2018 IEEE 59th Annual Symposium on Foundations of Computer Science

An End-to-end Argument in Mechanism Design (Prior-independent Auctions for Budgeted Agents)[†]

Yiding Feng
EECS Dept.
Northwestern University
Evanston, IL
yidingfeng2021@u.northwestern.edu

Jason D. Hartline
EECS Dept.
Northwestern University
Evanston, IL
hartline@northwestern.edu

***Abstract*—This paper considers prior-independent mechanism design, namely identifying a single mechanism that has near optimal performance on every prior distribution. We show that mechanisms with truthtelling equilibria, a.k.a., revelation mechanisms, do not always give optimal prior-independent mechanisms and we define the revelation gap to quantify the non-optimality of revelation mechanisms. This study suggests that it is important to develop a theory for the design of non-revelation mechanisms.**

Our analysis focuses on welfare maximization by single-item auctions for agents with budgets and a natural regularity assumption on their distribution of values. The all-pay auction (a non-revelation mechanism) is the Bayesian optimal mechanism; as it is prior-independent it is also the prior-independent optimal mechanism (a 1-approximation). We prove a lower bound on the prior-independent approximation of revelation mechanisms of 1.013 and that the clinching auction (a revelation mechanism) is a prior-independent $e \approx 2.714$ approximation. Thus the revelation gap for single-item welfare maximization with public budget agents is in $[1.013, e]$.

Some of our analyses extend to the revenue objective, position environments, and irregular distributions.

***Keywords*-Mechanism design, Approximation, Budgets**

I. Introduction

The *end-to-end principle* in distributed systems advocates environment-independent protocols (for the center) that push environment-specific complexity to the applications (the end points) that use the protocol (Saltzer et al., 1984). This principle enabled the Internet protocols designed for the workloads of the 1980s to continue to succeed with workloads of the 2010s. On the other hand, research in mechanism design (which governs the design of protocols for strategic agents and has application both in computer science and economics) almost exclusively adheres to the *revelation principle* which suggests the design of mechanisms where each agent's best strategy is to truthfully report her preferences. In revelation mechanisms the agents (the end points) have very a simple "report your true preference" strategies and the mechanism (the center) has the complex task of finding an outcome that both enforces this truthfulness property and also obtains a desirable outcome. Unsurprisingly, optimal revelation mechanisms tend to be complex and dependent on the environment. This paper demonstrates that the end-to-end argument has bite in mechanism design by showing that non-revelation mechanisms are strictly better than revelation mechanisms for a canonical mechanism design problem.

In prior-independent mechanism design, it is assumed that the agents' preferences are drawn from a distribution that is not known to the designer. The goal of prior-independent mechanism design is to identify mechanisms that are good approximations to the optimal mechanism that is tailored to the distribution of preferences. Specifically, a mechanism is sought to minimize the ratio of its expected performance to the expected performance of the optimal mechanism in worst case over distributions from which the preferences of the agents are drawn. This notion is a standard one that has been applied to revenue maximization (Dhangwatnotai et al., 2015; Fu et al., 2015; Allouah and Besbes, 2018), multi-dimensional mechanism design (Devanur et al., 2011; Roughgarden et al., 2012), makespan minimization (Chawla et al., 2013), mechanism design for risk-averse agents (Fu et al., 2013), and mechanism design for agents with interdependent values (Chawla et al., 2014a). In none of these scenarios is the optimal prior-independent mechanism known; cf. Fu et al. (2015) and Allouah and Besbes (2018).

The revelation principle suggests that if there is a mechanism with a good equilibrium outcome, there is a mechanism where truthtelling achieves the same outcome in a truthtelling equilibrium. Due to the rev-

[†]This work was supported by NSF CCF-1618502. The full version of this paper is available at https://arxiv.org/abs/1804.01977.

2575-8454/18/$31.00 ©2018 IEEE
DOI 10.1109/FOCS.2018.00046

elation principle, much of the theory of mechanism design is developed for revelation mechanisms, i.e., ones where truthtelling is an equilibrium. The proof of the revelation principle is simple: A revelation mechanism can simulate the equilibrium strategies in the non-revelation mechanism to obtain the same outcome as a truthtelling equilibrium, i.e., agents report true preferences to the revelation mechanism, it simulates the agent strategies in the non-revelation mechanism, and it outputs the outcome of the simulation. For Bayesian non-revelation mechanisms (where the agents' preferences are drawn from a prior distribution), the agents' equilibrium strategies are a function of the prior and thus the corresponding revelation mechanism constructed via the revelation principle is not prior-independent. Thus, the restriction to revelation mechanisms is not generally without loss for prior-independent mechanism design. Non-revelation mechanisms, on the other hand, are widely used in practice and frequently have easily to identify and natural equilibria (e.g., in rank-based auctions, see Chawla and Hartline, 2013). Our proof of a non-trivial *revelation gap* – that the prior-independent approximation factor of the best non-revelation mechanism is better than that of the best revelation mechanism – gives concrete motivation for a theory of mechanism design without the revelation principle.

It is not hard to invent pathological scenarios where there is a non-trivial revelation gap. Instead, this paper considers the canonical environment of welfare maximization for agents with budgets and shows such a gap even for distributions on preferences that satisfy a standard regularity property. Moreover, the environment in which we exhibit the revelation gap suggests the end-to-end principle: the agents can easily implement the optimal outcome in the equilibrium of a simple mechanism, while revelation mechanisms that satisfy the constraints must be complex and either prior-dependent or non-optimal.

Main Results: Our analysis focuses on welfare maximization in a canonical single-item environment with ex ante symmetric budget constrained agents, i.e., each agent's value is drawn independently and identically from an unknown distribution and the agent cannot make payments that exceed a known and identical budget (cf. Maskin, 2000). Our main treatment of this problem will make a simplifying assumption that the distribution follows a regularity property that implies that the Bayesian optimal mechanism has a nice form (Pai and Vohra, 2014). Our results require a symmetric environment, i.e., an i.i.d. distribution and identical budget. A number of our results extend to the objective of revenue (Laffont and Robert, 1996), to position environments as popularized as a model for ad auctions (Devanur et al., 2013), and beyond regular distributions (Devanur et al., 2013). For clarity, the main results are described first for welfare maximization, single-item environments, and regularly distributed agent values.

The main challenge in demonstrating a revelation gap is that it is difficult to identify prior-independent optimal mechanisms, cf. Fu et al. (2015). Though the question has been considered, the prior literature has rarely identified optimal prior-independent mechanisms for non-trivial environments. [1] Our non-trivial revelation-gap theorem follows from three results. First, the all-pay auction (from the literature, defined below) has a unique equilibrium that is Bayesian optimal and it is prior-independent. Second, we obtain a lower bound on the ability of a prior-independent revelation mechanism to approximate the Bayesian optimal mechanism by identifying the dominant strategy incentive compatible mechanism that is Bayesian optimal for the uniform distribution. The performance of this mechanism is strictly worse than that of the Bayesian optimal mechanism (which is Bayesian incentive compatible); specifically the gap is 1.013. Third, we show that the dominant strategy incentive compatible clinching auction (from the literature, defined below) is an $e \approx 2.72$ approximation to the Bayesian optimal mechanism. Combining the upper and lower bounds we see a revelation gap between 1.013 and e. [2] The first result follows naturally from the literature; the second and third results are the main technical contributions of the paper.

Three auctions are at the forefront of our study. The *all-pay auction* solicits bids, assigns the item to the highest agent, and charges all agents their bids. The *clinching auction* (Ausubel, 2004; Dobzinski et al., 2008; Goel et al., 2015) is an ascending price auction that can be thought of as allocating a unit measure of lottery tickets: a price is offered in each stage, each agent specifies the measure of tickets desired at the given price, each agent is allocated a number of tickets that is equal to the minimum of her demand

[1] A contemporaneous exception, Allouah and Besbes (2018) showed that the second-price auction is the prior-independent optimal revelation mechanism for revenue when the agents' values are distributed according to a monotone hazard rate distribution.

[2] Note that the lower bound of 1.013 is shown for uniform distributions on two agents while the upper bound of e is over all numbers of agents and (regular) distributions. Restricting to two agents and uniform distributions, the comparison between the Bayesian optimal dominant strategy incentive compatible mechanism and the clinching auction become much tighter, 1.013 versus 1.03. Improved analyses are needed for the general case.

and the measure of remaining tickets if this agent is only allowed to buy tickets after all other agents have bought as much as they desire first. [3] The *middle-ironed clinching auction* – which we identify as the optimal dominant strategy incentive compatible mechanism – behaves like the clinching auction except that values that fall within a middle range are ironed. The allocation that an agent in this middle range receives is the average over he original allocation of for middle range values in the clinching auction. This averaging results in the budget binding later and more efficient outcomes than in the original clinching auction.

The second step, mentioned above, is to obtain a lower bound on the prior-independent approximation of a revelation mechanism. Our analysis begins with the observation that a prior-independent revelation mechanism must be Bayesian incentive compatible for every distribution. For two agents, this condition is equivalent to being dominant strategy incentive compatible. We ask whether there a gap between the Bayesian optimal dominant strategy and Bayesian incentive compatible mechanism. The comparison between optimal dominant strategy and Bayesian incentive compatible mechanism is standard for multi-dimensional mechanism design problems, e.g., see Gershkov et al. (2013) and Yao (2017); we are unaware of previous studies of this phenomenon for single-dimensional agents with non-linear preferences. We answer this question positively by writing the dominant strategy mechanism design problem as a linear program and solving it by identifying a dual solution that proves the optimality of the middle-ironed clinching auction, cf. Pai and Vohra (2014) and Devanur and Weinberg (2017). The identified gap gives a lower bound on the approximation factor of the optimal prior-independent mechanism.

The third step, mentioned above, proves that the prior-independent approximation factor of the clinching auction auction is at most e and resolves in the affirmative an open question from Devanur et al. (2013). Our proof follows from a novel adaptation of a standard method for approximation results in mechanism design where an auction's performance is compared to the upper bound given by the ex ante relaxation, in this case, the welfare of the optimal mechanism that sells one item in expectation over the random draws of the agents' values (i.e., ex ante) rather than for all draws of the agents' values (i.e., ex post). This method was introduced by Chawla et al. (2007), formalized by Alaei (2014), generalized by Alaei et al. (2013), and employed in many subsequent analyses.

Extensions: A number of our results extend beyond regular distributions, single-item environments, and the welfare objective as described above. These extensions all require that the environment be symmetric, specifically, that the agents' values are independent and identically distributed and their budgets are identical.

For irregular distributions the welfare-optimal auction is not generally the all-pay auction; moreover, it does not generally have a prior-independent implementation. We prove that the all-pay auction is a prior-independent two approximation. Moreover, both the regular and irregular prior-independent optimality and approximation results for the one-item all-pay auction extend to the all-pay position auction.

The degradation of the approximation factor by a factor of two for irregular distributions extends to the single-item clinching auction which is an e approximation for regular distributions (as described above) and a $2e$ approximation for irregular distributions.

For the revenue objective with the appropriate definition of regularity (Laffont and Robert, 1996), the n-agent single-item all-pay auction is a prior-independent $n/(n-1)$ approximation to the revenue optimal auction (cf. Bulow and Klemperer, 1996).

Important Directions: The most general direction suggested by this work is for a systematic development of non-revelation mechanism design. Unfortunately, it is not generally helpful to do revelation mechanism design and then try to go from the suggested revelation mechanism to a practical and simple non-revelation mechanism. Papers working to develop a theory of non-revelation mechanism design include Chawla and Hartline (2013), which proves the uniqueness and optimality of equilibria in symmetric rank-based auctions; Chawla et al. (2014b, 2016), which gives data driven methods for optimizing non-revelation mechanisms in symmetric environments; and Hartline and Taggart (2016), which gives a theory for non-revelation sample complexity and the design of approximately optimal non-revelation mechanisms in asymmetric environments.

While the literature has many interesting approximation bounds for prior-independent mechanism design, rarely have prior-independent optimal mechanisms been identified. Moreover, the prior-independent approximation factors achievable tend to be surprising; for example, Fu et al. (2015) show that the second-price action is not the optimal prior-independent mechanisms for two-agent revenue maximization with agents with regularly

[3] For example, at a price of 0 all agents would want to buy all the tickets, but the agent that arrives last gets no tickets, thus no agents get any tickets at this price; the price increases.

distributed values. [4] The literature lacks general techniques for answering this question.

We have observed that there is a very simple prior-independent optimal mechanism for welfare maximization in symmetric environments for agents with identical budgets. This mechanism, namely the all-pay auction, achieves its optimal outcome in Bayes-Nash equilibrium. The general question of identifying prior-independent non-revelation mechanisms that optimize a desired objective, like welfare or revenue, needs to be asked with care. Without restrictions to this question, it is asked and answered in the literature on non-parametric implementation theory, see the survey of Jackson (2001). This literature shows that arbitrarily close approximations, called "virtual implementations", to the Bayesian optimal mechanism can be implemented by an uninformed principal. The mechanisms in this literature tend to be sequential – where agents interact in multiple rounds – and require agents to make reports about their own preferences and crossreports about their beliefs on other agents' preferences. Our perspective on these results is that they take the model of Bayes-Nash equilibrium too literally and the resulting cross-reporting mechanisms are both fragile and impractical. One approach for ruling out these mechanisms is to restrict attention to mechanism formats that are commonly occurring in practice. Specifically, in the general *winner-pays-bid* format: agents bid, an allocation rule maps bids to a set of winners, and all winners pay their bids; in the general *all-pay* format: agents bid, an allocation rule maps bids to a set of winners, and all agents pay their bids. There may be other restricted formats that are also interesting for specific scenarios, e.g., the seller-offer mechanisms that are prevalent as real estate exchange mechanisms (Niazadeh et al., 2014).

Finally, there are still many gaps in our understanding of auctions for identically distributed agents with common budgets. For welfare, the bounds in this paper show that the clinching auction's approximation factor for the welfare objective is in $[1.03, e]$ for regular distributions and $[2, 2e]$ for irregular distribution. Sharpening these bounds is an open question. Moreover, we conjecture that the clinching auction is also a prior-independent constant approximation for the revenue objective (with i.i.d. public-budget regular distributions). We also leave open a number of questions with regard to the Bayesian optimal dominant strategy incentive compatible mechanism for agents with budgets. We conjecture that the welfare optimality of the middle-ironed clinching auction extends from uniform distributions to regular distributions. We leave open the question of a similar result for the revenue objective, even for the special case of uniform distributions. There are specific issues that prevent straightforward generalization of our approach of using the dual to certify the optimality of the middle-ironed clinching auction for these questions.

Other Related Work: There is a significant area of research analyzing the performance of simple non-revelation mechanisms in equilibrium, a.k.a., the *price of anarchy*. Generally these mechanisms are prior-independent and the aim of the literature, e.g. Syrgkanis and Tardos (2013), is to demonstrate that they are approximately efficient. On the other hand, for welfare maximization in many of the studied environments, there is a DSIC revelation mechanism that is (exactly) efficient and, thus, there is no revelation gap. Though this literature focuses on the analysis rather than the design of mechanisms, two conclusions for mechanism design are: (a) that a simple *revenue covering property* is sufficient (Hartline et al., 2014), necessary (Dütting and Kesselheim, 2015), and potentially optimizable; and (b) that this property (and also a more general *smoothness property*) is closed under composition, i.e., when multiple independent mechanisms are run simultaneously (Syrgkanis and Tardos, 2013). For a surveys of these and other results see Roughgarden et al. (2017) and Chapter 6 of Hartline (2016).

For agents with budgets, approximation mechanisms have been studied from both a design and analysis perspective for the *liquid welfare* benchmark proposed by Chawla et al. (2011), Syrgkanis and Tardos (2013), and Dobzinski and Paes Leme (2014). The liquid welfare benchmark is the optimal surplus of a feasible allocation when each agent's contribution to the surplus is the minimum of her value for her allocation and her budget. These and subsequent papers show that simple mechanisms have welfare that approximate the liquid welfare benchmark. Unfortunately, when evaluated under the formal study of benchmarks for mechanism design developed by Hartline and Roughgarden (2008) and summarized in Chapter 7 of Hartline (2016), the liquid welfare does not satisfy a key property. Specifically, there are i.i.d. distributions where the expected welfare of the Bayesian optimal mechanism is arbitrarily larger than the expected optimal liquid welfare. This bound follows because liquid welfare is at most the sum of the agent budgets which can be arbitrarily close to zero and, in such cases, is unrelated to the welfare possible

[4] Allouah and Besbes (2018) show that with more restrictive monotone hazard rate distributions, the second-price auction is an optimal prior-independent revelation mechanism.

by a mechanism. [5] Thus, testing mechanisms for their worst-case approximation factor with respect to liquid welfare does not necessarily separate good mechanisms from bad mechanisms.

II. Preliminaries

Model for auctions with budgets: We consider mechanisms for agents with independent and identically distributed values and identical public budgets. The budget is denoted by B. For allocation $x \in [0, 1]$ and payment $p \in \mathbb{R}$, an agent with value $v \in \mathbb{R}$ has utility $vx - p$ if p is at most the budget B and utility $-\infty$ otherwise. In other words, the agent cannot under any circumstances pay more than her budget. The agents' values are drawn independently and identically from distribution F with density f and support $[0, h]$.

Denote the strategy function of an agent by $s(\cdot)$ where $s(v)$ is the bid made by the agent when her value is v. A bid profile is $\boldsymbol{b} = (b_1, \ldots, b_n)$. A mechanism is given by mapping from bids to allocations and payments which we will denote by $\tilde{\boldsymbol{x}}(\boldsymbol{b}) = (\tilde{x}_1(\boldsymbol{b}), \cdots, \tilde{x}_n(\boldsymbol{b}))$ and $\tilde{\boldsymbol{p}}(\boldsymbol{b}) = (\tilde{p}_1(\boldsymbol{b}), \cdots, \tilde{p}_n(\boldsymbol{b}))$. The outcome of the mechanism $(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{p}})$ and strategy profile $\boldsymbol{s} = (s_1, \ldots, s_n)$ on a profile of agent values $\boldsymbol{v}$ is denoted by allocation rule $\boldsymbol{x}(\boldsymbol{v}) = \tilde{\boldsymbol{x}}(\boldsymbol{s}(\boldsymbol{v}))$ and payment rule $\boldsymbol{p}(\boldsymbol{v}) = \tilde{\boldsymbol{p}}(\boldsymbol{s}(\boldsymbol{v}))$.

The auction designer typically faces a feasibility constraint that restricts the allocations that can be produced. For example, a single-item auction requires that the number of agents allocated is at most one, i.e., $\sum_i x_i(\boldsymbol{v}) \leq 1$. A position environment generalizes a single item auction and is given by a sequence of decreasing probabilities $\boldsymbol{w} = (w_1, \ldots, w_n)$ where without loss of generality the number of positions is equal to the number of agents. The probability that an agent is allocated if assigned to position j is w_j. A mechanism then can assign agents to positions (deterministically or randomly) and this process and the position probabilities induce allocation probabilities $x_i(\boldsymbol{v})$.

Basic auction theory: A Bayes-Nash equilibrium (BNE) in the mechanism $(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{p}})$ is a profile of agent strategies $\boldsymbol{s}$ where each s_i maps a value to a bid that is a best response to the other strategies and the common knowledge that values are drawn i.i.d. from distribution F. The mechanism $(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{p}})$ and strategy profile $\boldsymbol{s}$ induce, for each agent, an interim allocation rule $x_i(v) = \mathbf{E}_{\boldsymbol{v}_{-i}}[x_i(v, \boldsymbol{v}_{-i})]$. We will consider BNE only for symmetric distributions and symmetric rank-based auctions. In such auctions, Chawla and Hartline (2013) show that all equilibria are symmetric, thus it is without loss to drop the subscript and refer to the interim allocation rule and payment rule as (x, p). The Myerson (1981) characterization of BNE requires that (a) the interim allocation $x(v)$ is monotone non-decreasing and (b) the interim payment $p(v) = v \cdot x(v) - \int_0^v x(t)\, dt$. Condition (b) is known as the payment identity. A mechanism is Bayesian incentive compatible (BIC) if it induces a BNE where each agent's strategy is reporting her value truthfully. A mechanism is interim individually rational (IIR) if the interim utility of each agent is non-negative for all values.

A dominant strategy equilibrium (DSE) in the mechanism $(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{p}})$ is a profile of agent strategies $\boldsymbol{s}$ where each s_i maps a value to a bid that is a best response regardless of what other agents are doing. The characterization of DSE follows from the BNE characterization: (a) the allocation $x(v_i, \boldsymbol{v}_{-i})$ is monotone non-decreasing in v_i and (b) the payment $p(v_i, \boldsymbol{v}_{-i}) = v_i \cdot x(v_i, \boldsymbol{v}_{-i}) - \int_0^{v_i} x(t, \boldsymbol{v}_{-i}) dt$. A mechanism is dominant strategy incentive compatible (DSIC) if it induces a DSE where each agent's strategy is reporting her value truthfully. A mechanism is ex-post individually rational (ex-post IR) if each agent's utility is non-negative for all valuation profiles.

Optimal auctions with budgets: Laffont and Robert (1996) and Maskin (2000) for the revenue and welfare objectives, respectively, characterize the Bayesian optimal mechanisms for agents with i.i.d. distributions and identical public budgets. With the following regularity assumptions on the distribution, defined distinctly for revenue and welfare, the optimal mechanism has a nice form.

Definition II.1. *A single-dimensional public budget agent is* public-budget regular for welfare *if her cumulative distribution function $F(\cdot)$ is concave; she is* public-budget regular for revenue *if additionally $v - \frac{1-F(v)}{f(v)}$ is non-decreasing.*

The results of Laffont and Robert and Maskin can be reinterpreted, à la Alaei et al. (2013), as solving a single-agent interim optimization problem that is given by an interim constraint $x^{\star}(\cdot)$. An interim allocation is interim feasible under the interim constraint $x^{\star}(\cdot)$ if for all values $v \in [0, h]$, the probability of allocating the item to an agent with value greater than v with allocation rule $x(\cdot)$ is at most that with allocation rule $x^{\star}$, i.e. $\int_v^h x(t)\, dF(t) \leq \int_v^h x^{\star}(t)\, dF(t)$. In many cases solution to these interim optimization problems will take the form of the original constraint with ironed interval and reserve. Ironing on arbitrary interval $[v^{\dagger}, v^{\ddagger}]$

[5] For example, in a single-item environment with budgets identically equal to zero and agent values identically equal to one (trivially an i.i.d. distribution); the lottery, which allocates the item to a random agent for free, has welfare one while the liquid welfare is zero.

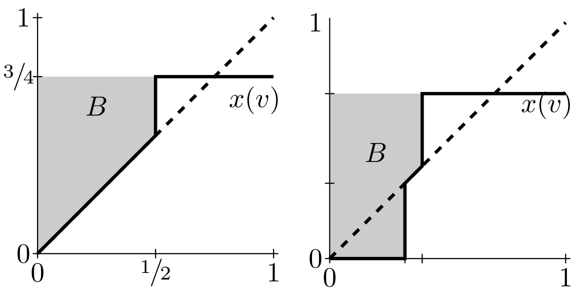

(a) Welfare maximization (b) Revenue maximization

Figure 1: Depicted are the interim allocation rules of the welfare-optimal and revenue-optimal mechanisms for two agents with uniform values on $[0, 1]$. In each figure the highest-bid-wins allocation rule is depicted with a dashed line.

corresponds to the distribution weighted averaging as follows, $x(v) = \int_{v^\dagger}^{v^\ddagger} x^\star(t)\, dF(t)$ for all $v \in [v^\dagger, v^\ddagger]$. Reserve pricing at value $v^\dagger$ corresponds to rejecting all values below $v^\dagger$ as follows, $x(v) = 0$ for all $v \in [0, v^\dagger]$. An important allocation constraint is that given by the *highest-bid-wins* allocation rule. The highest-bid-wins allocation rule for n agents and with values from cumulative distribution function F is $x^\star(v) = F^{n-1}(v)$, e.g., for two agents with uniform values it is $x^\star(v) = v$.

Theorem II.1 (Laffont and Robert, 1996; Maskin, 2000; Alaei et al., 2013)**.** *For public-budget regular i.i.d. agents and interim allocation constraint $x^\star(\cdot)$, the welfare-optimal single-agent mechanism allocates as by $x^\star(\cdot)$ except that values in $[v^\dagger, h]$ are ironed for some $v^\dagger$ and the revenue-optimal single-agent mechanism additionally reserve prices values in $[0, v^\ddagger]$ for some $v^\ddagger$; payments are given deterministically by the payment identity.*

For single-item environments, one possible implementation of Theorem II.1 is the all-pay auction. The all-pay auction has a unique Bayes-Nash equilibrium which is identical to outcome described in Theorem II.1 for the allocation constraint given by the highest-bid-wins allocation rule.

Definition II.2 (all-pay auction)**.** *The* all-pay auction *is the mechanism $(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{p}})$ where $\tilde{\boldsymbol{x}}(\cdot)$ allocates item to the agent with the highest bid with ties broken at random and $\tilde{\boldsymbol{p}}(\cdot)$ charges each agent their bid, i.e. $\tilde{p}_i(\boldsymbol{b}) = b_i$.*

Theorem II.2 (Maskin, 2000)**.** *For public-budget regular i.i.d. agents, the all-pay auction is welfare optimal.*

III. Welfare of the Clinching Auction

In this section, we study a prior-independent revelation mechanism called the clinching auction in single-item environments. Dobzinski et al. (2008) gave the following formulation of the clinching auction and characterized properties of its outcome. See Figure 3b.

Definition III.1 (clinching auction)**.** *The* clinching auction *maintains an allocation and price-clock starting from zero. The price-clock ascends continuously and the allocation and budget are adjusted as follows.*

1) *Agents whose values are less than price-clock are removed and their allocation is frozen.*
2) *The demand of any remaining agent is the remaining budget divided by the price-clock.*
3) *Each remaining agent clinches (and adds to their current allocation) an amount that corresponds to the largest fraction of their demand that can be satisfied when all other remaining agents are first given as much of their demand as possible (subject to the feasibility constraint).*
4) *The budget and allocation are updated to reflect the amount clinched in the previous step.*

Proposition III.1 (Dobzinski et al., 2008)**.** *For public-budget agents, the clinching auction always allocates all items, is ex-post IR, and is DSIC.*

We use the following approach to show that the clinching auction is an e-approximation for public-budget regular agents. We relax the feasibility constraint to an ex ante constraint and show that the optimal mechanism that sells to each agent with ex ante probability $1/n$ simply posts a price (of exactly B assuming that the budget binds) for a chance to win the item (Lemma III.2, below). This simple form of mechanism is closely approximated by the mechanism that divides the item into k lotteries of $1/k$ probability and sells them via the clinching auction (full details given subsequently). A key property of this *clinching auction with lotteries* is that with constant probability the budget does not bind. This property allows the clinching auction welfare to be compared to that of the ex ante relaxation.

Consider the welfare-optimal auction for n agents. Since agents are symmetric, each agent will win with ex ante probability exactly $\frac{1}{n}$. We replace the feasibility constraint, that ex post allocation cannot allocate more than one item (i.e. $\sum_{i \in [n]} \tilde{x}_i(\boldsymbol{v}) \leq 1$ for all $\boldsymbol{v}$), with a $\frac{1}{n}$ ex ante constraint that each agent cannot be allocated more than $\frac{1}{n}$ in expectation (i.e. $\mathbf{E}_v[x(v)] \leq {}^1\!/_n$). Ex ante optimal mechanisms for agents with public budgets were proposed and studied by Alaei et al. (2013).

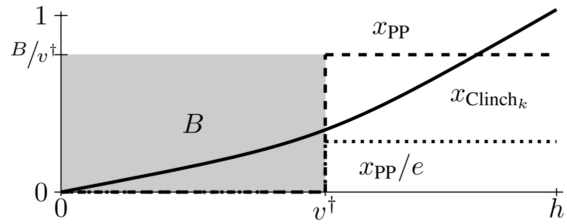

Figure 2: The allocation rules of the ex ante relaxation (dashed), an $1/e$-fraction of the ex ante relaxation (dotted), and the clinching auction with lotteries (solid) are depicted. The clinching auction with lotteries pointwise exceeds an $1/e$-fraction of the ex ante relaxation.

Lemma III.2 (Alaei et al., 2013)**.** *For i.i.d. public-budget regular agents with budget B, the* ex ante welfare-optimal mechanism *is either:*

i. Budget binds: Post the price B for allocation probability $\frac{B}{v^\dagger} \leq 1$ with $v^\dagger$ set to satisfy $\frac{1}{n} = \frac{B}{v^\dagger}(1 - F(v^\dagger))$. Values $v \geq v^\dagger$ select the lottery.

ii. Allocation probability binds: Post price $v^\dagger = F^{-1}(1 - \frac{1}{n})$ for allocation probability one.

We build the connection between the clinching auction and the ex ante optimal mechanism by considering an additional auction: the *clinching auction with lotteries*, denoted Clinch_k, which allocates k lotteries with winning probability $1/k$ per lottery, using the clinching auction framework under the same public budget. [6]

Lemma III.3 (a special case of Devanur et al., 2013)**.** *For valuation profile $\boldsymbol{v}$, budget B, $v_{(i)}$ denoting the i-th largest value, i denoting the largest integer such that $v_{(i)} \geq B \cdot i$, and $\kappa = \max(i, k)$; in the clinching auction with k lotteries the agents with highest $(\kappa - 1)$ values win with same probability in $[1/\kappa, 1/k]$ and the agent with the κ-th highest value wins with the remaining probability (so that the total probability is one).*

Lemma III.4 below shows that by selecting an appropriate k, the interim probability that an agent with value $v^\dagger$ wins in the clinching auction with lotteries is at least an e fraction of the interim probability that the agent (with value $v^\dagger$) wins in the ex ante relaxation. See Figure 2. The proof of the lemma is in the full version of the paper.

Lemma III.4. *For i.i.d. public budget agents, at value $v^\dagger$ defined in Lemma III.2, there exists $k \in [n]$, such that the interim allocation of the clinching auction with lotteries $x_{Clinch_k}(v^\dagger)$ is an e-approximation of the interim allocation of the ex ante optimal mechanism $x_{PP}(v^\dagger)$, i.e., $x_{Clinch_k}(v^\dagger) \geq {}^1\!/_e \cdot x_{PP}(v^\dagger)$.*

We now prove our main theorem about the approximation ratio for the clinching auction.

Theorem III.5. *For i.i.d. public-budget regular agents, the clinching auction is an e-approximation to the welfare-optimal mechanism.*

Proof: By Lemma III.2 the interim allocation rule of the ex ante optimal mechanism is a step function that steps at value $v^\dagger$. By Lemma III.4, at value $v^\dagger$, the allocation rule of the clinching auction with lotteries is an e-approximation to that of the ex ante optimal mechanism. The allocation rule of the clinching auction with lotteries is monotone, so its allocation rule is an e-approximation to that of the ex ante optimal mechanism at every value. Consequently, the expected welfare of the clinching auction with lotteries is at least an e-approximation to that of the ex ante relaxation. See Figure 2.

Finally, Lemma III.3 implies that for every ex post valuation profile, the welfare of the clinching auction is at least that of the clinching auction with lotteries. ■

For i.i.d. public-budget regular agents, the all-pay auction is optimal while the clinching auction is not; specifically, the budget binds for more valuation profiles in the clinching auction than in the all-pay auction. The following lemma results from searching over uniform distributions for the one with the worst ratio. See the full paper for its proof.

Lemma III.6. *There exists the instance of i.i.d. public-budget regular agents where the clinching auction is a 1.03-approximation of the welfare-optimal mechanism.*

IV. Bayesian Optimal DSIC Mechanisms

In Theorem II.2, the all-pay auction is welfare-optimal under public-budget regular distribution. Hence, applying the revelation principle to the all-pay auction, it produces a Bayesian optimal revelation mechanism. This mechanism is BIC but not DSIC. In this section, we characterize the optimal DSIC mechanism for two agents with uniformly distributed values. We obtain a lower bound on its approximation ratio with the BIC optimal mechanism.

We first introduce the middle-ironed clinching auction (for two agents).

Definition IV.1. *The two-agent* middle-ironed clinching auction *is parameterized by $v^\dagger \leq B$ and $v^\ddagger = 2B - v^\dagger$ and its outcome is highest-bid-wins on values less that*

[6]The clinching auction with k lotteries is a special case of the clinching auction in a position environment, specifically, with position weights $w_1 = \cdots = w_k = {}^1\!/_k$ and $w_{k+1} = \cdots = w_n = 0$.

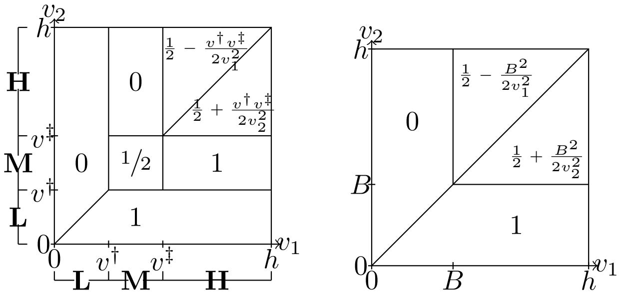

(a) Middle-ironed clinching auction (b) Clinching auction

Figure 3: The comparison of the allocation rule $x_1(v_1, v_2)$ for the middle-ironed clinching auction and the clinching auction. In the middle-ironed clinching auction, for the values in interval $\mathbf{M}$ can be thought as "ironed", i.e. an agent receives the same outcome for any value $v \in \mathbf{M}$.

$v^\dagger$, a fair lottery on values in $[v^\dagger, v^\ddagger]$, and the clinching auction on values exceeding $v^\ddagger$; a precise formulation for two-agents is given in Figure 3a.

For two-agents case, the middle-ironed clinching auction allocates the item efficiently except for value profiles in $\mathbf{MM}$ (both agents with values in $\mathbf{M}$) or $\mathbf{HH}$ (both agents with values in $\mathbf{H}$). For the value profile in $\mathbf{MM}$, it randomly allocates the item to one of the agent with probability $\frac{1}{2}$ with payment $\frac{v^\dagger}{2}$. For the value profile in $\mathbf{HH}$, it allocates the item such that the budget binds for the agent with higher value and the allocation rule depends on the lower value only. Figure 3b depicts the allocation rule of the clinching auction for comparison. The middle-ironed clinching auction can be implemented with an ascending price via a generalization of the clinching auction that allows for price jumps which we develop in the full version of the paper.

We will show that by selecting the proper thresholds $v^\dagger$ and $v^\ddagger$, the middle-ironed clinching auction is the Bayesian optimal DSIC mechanism for two agents with uniformly distributed values. An intuition behind the optimality of the middle-ironed clinching auction is as follows: Dobzinski et al. (2008) show that for two public budget agents, the clinching auction is the only Pareto optimal (i.e. there is no outcome which is weakly better for all agents and strictly better for one agent) and DSIC auction. Moreover, after the price increases past the point where the budget binds, a differential equation governs the allocation of any DSIC mechanism. Our goal is to optimize expected welfare rather than satisfy Pareto optimality. Sacrificing welfare for lower-valued agents by ironing can delay the budget from binding and enable greater welfare from higher-valued agents. From our proof of optimality, it is sufficient to only iron one region in the middle of value space.

Theorem IV.1. *For two public-budget agents with budget B and value uniformly drawn from $[0, h]$, Bayesian optimal DSIC mechanism is the middle-ironed clinching auction with some thresholds $v^\dagger$ and $v^\ddagger$.*

The approach of the proof is to write down our problem as a linear program (primal), assume the middle-ironed clinching auction to be the solution, and then construct the dual program with a dual solution which witnesses the optimality of the primal solution by complementary slackness. This approach is reminiscent of that of Pai and Vohra (2014) and Devanur and Weinberg (2017); however, our multi-agent DSIC constrained program presents novel challenges and for this reason we only solve the problem of two agents and uniform distributions. In this extended abstract, we only solve a discrete version of the problem. In the full paper, we additinally solve the continuous version as the limit of the discrete version.

Consider the value distribution with finite value space $[h] = \{1, 2, \ldots, h\}$ with equal probability each. We begin by writing down the optimization program for welfare maximization among all possible DSIC mechanism.

$$\begin{aligned} \sup_{(\boldsymbol{x},\boldsymbol{p})} \quad & \sum_{v_1, v_2 \in [h]} (v_1 \cdot x_1(v_1, v_2) + v_2 \cdot x_2(v_1, v_2)) \cdot \tfrac{1}{h} \cdot \tfrac{1}{h} \\ s.t. \quad & \\ & (\boldsymbol{x},\boldsymbol{p}) \text{ are DSIC, ex-post IR, and feasible,} \\ & (\boldsymbol{x},\boldsymbol{p}) \text{ is budget balanced.} \end{aligned}$$

By the characterization of dominant strategy equilibrium, we simplify this optimization program into a linear program as follows,

$$\begin{aligned} \max_{(\boldsymbol{x},\boldsymbol{p}) \geq 0} \quad & \sum_{v_1, v_2 \in [h]} v_1 \cdot x(v_1, v_2) \\ s.t. \quad & \\ & h \cdot x(h, v_2) - \textstyle\sum_{t=1}^{h} x(t, v_2) \leq B && \forall v_2 \in [h] \\ & x(v_1, v_2) + x(v_2, v_1) \leq 1 && \forall v_1, v_2 \in [h] \\ & x(v_1, v_2) \leq x(v_1 + 1, v_2) && \forall v_1 \in [h-1], v_2 \in [h] \end{aligned}$$

where we assume $x_1(a, b) = x_2(b, a) = x(a, b)$ for all $a, b \in [h]$, since it is an agent-symmetric program. [7]

Additionally, we relax the monotonicity constraint by replacing it with $x(v_1, v_2) \leq x(h, v_2)$ which is common for Bayesian mechanism design with public budget agents.

$$x(v_1, v_2) \leq x(h, v_2) \quad \text{for all } v_1 \in [h-1], v_2 \in [h]$$

[7] Note that the program in terms of $x(a, b)$ is asymmetric.

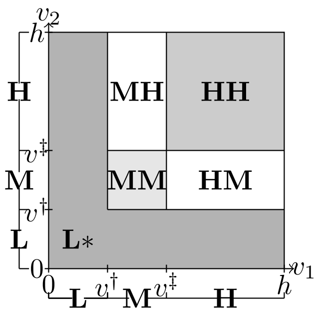

Figure 4: We partition the dual variables into $\mathbf{L}*$ (at least one agent with value in $\mathbf{L}$), $\mathbf{HH}$ (both agents with values in $\mathbf{H}$), $\mathbf{MM}$ (both agents with values in $\mathbf{M}$), $\mathbf{MH}$ and $\mathbf{HM}$ (one agent with value in $\mathbf{M}$ and the other with value in $\mathbf{H}$) five regions.

The corresponding dual program can be written as follows. Let $\{\Lambda(v_2)\}_{v_2\in[h]}$ denote the dual variables for budget constraints; $\{\beta(v_1,v_2)\}_{v_1,v_2\in[h]}$ denote the dual variables for feasibility constraints (for simplicity, we use both $\beta(v_1,v_2)$ and $\beta(v_2,v_1)$ to denote the same dual variable); and $\{\mu(v_1,v_2)\}_{v_1\in[h-1],v_2\in[h]}$ denote the dual variables for monotonicity constraints. The dual program is,

$$\min_{(\boldsymbol{\Lambda},\boldsymbol{\beta},\boldsymbol{\mu})\geq 0} \sum_{v_2\in[h]} B\cdot\Lambda(v_2)+\frac{1}{2}\sum_{v_1,v_2\in[h]}\beta(v_1,v_2)$$

$$s.t.$$

$$-\Lambda(v_2)+\beta(v_1,v_2)+\mu(v_1,v_2)\geq v_1 \quad \forall v_1\in[h-1], v_2\in[h],$$

$$(h-1)\Lambda(v_2)+\beta(h,v_2)-\textstyle\sum_{t=1}^{h-1}\mu(t,v_2)\geq h \qquad \forall v_2\in[h]$$

The plan to solve the program is as follows. For each possible thresholds $v^\dagger, v^\ddagger$ chosen in the middle-ironed clinching auction, we first construct a solution in dual which satisfies the complementary slackness with this middle-ironed clinching auction as a solution in primal. These induced dual solutions may be infeasible. Next, we will show that there exists a pair of thresholds $v^\dagger, v^\ddagger$ which induces a feasible dual solution. This feasible dual solution witnesses the optimality of the middle-ironed clinching auction.

We will partition the dual variables into following five regions ($\mathbf{L}*$, $\mathbf{MM}$, $\mathbf{HH}$, $\mathbf{MH}$ and $\mathbf{HM}$) as in Figure 4; and construct the dual solution for them separately. We denote $\boldsymbol{\lambda}$ as the discrete derivative of the dual variable $\boldsymbol{\Lambda}$, i.e. $\lambda(v)=\Lambda(v)-\Lambda(v+1)$.

$\boldsymbol{\Lambda}$ in L: Since the budget constraints do not bind, by complementary slackness,

$$\Lambda(v)=0 \text{ for all } v\in\mathbf{L}.$$

$\boldsymbol{\beta},\boldsymbol{\mu}$ in L*: Let (v,v') be a value profile in region $\mathbf{L}*$ such that $v\geq v'$. By complementary slackness on $x(v,v')$, $\beta(v,v')+\mu(v,v')-\Lambda(v')=v$ if $v<h$; $\beta(v,v')-\sum_{t=1}^{h-1}\mu(t,v')+(h-1)\Lambda(v')=v$ otherwise (i.e. $v=h$). We let

$$\beta(v,v')=v \text{ and } \mu(v,v')=0.^8$$

Since the relaxed monotonicity constraint does not bind at $x(v',v)$, i.e. $x(v',v)<x(h,v)$, the corresponding dual variable is

$$\mu(v',v)=0.$$

$\boldsymbol{\beta},\boldsymbol{\mu}$ in HH: Let (v,v') be a value profile in region $\mathbf{HH}$ such that $v\geq v'$. Since both agents win with non-zero probability, by complementary slackness on $x(v,v')$ and $x(v',v)$, the corresponding dual constraints bind. Since the relaxed monotonicity constraint does not bind at $x(v',v)$, the monotonicity dual variable is

$$\mu(v',v)=0.$$

The binding dual constraint of $x(v',v)$ is $\beta(v',v)-\Lambda(v)+\mu(v',v)=v'$. Hence,

$$\beta(v,v')=\beta(v',v)=v'+\Lambda(v).^9$$

The binding dual constraint of $x(v,v')$ is $\beta(v,v')-\Lambda(v')+\mu(v,v')=v$. Note that the relaxed monotonicity constraint is tight for (v,v'). Hence,

$$\mu(v,v')=v-v'+\Lambda(v')-\Lambda(v).$$

Here we write $\boldsymbol{\beta},\boldsymbol{\mu}$ as terms of $\boldsymbol{\Lambda}$. In the next paragraph, we will solve for $\boldsymbol{\Lambda}$.

$\boldsymbol{\Lambda}$ in H: Let $v\in\mathbf{H}$. Consider the binding dual constraint of $x(h,v)$, $(h-1)\Lambda(v)+\beta(h,v)-\sum_{t=1}^{h-1}\mu(t,v)=h$. Notice that by complementary slackness, $\mu(t,v)=0$ for all $t\leq v$. Plugging $\boldsymbol{\beta}$ and $\boldsymbol{\mu}$ as terms of $\boldsymbol{\Lambda}$ into the these dual constraints of $x(h,v)$, we can solve for $\boldsymbol{\Lambda}$ as

$$\lambda(v)=\frac{h-v}{v} \text{ for all } v\in\mathbf{H} \text{ and } \Lambda(h)=0.^{10}$$

$\boldsymbol{\beta},\boldsymbol{\mu}$ in MM and $\boldsymbol{\Lambda}$ in M: Let (v,v') be a value profile in region $\mathbf{MM}$ such that $v\geq v'$. Since the relaxed monotonicity constraints do not bind for either $x(v,v')$ or $x(v',v)$, the corresponding dual variables are

$$\mu(v,v')=\mu(v',v)=0.$$

[8] An intuition here is: $\boldsymbol{\mu}$ are the dual variables for the relaxed monotonicity constraint and can be thought as indicators of ironing. Though the monotonicity constraint binds, this is not because of ironing but binding allocation (i.e. $x(\cdot)\leq 1$). Therefore, we set $\boldsymbol{\mu}$ as zero.

[9] Recall that $\beta(v,v')$ and $\beta(v',v)$ denote the same dual variable.

[10] Recall that $\boldsymbol{\lambda}$ is the discrete derivative of dual variables $\boldsymbol{\Lambda}$, so $\Lambda(v)=\sum_{t=v}^{h-1}\lambda(v)$.

The binding dual constraints of $x(v, v')$ implies $\beta(v, v') = v' + \Lambda(v)$. On the other hand, the binding dual constraints of $x(v', v)$ implies $\beta(v', v) = v + \Lambda(v')$. Recall that $\beta(v, v')$ and $\beta(v', v)$ denote the same variable, hence,

$$\lambda(v) = -1 \text{ for all } v \in \mathbf{M} \setminus \{v^{\ddagger} - 1\},^{11}$$
$$\beta(v, v') = \Lambda(v^{\ddagger}) + \lambda(v^{\ddagger} - 1) + v + v' - v^{\ddagger}.$$

$\boldsymbol{\beta}, \boldsymbol{\mu}$ in MH and HM: Let (v, v') be a value profile in region $\mathbf{HM}$ such that $v > v'$. With the similar argument for region $\mathbf{HH}$,

$$\mu(v', v) = 0 \text{ and } \mu(v, v') = v - v' + \Lambda(v') - \Lambda(v),$$
$$\beta(v, v') = v' + \Lambda(v) \text{ if } v < h.$$

Plugging the above expressions for $\boldsymbol{\mu}$ into the binding dual constraint of $x(h, v')$,

$$\beta(h, v') = (h-1)(v^{\ddagger} - v') + 1 + (v^{\ddagger} - 1)\lambda(v^{\ddagger} - 1).$$

With the analysis above, we construct the dual solution which satisfies complementary slackness with the middle-ironed clinching auction as a solution in primal.

Lemma IV.2. *For the middle-ironed clinching auction with arbitrary thresholds $v^{\dagger}$ and $v^{\ddagger}$, the constructed dual solution satisfies the complementary slackness.*

Proof: The complementary slackness is directly implied by the construction. ■

Though the this dual solution satisfies the complementary slackness, it may be infeasible. Therefore, we argue that there exists some thresholds $v^{\dagger}$, $v^{\ddagger}$ and $\lambda(v^{\ddagger} - 1)$ under which the dual solution is feasible. See the full version for the proof of the following lemma.

Lemma IV.3. *There exists $v^{\dagger}, v^{\ddagger}$ and $\lambda(v^{\ddagger} - 1)$ such that the constructed dual solution is feasible.*

The construction of the dual solution which satisfies feasibility and complementary slackness witnesses the optimality of the middle-ironed clinching auction. We offer the following discrete version of Theorem IV.1.

Theorem IV.4. *For two public-budget agents with value uniformly distributed from $\{1, \ldots, h\}$, the Bayesian optimal DSIC mechanism is the middle-ironed clinching auction for some thresholds $v^{\dagger}$ and $v^{\ddagger}$.*

Based on Theorem IV.1, the performance of the welfare-optimal DSIC and BIC mechanisms can be compared.

[11] Complementary slackness does not pin down $\lambda(v^{\ddagger} - 1)$. We leave it as a variable and identify it later when we choose the thresholds $v^{\dagger}$, $v^{\ddagger}$ to ensure that the dual solution is feasible.

Lemma IV.5. *There exists the instance of public-budget regular agents where the welfare-optimal DSIC mechanism is a 1.013-approximation to the welfare-optimal BIC mechanism.*

Proof: Consider two agents with values drawn uniformly from $[0, h]$ where $h \geq 5.5$ and the budget $B = 1$. By Theorem IV.1, the welfare-optimal DSIC mechanism in this case is the middle-ironed clinching auction with $v^{\dagger} = 0$ and $v^{\ddagger} = 2$. The welfare-optimal BIC mechanism is the all-pay auction (applying the revelation principle). By computing the welfare for both mechanisms under this distribution and setting $h = 5.5$, the ratio is optimized as 1.013. ■

V. Revelation Gap

In the literature, prior-independent mechanisms have been shown to approximate the Bayesian optimal mechanism for many objectives. Interestingly, except when the solution is trivial, none of the approximation mechanisms in the literature are known to be optimal. The formal question of optimal prior-independent mechanism design is the following:

$$\beta = \min_{\mathcal{M} \in \text{MECHS}} \max_{F \in \text{DISTS}} \frac{\mathbf{E}_{\boldsymbol{v} \sim F}[\text{OPT}_F(\boldsymbol{v})]}{\mathbf{E}_{\boldsymbol{v} \sim F}[\mathcal{M}(\boldsymbol{v})]}.$$

In this definition, $\mathbf{E}_{v \sim F}[\mathcal{M}(v)]$ is the equilibrium performance of mechanism $\mathcal{M}$ on distribution F and OPT_F is the optimal mechanism for given objective on distribution F. Importantly in this definition, the family of mechanism MECHS may be restricted from all mechanisms and the family of distribution DISTS may be restricted from all distributions.

As discussed in the introduction, the revelation principle is not without loss for prior-independent mechanism design. Based on this idea, we introduce the concept of revelation gap.

Definition V.1. *The* revelation gap *is the ratio of the prior-independent approximation of incentive compatible mechanisms to the prior-independent approximation of (generally non-revelation) mechanisms.*

In this section, we consider welfare maximization with public-budget regular agents and, with the results from previous sections, we give upper and lower bounds on the revelation gap. The following is a corollary of Theorem III.5, which follows from observing that the clinching auction is a prior-independent revelation mechanism.

Corollary V.1. *For i.i.d. public-budget regular agents, the revelation gap for welfare maximization is at most e.*

For the lower bound, the welfare-optimal DSIC mechanism for two agent with uniformly distributed values was identified in Section IV. Note that for two-agent environments, the DSIC and ex-post IR constraints are equivalent to prior-independent BIC and IIR constraints. [12] With more than two agents, this equivalence may not hold.

Lemma V.2. *For two i.i.d. agents, a mechanism is Bayesian incentive compatible and interim individual rational for all i.i.d. distributions if and only if it is dominant strategy incentive compatible and ex-post individually rational.*

Proof: The direction that DSIC implies BIC for all i.i.d. distribution is trivial by the definition. To show the other direction, for arbitrary value v, consider the distribution which puts the whole mass on v. These distributions break the interim constraints in BIC into the ex-post constraints in DSIC for every valuation profile. Hence, BIC for all i.i.d. distribution implies DSIC for two agents setting. ■

Theorem V.3. *For i.i.d. public-budget regular agents, the revelation gap for welfare maximization is at least* 1.013.

Proof: As described in Section II, the all-pay auction is prior-independent. Moreover, Theorem II.2 states that it is welfare-optimal for public budget regular agents. Thus, the (non-revelation) all-pay auction is a prior-independent 1-approximation.

Theorem IV.1 states that the middle-ironed clinching auction is the Bayesian optimal DSIC mechanism for two agents with values drawn uniformly from $[0, h]$. Lemma V.2 shows that, in two-agent environments, DSIC is equivalent to BIC for all i.i.d. distributions. Lemma IV.5 shows that there is a uniform distribution where the middle-ironed clinching auction is a 1.013 approximation. The optimal prior-independent revelation mechanism on its worst-case public-budget regular distribution is no better. ■

VI. Other Results

Proof of the following results can be found in the full version of the paper.

Theorem VI.1. *For i.i.d. public-budget agents, the all-pay auction is a 2-approximation to the welfare-optimal mechanism (regularity is not assumed).*

Theorem VI.2. *For i.i.d. public-budget agents, the clinching auction is a* $2e$*-approximation to the welfare-optimal mechanism (regularity is not assumed).*

Theorem VI.3. *For i.i.d. public-budget regular agents, the welfare-optimal winner-pays-bid mechanism is the highest-bid-wins mechanism, i.e. the first-price auction.*

Lemma VI.4. *For* n *i.i.d. public-budget agents, the first-price auction is at best an* $\left(\frac{1}{4}n - o(n)\right)$*-approximation to the all-pay auction.*

Theorem VI.5. *For* $n \geq 2$ *i.i.d. public-budget regular agents, the all-pay auction is an* $\frac{n}{n-1}$*-approximation to the revenue-optimal mechanism.*

References

Alaei, S. (2014). Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972.

Alaei, S., Fu, H., Haghpanah, N., and Hartline, J. (2013). The simple economics of approximately optimal auctions. In *Proc. 54th IEEE Symp. on Foundations of Computer Science*.

Allouah, A. and Besbes, O. (2018). Prior-independent optimal auctions. In *Proceedings of the 19th ACM Conference on Economics and Computation*, pages 503–503.

Ausubel, L. M. (2004). An efficient ascending-bid auction for multiple objects. *The American Economic Review*, 94(5):1452–1475.

Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *The American Economic Review*, 86(1):180–194.

Chawla, S., Fu, H., and Karlin, A. (2014a). Approximate revenue maximization in interdependent value settings. In *Proceedings of the 15th ACM conference on Economics and computation*, pages 277–294.

Chawla, S. and Hartline, J. (2013). Auctions with unique equilbria. In *Proc. 14th ACM Conf. on Electronic Commerce*, pages 181–196.

Chawla, S., Hartline, J., and Kleinberg, R. (2007). Algorithmic pricing via virtual valuations. In *Proc. 8th ACM Conf. on Electronic Commerce*.

Chawla, S., Hartline, J., and Nekipelov, D. (2014b). Mechanism design for data science. In *Proceedings of the 15th ACM conference on Economics and computation*, pages 711–712.

Chawla, S., Hartline, J., and Nekipelov, D. (2016). A/b testing of auctions. In *Proceedings of the 17th ACM Conference on Economics and Computation*, pages 19–20.

Chawla, S., Hartline, J. D., Malec, D., and Sivan, B. (2013). Prior-independent mechanisms for scheduling. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 51–60.

Chawla, S., Malec, D. L., and Malekian, A. (2011). Bayesian mechanism design for budget-constrained agents. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 253–262.

Devanur, N., Ha, B., and Hartline, J. (2013). Prior-free auctions for budgeted agents. In *Proc. 14th ACM Conf. on Electronic Commerce*, pages 287–304.

Devanur, N. R., Hartline, J. D., Karlin, A. R., and Nguyen, C. T. (2011). Prior-independent multi-parameter mechanism design. In *WINE*, pages 122–133.

Devanur, N. R. and Weinberg, S. M. (2017). The optimal mechanism for selling to a budget constrained buyer: The general case. In *Proceedings of the 18th ACM Conference on Economics and Computation*, pages 39–40.

Dhangwatnotai, P., Roughgarden, T., and Yan, Q. (2015). Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333.

Dobzinski, S., Lavi, R., and Nisan, N. (2008). Multi-unit auctions with budget limits. In *Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on*, pages 260–269.

Dobzinski, S. and Paes Leme, R. (2014). Efficiency guarantees in auctions with budgets. In *International Colloquium on Automata, Languages, and Programming*, pages 392–404.

Dütting, P. and Kesselheim, T. (2015). Algorithms against anarchy: Understanding non-truthful mechanisms. In *Proceedings of the 16th ACM Conference on Economics and Computation*, pages 239–255.

Fu, H., Hartline, J., and Hoy, D. (2013). Prior-independent auctions for risk-averse agents. In *Proceedings of the 14th ACM conference on Electronic commerce*, pages 471–488.

Fu, H., Immorlica, N., Lucier, B., and Strack, P. (2015). Randomization beats second price as a prior-independent auction. In *Proceedings of the 16th ACM Conference on Economics and Computation*, pages 323–323.

Gershkov, A., Goeree, J. K., Kushnir, A., Moldovanu, B., and Shi, X. (2013). On the equivalence of bayesian and dominant strategy implementation. *Econometrica*, 81(1):197–220.

Goel, G., Mirrokni, V., and Leme, R. P. (2015). Polyhedral clinching auctions and the adwords polytope. *Journal of the ACM*, 62(3):18.

Hartline, J., Hoy, D., and Taggart, S. (2014). Price of anarchy for auction revenue. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 693–710.

Hartline, J. and Roughgarden, T. (2008). Optimal mechanism design and money burning. In *Proc. 39th ACM Symp. on Theory of Computing*.

Hartline, J. and Taggart, S. (2016). Non-revelation mechanism design. *arXiv preprint arXiv:1608.01875*.

Hartline, J. D. (2016). Mechanism design and approximation. Under preparation.

Jackson, M. (2001). A crash course in implementation theory. *Social Choice and Welfare*, 18:655–708.

Laffont, J.-J. and Robert, J. (1996). Optimal auction with financially constrained buyers. *Economics Letters*, 52(2):181–186.

Maskin, E. (2000). Auctions, development, and privatization: Efficient auctions with liquidity-constrained buyers. *European Economic Review*, 44(4–6):667–681.

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

Niazadeh, R., Yuan, Y., and Kleinberg, R. (2014). Simple and near-optimal mechanisms for market intermediation. In *International Conference on Web and Internet Economics*, pages 386–399.

Pai, M. M. and Vohra, R. (2014). Optimal auctions with financially constrained buyers. *Journal of Economic Theory*, 150:383–425.

Roughgarden, T., Syrgkanis, V., and Tardos, E. (2017). The price of anarchy in auctions. *Journal of Artificial Intelligence Research*, 59:59–101.

Roughgarden, T., Talgam-Cohen, I., and Yan, Q. (2012). Supply-limiting mechanisms. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 844–861.

Saltzer, J. H., Reed, D. P., and Clark, D. D. (1984). End-to-end arguments in system design. *ACM Transactions on Computer Systems (TOCS)*, 2(4):277–288.

Syrgkanis, V. and Tardos, E. (2013). Composable and efficient mechanisms. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 211–220.

Yao, A. (2017). Dominant-strategy versus bayesian multi-item auctions: Maximum revenue determination and comparison. In *Proceedings of the 18th ACM Conference on Economics and Computation*, pages 3–20.