

A Duality-Based Unified Approach to Bayesian Mechanism Design

YANG CAI

McGill University

and

NIKHIL R. DEVANUR

Microsoft Research

and

S. MATTHEW WEINBERG

Princeton University

In this letter we briefly survey our recent work [Cai et al. 2016]. In it, we provide a new duality theory for Bayesian mechanism design which is quite general, and applies for any objective the designer wishes to optimize, and for arbitrary agent valuations. We then apply our theory to auction design settings with many independent buyers who have independent values for many items, and are able to provide a unified proof of several recent exciting works on this front [Hart and Nisan 2012; Li and Yao 2013; Babaioff et al. 2014; Yao 2015; Chawla et al. 2007; Chawla et al. 2010; Chawla et al. 2015]. These works all show that simple mechanisms are approximately optimal in various settings. In some cases, our principled approach yields greatly improved approximation ratios as well.

Categories and Subject Descriptors: J.4 [Social and Behavioral Science]: Economics

General Terms: Algorithms, Design, Economics, Theory

Additional Key Words and Phrases: Bayesian Mechanism Design, Simple Mechanisms, Revenue, Approximation, Duality

1. INTRODUCTION

In this letter, we briefly overview our recent paper [Cai et al. 2016]. In it, we present a new duality theory for Bayesian mechanism design that applies to settings with multiple bidders with arbitrary valuation functions, and where the designer's objective may be some arbitrary function of the bidder types and outcome selected. For the sake of exposition, in this letter we focus on a special case that has received much attention recently: there are n bidders and m items, bidder i 's value for item j (t_{ij}) is drawn independently from D_{ij} , and the designer wishes to find the revenue-optimal Bayesian Incentive Compatible auction. Such bidders might be *unit-demand* ($t_i(S) = \max_{j \in S} \{t_{ij}\}$), *additive* ($t_i(S) = \sum_{j \in S} \{t_{ij}\}$), or more generally \mathcal{I} -additive ($t_i(S) = \max_{S' \subseteq S, S' \in \mathcal{I}} \{\sum_{j \in S'} t_{ij}\}$, for some $\mathcal{I} \subseteq 2^{[m]}$). We use the term *matroid-additive* if \mathcal{I} is a matroid, and *downwards-closed-additive* if \mathcal{I} is downwards-closed. Recent work further generalizes this valuation space to arbitrary XOS or subadditive functions with independent items. We omit formal definitions of these in this note, referring the reader to [Rubinstein and Weinberg

Authors' addresses: cai@cs.mcgill.ca, nikdev@microsoft.com, smweinberg@princeton.edu

2015].

Even in the simplest of these settings (say, one additive or unit-demand bidder and two items), the optimal mechanism is known to be prohibitively complex (e.g. [Hart and Reny 2012]), so much recent attention has shifted instead to designing simple mechanisms that are approximately optimal. Starting from seminal works of Chawla et al. [Chawla et al. 2007] for unit-demand bidders, and Hart and Nisan [Hart and Nisan 2012] for additive bidders, we now know that simple mechanisms are approximately optimal in quite general settings. Proofs of these results were already quite principled, however the techniques developed were largely disjoint and there seem to be significant technical barriers to extending them beyond unit-demand or additive bidders in multi-bidder settings. In our paper, we apply our new duality theory to these problems and derive a unified proof for many works in this series. These results are important for two reasons: first, our duality framework provides more explicit bounds than previous techniques, which allows for improved approximation ratios and extensions to broader settings. Second, because our proofs all follow the same framework, we gain a considerable amount of previously undiscovered intuition about these problems.

We summarize these results below. The first four bullets are proved in [Cai et al. 2016]. The last three are proved in a manuscript by one of the authors and his collaborator [Cai and Zhao 2016]. Informal definitions of the simple mechanisms used are included in footnotes below, and we refer the reader to our paper [Cai et al. 2016] for more formal definitions.

THEOREM 1. *There is a canonical dual solution inducing the following results:*

In [Cai et al. 2016],

- $\text{SREV} \geq \text{OPT}/4$, when there is a single unit-demand buyer (matches [Chawla et al. 2010]+ [Chawla et al. 2015]).¹
- $\max\{\text{SREV}, \text{BREV}\} \geq \text{OPT}/6$, when there is a single additive buyer (matches [Babaioff et al. 2014]).²
- $\text{POSTREV} \geq \text{OPT}/24$, when there are many unit-demand buyers (improved from 33.75 [Chawla et al. 2010]+ [Chawla et al. 2015]).³
- $\max\{\text{SREV}, \text{BVCG}\} \geq \text{OPT}/8$, when there are many additive buyers (improved from 69 [Yao 2015]).⁴

In [Cai and Zhao 2016] (results below use the definition of “subadditive/XOS over independent items” posed in [Rubinstein and Weinberg 2015]),

- $\max\{\text{SREV}, \text{BREV}\} \geq \text{OPT}/64$, when there is a single subadditive buyer (improved from 338 [Rubinstein and Weinberg 2015]).

¹SREV denotes the optimal revenue attainable by posting a price on each item separately.

²BREV denotes the optimal revenue attainable by posting a price only on the grand bundle of all items.

³POSTREV denotes the optimal revenue attainable by a posted-price mechanism.

⁴For multiple bidders, SREV denotes the optimal revenue attainable by running Myerson’s auction separately for each item. BVCG denotes the revenue attainable by running the VCG mechanism with an additional per-bidder entry fee.

- BPOSTREV $\geq OPT/70$, when there are many downwards-closed-additive buyers. (improved from 133 just for matroid-additive buyers [Chawla and Miller 2016]).⁵
- BPOSTREV $\geq OPT/268$, when there are many XOS buyers.

In the remainder of this letter we give a high level overview as to how our duality theory is derived and applied to this setting. We also note that other duality theories for Bayesian mechanism design have been developed in recent works of [Giannakopoulos and Koutsoupias 2014; 2015; Daskalakis et al. 2014; 2015; Haghpanah and Hartline 2015]. We refer the reader to our paper for a more detailed comparison of our work to these, but note here that ours is the first to extend to multiple bidders and valuation functions beyond additive or unit-demand, and also the first to be successfully applied to obtain approximation results, rather than a characterization of the exact optimal mechanism in simple settings.

2. A LINEAR PROGRAM AND OUR DUALITY THEORY

Similar to a recent series of works that develops black-box reductions for Bayesian mechanism design [Cai et al. 2012a; 2012b; 2013a; 2013b], we begin with an LP formulation. To write the LP, we first introduce the concept of a *reduced form*. For a mechanism M , the reduced form stores variables $\pi_{ij}^M(t)$, which is the probability that bidder i receives item j when reporting type t to M (over any randomness in M , as well as in D_{-i}), as well as $p_i^M(t)$, which is the expected price bidder i pays when reporting type t to M (over the same randomness). We can also go in the other direction, and say that a reduced form π is *feasible* if there exists some mechanism M (that allocates each item at most once on every profile⁶) such that $\pi = \pi^M$. The space of feasible reduced forms is convex, and depends on D , so we denote it by $P(D)$. The LP simply optimizes expected revenue over all truthful and feasible reduced forms.

To proceed, we'll introduce a variable $\lambda_i(t, t')$ for each of the Bayesian Incentive Compatible (BIC) constraints, and take the partial Lagrangian of LP 1 by Lagrangifying all BIC constraints. Every set of Lagrangian multipliers/dual variables induces a (possibly infinite) upper bound on the optimal solution to LP 1. Of course, only finite upper bounds are useful. The first step is to characterize exactly which dual solutions induce a finite upper bound (we call them *useful*):

LEMMA 1 USEFUL DUAL VARIABLES. *A set of dual variables is useful iff for each bidder i , λ_i forms a valid flow, i.e., iff the following satisfies weak flow conservation (flow in \geq flow out) at all nodes except the source:*

- Nodes: A super source s , along with a node t for every type $t \in T$.
- $f_i(t)$ flow from s to t , for all $t \in T$.
- $\lambda_i(t, t')$ flow from t to t' for all $t, t' \in T$.

⁵BPOSTREV denotes the optimal revenue attainable by a posted-price mechanism with an additional per-bidder entry fee.

⁶This is for the case of additive bidders. In the case of unit-demand bidders, M must also allocate each bidder at most one item. In the case of \mathcal{I} -additive bidders, M must also allocate each bidder a set in \mathcal{I} . A generalization of the reduced form is required to handle valuation functions beyond \mathcal{I} -additive, see [Cai et al. 2016].

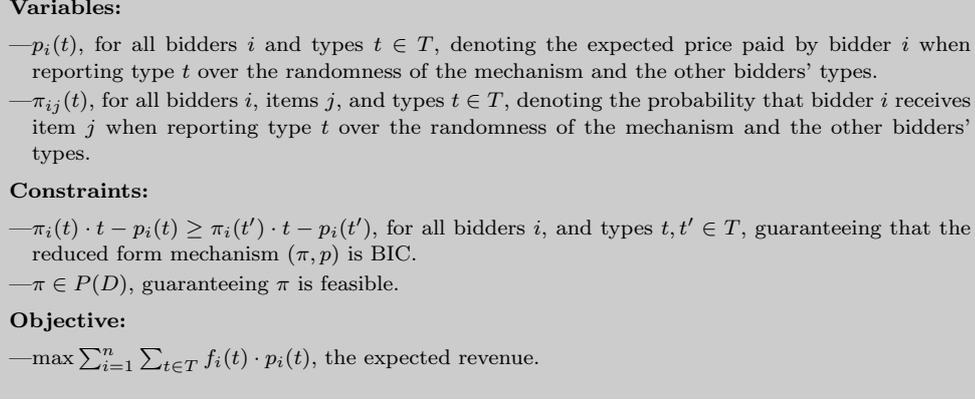


Fig. 1. A Linear Program (LP) for Revenue Optimization.

Lemma 1 tightly characterizes useful dual solutions as those that induce a valid flow in a graph where the nodes correspond to possible bidder types. The way to interpret these dual solutions is through complementary slackness: t sending flow to t' in λ_i corresponds to the incentive constraint between t and t' binding for bidder i . For every useful dual solution, we can also define a corresponding virtual value function. The definition might seem arbitrary until the theorem stated afterwards.

DEFINITION 1 VIRTUAL VALUE FUNCTION. *For each λ , we define a corresponding virtual value function $\Phi^\lambda(\cdot)$, such that for every bidder i , every type $t \in T$, $\Phi_i^\lambda(t) = t - \frac{1}{f_i(t)} \sum_{t' \in T} \lambda_i(t', t)(t' - t)$. Note that $\Phi_i^\lambda(t)$ is an m -dimensional vector (like all $t \in T$).*

THEOREM 2 VIRTUAL WELFARE \geq REVENUE. *For any set of useful duals λ and any BIC mechanism $M = (\pi, p)$, the revenue of M is no larger than the virtual welfare of π w.r.t. $\Phi^\lambda(\cdot)$.*

$$\sum_{i=1}^n \sum_{t \in T} f_i(t) \cdot p_i(t) \leq \sum_{i=1}^n \sum_{t \in T} \sum_{j=1}^m f_i(t) \cdot \pi_{ij}(t) \cdot \Phi_{ij}^\lambda(t)$$

Let λ^ be the optimal dual variables and $M^* = (\pi^*, p^*)$ be the revenue optimal BIC mechanism, then the expected virtual welfare with respect to Φ^{λ^*} under π^* equals the expected revenue of M^* .*

Theorem 2 provides our duality theory. Essentially, we are claiming that every useful dual solution can be interpreted as a virtual valuation function, which upper bounds the optimal attainable revenue. So the dual problem to revenue maximization is a search for virtual valuation functions. In the next section, we show that there is a canonical way of setting the dual variables such that the corresponding optimal virtual welfare attainable by any feasible allocation (not necessarily BIC) is a good upper bound - this is important as analyzing the space of feasible allocations is *drastically* less complex than the space of BIC mechanisms.

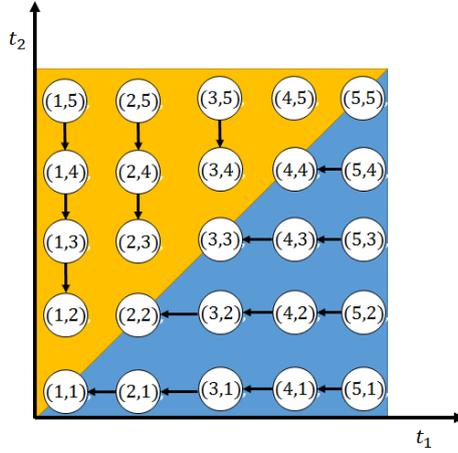


Fig. 2. An example of λ for one bidder and two items.

3. A CANONICAL DUAL AND DECOMPOSITION

Here, we briefly show two examples of how to use our duality theory. Let's begin with a somewhat trivial upper bound. Observe that setting $\lambda_i(t, t') = 0$ for all i, t, t' results in a useful dual solution. Plugging this in to Definition 1, we get $\Phi_i^\lambda(t) = t$. So by Theorem 2, we get that the revenue of any BIC mechanism is at most its welfare. Of course, this is a trivial upper bound and we did not need a duality theory to prove it, but perhaps it is an illustrative example.

Let's move on now to a canonical dual solution for a single bidder that induces the upper bound used in the first two bullets of Theorem 1.⁷ For simplicity, we assume the type space is a m -dimensional bounded integer lattice $\times_{j \in [m]} [H_j]$. It will soon be clear to the readers that the same flow can be easily extended to the general case. Consider the following implicitly defined flow: for all valuation vectors t , let $j^*(t) = \operatorname{argmax}_j \{t_j\}$ (break ties lexicographically). Let $t' = (t_{j^*(t)} + 1; t_{i, -j^*(t)})$, i.e. the type that takes t 's favorite item and increases the value by 1. Then t' sends all of its incoming flow into t and it is the only type that sends flow into t . Note that t' might not exist for all t , in which case these types receive no incoming flow except from the super source s . Note that every type sends outgoing flow to at most one other type, but again that not all types send outgoing flow. Figure 2 shows a diagram of this flow. So what Φ^λ does this flow induce? It is not immediate to see, but some inductive calculation yields that: $\Phi_j^\lambda(t) = t_j$, for all $j \neq j^*(t)$, and $\Phi_{j^*(t)}^\lambda(t) = \varphi_{j^*(t)}(t_{j^*(t)})$, where $\varphi_j(t_j) = t_j - \frac{1 - F_j(t_j)}{f_j(t_j)}$, Myerson's single-dimensional virtual value for D_j [Myerson 1981]. This gives us a stronger upper bound on the optimal attainable revenue by any BIC mechanism: it is no more than the welfare of all *non-favorite* items, plus the revenue of a related single-dimensional "copies" setting first studied by [Chawla et al. 2007]. In the paper, we show how to get a

⁷Technically, the presented dual is only for the case that each D_j is regular. A small modification is necessary for irregular distributions.

constant factor of this upper bound using the simple mechanisms prescribed above.

Here is some further intuition behind this bound: the optimal attainable welfare is a trivial upper bound, and can be observed simply by ignoring all BIC constraints except for individual rationality. Our canonical bound is also a relaxation: we are ignoring all BIC constraints *except* for where the bidder considers minimally underreporting their value for their favorite item. This turns out to be a surprisingly useful relaxation, which induces all of the bounds in Theorem 1 after going through our duality theory.

REFERENCES

- BABAIOFF, M., IMMORLICA, N., LUCIER, B., AND WEINBERG, S. M. 2014. A Simple and Approximately Optimal Mechanism for an Additive Buyer. In *the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- CAI, Y., DASKALAKIS, C., AND WEINBERG, S. M. 2012a. An Algorithmic Characterization of Multi-Dimensional Mechanisms. In *the 44th Annual ACM Symposium on Theory of Computing (STOC)*.
- CAI, Y., DASKALAKIS, C., AND WEINBERG, S. M. 2012b. Optimal Multi-Dimensional Mechanism Design: Reducing Revenue to Welfare Maximization. In *the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- CAI, Y., DASKALAKIS, C., AND WEINBERG, S. M. 2013a. Reducing Revenue to Welfare Maximization : Approximation Algorithms and other Generalizations. In *the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- CAI, Y., DASKALAKIS, C., AND WEINBERG, S. M. 2013b. Understanding Incentives: Mechanism Design becomes Algorithm Design. In *the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- CAI, Y., DEVANUR, N. R., AND WEINBERG, S. M. 2016. A duality based unified approach to bayesian mechanism design. In *the 48th Annual ACM Symposium on Theory of Computing (STOC)*.
- CAI, Y. AND ZHAO, M. 2016. Simple Mechanisms for Subadditive Buyers via Duality. *Manuscript*.
- CHAWLA, S., HARTLINE, J. D., AND KLEINBERG, R. D. 2007. Algorithmic Pricing via Virtual Valuations. In *the 8th ACM Conference on Electronic Commerce (EC)*.
- CHAWLA, S., HARTLINE, J. D., MALEC, D. L., AND SIVAN, B. 2010. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *the 42nd ACM Symposium on Theory of Computing (STOC)*.
- CHAWLA, S., MALEC, D. L., AND SIVAN, B. 2015. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior* 91, 297–317.
- CHAWLA, S. AND MILLER, J. B. 2016. Mechanism design for subadditive agents via an ex-ante relaxation. *CoRR abs/1603.03806*.
- DASKALAKIS, C., DECKELBAUM, A., AND TZAMOS, C. 2014. The complexity of optimal mechanism design. In *the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- DASKALAKIS, C., DECKELBAUM, A., AND TZAMOS, C. 2015. Strong duality for a multiple-good monopolist. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*. 449–450.
- GIANNAKOPOULOS, Y. AND KOUTSOPIAS, E. 2014. Duality and optimality of auctions for uniform distributions. In *ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014*. 259–276.
- GIANNAKOPOULOS, Y. AND KOUTSOPIAS, E. 2015. Selling two goods optimally. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part II*. 650–662.
- HAGHPANAH, N. AND HARTLINE, J. 2015. Reverse mechanism design. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*.

- HART, S. AND NISAN, N. 2012. Approximate Revenue Maximization with Multiple Items. In *the 13th ACM Conference on Electronic Commerce (EC)*.
- HART, S. AND RENY, P. J. 2012. Maximal revenue with multiple goods: Nonmonotonicity and other observations. *Discussion Paper Series dp630, The Center for the Study of Rationality, Hebrew University, Jerusalem*.
- LI, X. AND YAO, A. C.-C. 2013. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences* 110, 28, 11232–11237.
- MYERSON, R. B. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1, 58–73.
- RUBINSTEIN, A. AND WEINBERG, S. M. 2015. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*. 377–394.
- YAO, A. C. 2015. An n -to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*. 92–109.