 item, 2-bidder auction
- Bidder 1 values the item at $1, bidder 2 at $1,000
- The standard Vickrey auction will sell the item for $1
- A smart auctioneer would put a reserve price close to $1,000
- The auctioneer loses if the actual value is $900 or $10,000
- Revenue maximization depends on prior knowledge

Bayesian Auctions

We add prior distributions on the players' valuations to the auction.

$$v_i \sim D_i$$

and the goal is to design a truthful mechanism maximizing $E[\sum_i p_i]$

Single Item Example

Example (Algorithmic Game Theory, Chapter 13)

Consider the 2-bidder 1-item auction in which $v_i \sim \mathit{uniform}([0, 1])$.
Let the max and min prices be $v_{\max}$ and $v_{\min}$.

- If $v_{\max} < 1/2$, don't allocate
- Otherwise, allocate to the higher bidder for $\max(1/2, v_{\min})$

Single Item Example

Example (Algorithmic Game Theory, Chapter 13)

Consider the 2-bidder 1-item auction in which $v_i \sim uniform([0,1])$.
Let the max and min prices be $v_{\max}$ and $v_{\min}$.

- If $v_{\max} < 1/2$, don't allocate
- Otherwise, allocate to the higher bidder for $\max(1/2, v_{\min})$

- This is truthful because it is a Vickrey auction

Single Item Example

Example (Algorithmic Game Theory, Chapter 13)

Consider the 2-bidder 1-item auction in which $v_i \sim uniform([0,1])$.
Let the max and min prices be $v_{\max}$ and $v_{\min}$.

- If $v_{\max} < 1/2$, don't allocate
- Otherwise, allocate to the higher bidder for $\max(1/2, v_{\min})$

- This is truthful because it is a Vickrey auction
- The expected revenue is $5/12$

Single Item Example

Example (Algorithmic Game Theory, Chapter 13)

Consider the 2-bidder 1-item auction in which $v_i \sim uniform([0,1])$.
Let the max and min prices be $v_{\max}$ and $v_{\min}$.

- If $v_{\max} < 1/2$, don't allocate
- Otherwise, allocate to the higher bidder for $\max(1/2, v_{\min})$

- This is truthful because it is a Vickrey auction
- The expected revenue is $5/12$
- The standard Vickrey auction has expected revenue $4/12$

Single Item Example

Example (Algorithmic Game Theory, Chapter 13)

Consider the 2-bidder 1-item auction in which $v_i \sim \textit{uniform}([0,1])$.
Let the max and min prices be $v_{\max}$ and $v_{\min}$.

- If $v_{\max} < 1/2$, don't allocate
- Otherwise, allocate to the higher bidder for $\max(1/2, v_{\min})$

- This is truthful because it is a Vickrey auction
- The expected revenue is $5/12$
- The standard Vickrey auction has expected revenue $4/12$
- For more complicated distributions, the expected revenue is maximized via a virtual auction (R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 1981)

The Problem

While mechanism design is solved for a single item, it remains largely open for multiple items

The Problem

While mechanism design is solved for a single item, it remains largely open for multiple items

Example (Unsolved Revenue Maximization Problem)

- 2 bidders
- 2 items
- Each bidder wants at most 1 item
- Item values are drawn from a known distribution

The Problem

While mechanism design is solved for a single item, it remains largely open for multiple items

Example (Unsolved Revenue Maximization Problem)

- 2 bidders
- 2 items
- Each bidder wants at most 1 item
- Item values are drawn from a known distribution

Maybe revenue maximization is computationally hard?

Price menus

The Problem

2 bidders, 2 items, allocate at most 1 item to each bidder to maximize revenue with known priors

All truthful mechanisms have the following form:

1. Players reveal their values
2. Based on each player's value, the mechanism determines item prices for the other player
3. Each player is given at most 1 item in order to maximize his utility under the price regime
4. Players are charged according to the prices from step 2

Price menu functions

So we want two price menu functions P_1, P_2 such that:

- For every pair of values $(v_{11}, v_{12}), (v_{21}, v_{22})$ there is an allocation such that:
  - Each player's utility is maximized
  - Both players don't get the same item
- Revenue is maximized

Price menu functions

So we want two price menu functions P_1, P_2 such that:

- For every pair of values $(v_{11}, v_{12}), (v_{21}, v_{22})$ there is an allocation such that:
  - Each player's utility is maximized
  - Both players don't get the same item
- Revenue is maximized

Note

Given P_1, we can find the function P_2 maximizing the revenue for any discrete distribution. For each pair (v_{21}, v_{22}), try all relevant prices and calculate the expected revenue, then choose the max.

Nash?

Note

Given P_1, we can find a best P_2.

So we have a table of numbers, and want to find a pair of strategies P_1, P_2 based on these numbers that are best responses to each other.

Nash?

Note

Given P_1, we can find a best P_2.

So we have a table of numbers, and want to find a pair of strategies P_1, P_2 based on these numbers that are best responses to each other.

Conjecture

Finding a pair P_1, P_2 in equilibrium is PPAD-hard.

No PPAD Hardness

Lemma

A pair of price menu functions P_1, P_2 which are best responses to each other can be found in polynomial time.

Proof.

- Let $P^*(u, v) = (\infty, \infty)$
- Let P_2 be a best response to P^*
- Let P_1 be a best response to P_2

□

No PPAD Hardness

Lemma

A pair of price menu functions P_1, P_2 which are best responses to each other can be found in polynomial time.

Proof.

- Let $P^*(u, v) = (\infty, \infty)$
- Let P_2 be a best response to P^*
- Let P_1 be a best response to P_2
- P_2 gets maximum possible revenue from player 1, so it is a best response to anything, including P_1 □

No PPAD Hardness

Lemma

A pair of price menu functions P_1, P_2 which are best responses to each other can be found in polynomial time.

Proof.

- Let $P^*(u, v) = (\infty, \infty)$
- Let P_2 be a best response to P^*
- Let P_1 be a best response to P_2
- P_2 gets maximum possible revenue from player 1, so it is a best response to anything, including P_1 □

So we'll have to try for NP-hardness

Outline

Summary

Results so far

- A $\min(n, \sqrt{m})$ bound on maximal-in-range submodular auctions
- Several results relating to public projects, including a troubling bound of $\sqrt{m}$ for a truthful 1-player public project

Hopeful future results

- A $\min(n, \sqrt{m})$ bound on truthful submodular public projects
- Player knowledge and/or bounded rationality can be used to circumvent issues with truthfulness
- Revenue maximization is NP-hard even for a constant number of bidders and items and simple valuations