Penalty Bidding Mechanisms for Allocating Resources and Overcoming Present Bias*

Hongyao Ma
Harvard University
Cambridge, MA, US
hma@seas.harvard.edu

Reshef Meir
Technion - Israel Institute of Technology
Haifa, Israel
reshefm@ie.technion.ac.il

David C. Parkes
Harvard University
Cambridge, MA, US
parkes@eecs.harvard.edu

Elena Wu-Yan
Facebook, Inc.
Menlo Park, CA, US
elenawy@fb.com

ABSTRACT

From skipped exercise classes to last-minute cancellation of dentist appointments, underutilization of reserved resources abounds. Likely reasons include uncertainty about the future, further exacerbated by present bias. In this paper, we unite resource allocation and commitment devices through the design of contingent payment mechanisms, and propose the *two-bid penalty- bidding mechanism*. This extends an earlier mechanism proposed by Ma et al. [21], assigning the resources based on willingness to accept a no-show penalty, while also allowing each participant to increase her own penalty in order to counter present bias. We establish a simple dominant strategy equilibrium, regardless of an agent's level of present bias or degree of "sophistication". Via simulations, we show that the proposed mechanism substantially improves utilization and achieves higher welfare and better equity in comparison with mechanisms used in practice and mechanisms that optimize welfare in the absence of present bias.

KEYWORDS

mechanism design, contingent payments, present bias

ACM Reference Format:
Hongyao Ma, Reshef Meir, David C. Parkes, and Elena Wu-Yan. 2020. Penalty Bidding Mechanisms for Allocating Resources and Overcoming Present Bias. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 9 pages.

1 INTRODUCTION

"It was a disaster," recalled Matt Juszczak, co-founder of Turnstyle Cycle and Bootcamp, a fitness company that offers cycling and bootcamp classes across five studios in the Boston area. "When we opened our first indoor cycling location in Boston's Back Bay, we saw 40 to 50 no-shows and late cancels in an average day— that's over 15,000 in a year!"[1] Like many well-known exercise studios, Turnstyle allowed customers to reserve class spots several days in advance with a first-come-first-serve reservation system. However, ambitious customers, overestimating the amount of time in their schedules or their desire to exercise in the future, often snag a spot only to ultimately cancel last-minute or simply not show up.

Similarly, at the gymnasium of the university to which some authors of this paper are affiliated, the squash courts used to allow members to reserve time-slots to play squash up to seven days ahead of time. Now three days, the gym operators still warn upon a reservation that "failure to cancel or show-up for your reservation may lead to loss of reservation privileges." For other examples, organizers of free events report to Eventbrite that their no-show rate can be as high as 50%,[2] and even for prepaid events organized through Doorkeeper, the fraction of no-shows can be 20%.[3] Studies of outpatient clinics report that no-shows can range from 23-34%, with no-shows costing an estimated 14% of daily revenue as well as impacting efficiency [24].

Common to all these examples is the presence of uncertainty, self-interest, and down-stream decisions by participants, together with the interest of the planner (gym manager, event organizer, health clinic) in a resource being used and not wasted. Beyond revenue and efficiency motivations, utilization can have positive externality in and of itself, cycling studio members gaining motivation from fellow bikers, for example. Complicating the problem is *present bias*, often phrased as the constant struggle between our current and future selves [20, 26]. It is easy to imagine that at the beginning of the week, someone might prefer a spin class over watching TV on Friday, reserving a spot, but by the time Friday comes around preferring to just watch TV.

Recognizing the problem of low utilization, many reservation systems charge a penalty for no-show. Turnstyle has started to charge a $20 penalty for missing a class,[4] patients who miss appointments at hospitals may need to pay a fee,[5] and organizers of some conferences collect a deposit that is returned only to students who actually attended talks.[6] These approaches can be viewed as

*The full version of this paper is available at https://arxiv.org/abs/1906.09713.
[1] https://business.mindbody.io/education/blog/tips-reduce-no-shows-and-late-cancels-your-fitness-business, visited September 1, 2018.
[2] https://www.eventbrite.com/blog/asset/ultimate-way-reduce-no-shows-free-events/, visited May 6, 2019.
[3] https://www.doorkeeper.jp/event-planning/increasing-participants-decreasing-no-shows?locale=en, visited May 6, 2019.
[4] https://kb.turnstylecycle.com/policies/what-is-the-late-cancel-no-show-policy, visited May 6, 2019.
[5] https://huhs.harvard.edu/sites/default/files/HUHS%20Missed%20Appointment%20Fee%20Appeal%20Form.pdf, visited May 10th, 2018.
[6] https://risingstarsasia.org/guidelines.php, visited May 10th, 2019.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

ad-hoc, first-come-first serve schemes, for some choice of no-show penalty: a penalty that is too small is not effective, whereas a penalty that is too big will drive away participation in the scheme.

In recent work, Ma et al. [21] model participants' future value from using a resource as a random variable, and propose the *contingent second price mechanism* (CSP). The mechanism elicits from each participant a bid on the highest no-show penalty she is willing accept, assigns the resources to the highest bidders, and charges the highest losing bid as the penalty. The bids provide a good signal for participants' reliability, and the CSP mechanism provably optimizes utilization in dominant strategy equilibrium among a large family of mechanisms. With present bias, however, charging the highest losing bid as penalty no longer guarantees truthfulness: a rational participant always prefers smaller penalties, but a present-biased participant may favor larger penalties when a stronger commitment device is more effective in overcoming myopia.

1.1 Our results

In this paper, we unite through contingent payment mechanisms the allocation of scarce resources under uncertainty, and the design of commitment devices— techniques that aim to overcome present bias and to fulfill a plan for desired future behavior.

We generalize the model proposed in Ma et al. [21], decomposing an agent's value for a resource into the *immediate value* and the *future value*. The immediate value is a random variable, and the value experienced at the time of using the resource (modeling for example the opportunity cost and present pain of going to the gym). The future value is not gained until some future time (consider, for example the future benefit from better health). We incorporate the standard quasi-hyperbolic discounting model for time-inconsistent preferences [20, 26], such that when an agent is making a decision on whether to use a resource, the future value is discounted by a present bias factor. Agents may also have different levels of sophistication in regard to their level of self-awareness, modeled by agents' belief on their own present bias factor— a *naive agent* believes she does not discount the future, a *sophisticated agent* knows her bias factor precisely and is able to perfectly forecast her future actions, and a *partially naive agent* resides somewhere in between [26, 27].

In period 0, an agent's private information is the distribution of the immediate value, the (fixed) future value, and what she believes to be her present bias factor. A mechanism elicits information from each agent, assigns each of $m \geq 1$ resources, and may determine both a *base payment* that an assigned agent always pays, as well as a *penalty* for each assigned agent in the event of a no-show. In period 1, each assigned agent learns her immediate value, and with knowledge of the penalty and future value, decides (under the influence of present bias) whether or not to use the resource.

The *two-bid penalty-bidding mechanism* (2BPB) works as follows. In period 0, the mechanism elicits a bid from each agent, representing the highest penalty she is willing to accept for no show, and assigns the resources to the m highest bidders. To address the non-monotonicity of agent's expected utility in the penalty, the mechanism asks each assigned agent to report a penalty weakly higher than the $m + 1^{\text{th}}$ bid, representing the actual amount she would like to be charged in the case of a no-show (thereby operating also as a commitment device).

Given the option to choose an optimal level of commitment weakly above the highest losing bid, it is a dominant strategy under the 2BPB mechanism for each agent to bid her maximum acceptable no-show penalty, regardless of her immediate value distribution, future value, level of present bias, or degree of sophistication (Theorem 3.5). While naive agents do not see the value of commitment and generally do not take any commitment device when offered [5, 6], the 2BPB mechanism is still able to help reducing the loss of welfare and utilization due to no show, since a commitment device is designed through the mechanism, and is an integral part of the system. We also prove that the mechanism satisfies voluntary participation, and runs without a budget deficit.

We show via simulation that the two-bid penalty bidding mechanism not only improves utilization, but also achieves higher social welfare than the standard $m + 1^{\text{th}}$ price auction, which is welfare-optimal for settings without present bias. The mechanism also outperforms a family of mechanisms widely used in practice, which assign resources first-come-first-served and charge a fixed no-show penalty. Moreover, in a population where agents have different levels of present bias, the more biased agents benefit more than the less biased agents under two-bid penalty bidding. This results in better equity compared with the outcome under the $m + 1^{\text{th}}$ price auction, where the most biased agents gain little or no welfare.

1.2 Related Work

To the best of our knowledge, this current paper is the first to study resource assignment in the presence of uncertainty and present bias. The closest related work is on the design of mechanisms to improve resource utilization where agents have uncertain future values [21–23]. The proposed mechanisms, however, no longer have dominant strategy equilibrium for present-biased agents. This present work builds on Ma et al. [21], generalizing the model to incorporate present bias, and makes use of two-bid penalty bidding to align incentives. Crucially, the mechanism does not need any knowledge about agents' level of bias or value distributions.

Contingent payments have arisen in the past in the context of oil drilling license auctions [16], royalties [7, 10], ad auctions [29], and selling a firm [12]. Payments that are contingent on some observable world state also play the role of improving revenue as well as hedging risk [28]. In our model, in contrast, payments are contingent on agents' own downstream decisions and serve the role of commitment devices. In regard to auctions in which actions take place after the time of contracting, Atakan and Ekmekci [2] study auctions where the value of taking each action depends on the collective actions by others, but these actions are taken before rather than after observing the world state. Courty and Li [9] study the problem of revenue maximization in selling airline tickets, where passengers have uncertainty about their value for a trip, and may decide not to take a trip after realizing their actual values. The type space considered there is effectively one-dimensional, and present bias is not considered.

Laibson [20] introduced the quasi-hyperbolic discounting for modeling time-inconsistent decision making, where in addition to exponential discounting, all future utilities are discounted by an additional present bias factor. O'Donoghue and Rabin [26] classify present-biased agents into naive agents (unaware of present bias) and sophisticated agents (fully aware), and find that naive agent procrastinate immediate-cost activities and do immediate-reward activities too soon, while sophistication lessens procrastination but intensifies the doing-too-soon. O'Donoghue and Rabin [27] also study how the role of choice affects procrastination, and introduce the idea of a partially naive agent, who is aware of present bias but underestimates the degree of this bias. Researchers have also attempted to estimate the present bias factor in the real world, however, there has not been consensus about this [3, 8, 13] .

Researchers have also examined various kinds of *commitment devices* to mitigate present bias. Giné et al. [14], for example, offer smokers a savings account that forfeits deposits to a charity if the they fail a urine test for nicotine. By bundling a "want" activity (listening to one's favorite audio book) with a "should" activity (going to the gym), Milkman et al. [25] evaluate the effectiveness of temptation bundling as a commitment device to tackle two self-control problems at a time. See also Laibson [20] and Beshears et al. [4]. In a different setting, Kleinberg and Oren [18] consider how to modify the sequencing of tasks available to individuals in order to help a present-biased agent adopt a more optimal sequence of tasks. This work is later extended to consider sophisticated agents, the interaction between present bias and sunk-cost bias, and agents whose present bias factors are uncertain [15, 17, 19]. There are no uncertain values or costs in these models, and no contention for limited resources.

2 PRELIMINARIES

We first introduce the model for the assignment of m homogeneous resources, leaving a discussion of the generalization to heterogeneous resources to Section 3.1. There is a set of agents $N = \{1, \ldots, n\}$ and three time periods. In period 0, when resources need to be assigned, the value of each agent $i \in N$ for using a resource is uncertain, represented by $V_i = V_i^{(1)} + v_i^{(2)}$. The period 1 *immediate value*, $V_i^{(1)}$ is a random variable with cumulative distribution function F_i, whose exact (and potentially negative) value is not realized until period 1. This models, for example, the opportunity cost and present pain of going to the gym. The period 2 *future value* $v_i^{(2)} \geq 0$ models the expected future benefit for agent i (e.g. the future benefit from better health), if she uses a resource in period 1.

Agents are present-biased, such that at any point of time, agent i discounts her utility from all future periods by a factor of $\beta_i \in [0, 1]$. Agents may not be fully aware of this bias, however, and agent i believes that when making decisions, she will discount her future utility by a factor of $\tilde{\beta}_i \in [\beta_i, 1]$. An agent with $\tilde{\beta}_i = \beta_1 = 1$ is *rational* and does not discount her future utility. An agent with $\tilde{\beta}_i = \beta_i < 1$ is said to be *sophisticated*, and fully aware of the degree of her present bias. An agent with $\beta_i < 1$ and $\tilde{\beta}_i = 1$ is said to be *naive*, believing that she will make rational decisions in the future, and an agent with $\tilde{\beta}_i \in (\beta_i, 1)$ is said to be *partially naive*.

Let $\theta_i = (F_i, v_i^{(2)}, \beta_i, \tilde{\beta}_i)$ denote agent i's *type*, and $\theta = (\theta_1, \ldots, \theta_n)$ denote a type profile. The tuple $\tilde{\theta}_i = (F_i, v_i^{(2)}, \tilde{\beta}_i)$ is agent i's private information at period 0, when the assignment of resources is determined. Each allocated agent privately learns the realization $v_i^{(1)}$ of $V_i^{(1)}$ and then decides whether to use the resource at period 1. Define $V_i^+ \triangleq \max\{V_i, 0\}$. Following Ma et al. [21], we make the following assumptions about F_i for each $i \in N$:

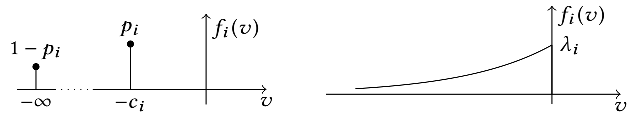

(a) The (c_i, p_i) model. **(b) The exponential model.**

Figure 1: Distributions of $V_i^{(1)}$ under the (c_i, p_i) and the exponential type models.

(A1) $\mathbb{E}\left[V_i^+\right] > 0$, which means that a rational agent gets positive value from using the resource with non-zero probability, thus the *option* to use the resource has positive value.

(A2) $\mathbb{E}\left[V_i^+\right] < +\infty$, meaning that agents do not get infinite expected utility from the option to use the resource, thus would not be willing to pay an unboundedly large payment for it.

(A3) $\mathbb{E}[V_i] < 0$, meaning that being forced to always use the resource regardless of what happens is not favorable, so that no agent would accept any unboundedly large no-show penalty for the right to use a resource.[7]

We now provide a few examples of models for agent types.

Example 2.1 ((c_i, p_i) model). The future value for agent i for using a resource is $v_i^{(2)} = w_i > 0$, however, she is able to do so only with probability $p_i \in (0, 1)$, and at a period 1 opportunity cost modeled by $V_i^{(1)} = -c_i$. With probability $1 - p_i$, agent i is unable to show up to use the resource. This hard constraint can be modeled as $V_i^{(1)}$ taking value $-\infty$ with probability $1 - p_i$. See Figure 1a. We have $\mathbb{E}\left[V_i^+\right] = (w_i - c_i)p_i > 0$, and $\mathbb{E}[V_i] = -\infty$ thus assumptions (A1)-(A3) are satisfied.

Example 2.2 (Exponential model). The opportunity cost for an agent to use the resource in period one is an exponentially distributed random variable with parameter λ_i, (i.e. $-V_i^{(1)} \sim \text{Exp}(\lambda_i)$). If the agent used a resource, she gains a future utility of $v_i^{(2)} = w_i > 0$. See Figure 1b. The expectation of $\mathbb{E}\left[V_i^{(1)}\right]$ is λ_i^{-1}, thus $\mathbb{E}[V_i] = \lambda_i^{-1} + w_i$ and (A1)-(A3) are satisfied when $w_i < \lambda_i^{-1}$.

2.1 Two-Period Mechanisms

We consider *two-period mechanisms*, denoted as $\mathcal{M} = (\mathcal{R}, x, s, t)$, and following the timeline suggested by Ma et al. [21]. The mechanisms can, in general, involve both a base payment that an agent will pay irrespective of her utilization decision as well as a penalty for no show. The mechanisms are defined for a general message space $\mathcal{R}$ for reports, and with allocation rule x, and with each agent i facing a base payment $s_i(r)$ and a penalty $t_i(r)$.

The timeline for a two-period mechanism is as follows:

Period 0:

- Each agent i reports $r_i \in \mathcal{R}$ to the mechanism based on knowledge of $\tilde{\theta}_i$.

[7] Regardless of the degree of present bias or sophistication, an agent for which (A3) is violated is willing to accept a 1 billion dollar no-show penalty, (almost) always use the resource, and get a non-negative utility in expectation.

- The mechanism allocates the right to use the resources to a subset of agents, $A \subseteq N$, with $|A| \leq m$, thus $x_i(r) = 1$ for all $i \in A$ and $x_i(r) = 0$ for all $i \notin A$.
- For each agent $i \in N$, the mechanism determines a base payment $s_i(r)$ that the agent will pay for sure. For each assigned agent $i \in A$, the mechanism determines an additional penalty $t_i(r)$ that will be charged for a no show.

Period 1:

- The mechanism collects base payment $s_i(r)$ from each agent.
- Each allocated agent $i \in A$ privately observes the realized immediate value $v_i^{(1)}$, and decides whether to use the resource based on $v_i^{(1)}$, the future value $v_i^{(2)}$, and the no-show penalty $t_i(r)$.
- The mechanism collects penalty $t_i(r)$ from any agent $i \in A$ who is a no show.

Example 2.3 ($(m+1)^{\text{th}}$ price auction). The standard $m + 1^{\text{th}}$ price auction for assigning m resources can be described as a two-period mechanism, where the report space is $\mathcal{R} = \mathbb{R}$. Ordering agents in decreasing order of their reports, s.t. $r_1 \geq r_2 \geq \ldots r_n$, the allocation rule is $x_i(r) = 1$ for all $i \leq m$, $x_i(r) = 0$ for $i > m$. Each allocated agent is charged $s_i(r) = r_{m+1}$, and all other payments are zero. The $m + 1^{\text{th}}$ price auction does not make use of any penalties.

Example 2.4 (Generalized contingent second price mechanism). The *generalized contingent second price* (GCSP) mechanism [21] for assigning m homogeneous resources collects a single bid from each agent, allocates the right to use resource to the m highest bidders, and charges the $m + 1^{\text{th}}$ highest bid, *but only if an allocated agent fails to use the resource.* Formally, $\mathcal{R} = \mathbb{R}$. Ordering the agents s.t. $r_1 \geq r_2 \geq \ldots r_n$ (breaking ties randomly), we have $x_i(r) = 1$ for $i \leq m$, $x_i(r) = 0$ for $i > m$, $t_i(r) = \max_{i' \notin A} r_{i'}$, and all other payments are 0.

We assume risk-neutral, expected-utility maximizing agents, but with quasi-hyperbolic discounting for future utilities. Each assigned agent i faces a *two part payment* (z, y), where z is the *penalty* the agent pays in period 1 in the case of no-show, and y is the *base payment* the agent always pays in period 1. When period 1 arrives and the agent learns the realized immediate value $v_i^{(1)}$, she discounts the future by β_i, and makes decisions as if she will gain utility $v_i^{(1)} - y + \beta_i v_i^{(2)}$ from using the resource, and $-y - z$ from not using the resource. Based on this, the agent uses the resource if and only if

$$v_i^{(1)} - y + \beta_i v_i^{(2)} \geq -y - z \Leftrightarrow v_i^{(1)} \geq -z - \beta_i v_i^{(2)}, \tag{1}$$

breaking ties in favor of using the resource. Let $\mathbb{1}\{\cdot\}$ be the indicator function, and define $u_i(z)$, the expected utility of the agent when facing penalty z, as

$$\begin{aligned} u_i(z) \triangleq & \mathbb{E}\left[(V_i^{(1)} + v_i^{(2)})\mathbb{1}\left\{V_i^{(1)} \geq -z - \beta_i v_i^{(2)}\right\}\right] \\ & - z\mathbb{P}\left[V_i^{(1)} < -z - \beta_i v_i^{(2)}\right]. \end{aligned} \tag{2}$$

The *actual* expected utility of an allocated agent facing a two-part payment (z, y) is $u_i(z) - y$. Under a two-period mechanism $\mathcal{M}$, given report profile r, agent i's expected utility is $x_i(r)u_i(t_i(r)) - s_i(r)$.

An agent believes that she will make decisions as if she has present-bias factor $\tilde{\beta}_i$, and will decide to use the resource in period 1 if and only if

$$v_i^{(1)} \geq -z - \tilde{\beta}_i v_i^{(2)}. \tag{3}$$

Therefore, an agent's *subjective expected utility* given penalty z is

$$\begin{aligned} \tilde{u}_i(z) \triangleq & \mathbb{E}\left[(V_i^{(1)} + v_i^{(2)})\mathbb{1}\left\{V_i^{(1)} \geq -z - \tilde{\beta}_i v_i^{(2)}\right\}\right] \\ & - z\mathbb{P}\left[V_i^{(1)} < -z - \tilde{\beta}_i v_i^{(2)}\right]. \end{aligned} \tag{4}$$

We call $\tilde{u}_i(z)$ the *subjective expected utility function.* For sophisticated agents who are able to perfectly predict their future decisions (i.e. $\tilde{\beta}_i = \beta_i$), $\tilde{u}_i(z)$ and $u_i(z)$ coincide.

We assume that if allocated, agents' decisions in period 1 are influenced by their present bias, but are otherwise rational. The interesting question is to study an agent's incentives regarding reports in period 0, which are made based on subjective expected utility $\tilde{u}_i(z) - y$. For any vector $g = (g_1, \ldots, g_n)$, we denote $g_{-i} \triangleq (g_1, \ldots, g_{i-1}, g_{i+1}, \ldots, g_n)$.

Definition 2.5 (Dominant strategy equilibrium). A two-period mechanism has a *dominant strategy equilibrium* (DSE) if for each agent $i \in N$, for any type θ_i satisfying (A1)-(A3), there exists a report $r_i^* \in \mathcal{R}$ s.t. $\forall r_i \in \mathcal{R}$, $\forall r_{-i} \in \mathcal{R}^{n-1}$, $x_i(r_i^*, r_{-i})\tilde{u}_i(t_i(r_i^*, r_{-i})) - s_i(r_i^*, r_{-i}) \geq x_i(r_i, r_{-i})\tilde{u}_i(t_i(r_i, r_{-i})) - s_i(r_i, r_{-i})$.

Let $r^*(\theta) = (r_1^*, \ldots, r_n^*)$ denote a report profile under a DSE given type profile θ.

Definition 2.6 (Voluntary participation). A two-period mechanism satisfies *voluntary participation* (VP) if for each agent $i \in N$, for any type θ_i satisfying (A1)-(A3), and any report profile $r_{-i} \in \mathcal{R}^{n-1}$, $x_i(r_i^*, r_{-i})\tilde{u}_i(t_i(r_i^*, r_{-i})) - s_i(r_i^*, r_{-i}) \geq 0$.

VP requires that each agent has non-negative subjective expected utility under her dominant strategy, given that she makes present-biased but otherwise rational decisions in period 1 if allocated, regardless of the reports made by the rest of the agents.

The expected revenue of a two-period mechanism $\mathcal{M}$ is the total expected payment made by the agents in the DSE, assuming present-biased but otherwise rational decisions in period 1:

$$rev_{\mathcal{M}}(\theta) \triangleq \sum_{i \in N} \left(s_i(r^*) + x_i(r)t_i(r^*)\mathbb{P}\left[V_i^{(1)} < -t_i(r^*) - \beta_i v_i^{(2)}\right]\right).$$

Definition 2.7 (No deficit). A two-period mechanism satisfies *no deficit* (ND) if, for any type profile θ that satisfies (A1)-(A3), the expected revenue is non-negative: $rev_{\mathcal{M}}(\theta) \geq 0$.

The *utilization* achieved by mechanism $\mathcal{M}$ is the expected number of resources used by the assigned agents in the DSE:

$$ut_{\mathcal{M}}(\theta) \triangleq \sum_{i \in N} x_i(r^*)\mathbb{P}\left[V_i^{(1)} \geq -t_i(r^*) - \beta_i v_i^{(2)}\right].$$

The expected *social welfare* achieved by mechanism $\mathcal{M}$ is the total expected value derived by agents from using the resources:

$$sw_{\mathcal{M}}(\theta) \triangleq \sum_{i \in N} x_i(r^*)\mathbb{E}\left[(V_i^{(1)} + v_i^{(2)})\mathbb{1}\left\{V_i^{(1)} \geq -t_i(r^*) - \beta_i v_i^{(2)}\right\}\right].$$

Our objective is to design mechanisms that maximize expected social welfare. We do not consider monetary transfers as part of the social welfare function. The reason $t_i(r^*)$ appears in $sw_{\mathcal{M}}(\theta)$is that it affects decisions of the allocated agents in period 1.

3 THE TWO-BID PENALTY BIDDING MECHANISM

In this section, we introduce the two-bid penalty bidding mechanism, and prove that agents have simple dominant strategies, regardless of their value distributions, levels of present bias, or degrees of sophistication.

Definition 3.1 (Two-bid penalty bidding mechanism (2BPB)). The *two-bid penalty bidding mechanism* collects bids $\bar{b} = (\bar{b}_i, \dots, \bar{b}_n)$ from agents in period 0, and reorders agents s.t. $\bar{b}_1 \geq \bar{b}_2 \geq \cdots \geq \bar{b}_n$ (breaking ties randomly).

- Allocation rule: $x_i(b) = 1$ for $i \leq m$, $x_i(b) = 0$ for $i > m$.
- Payment rule: the mechanism announces $\bar{b}_{m+1}$, elicits a second bid $\underline{b}_i \geq \bar{b}_{m+1}$ from each assigned agent $i \leq m$, and sets $t_i(b) = \underline{b}_i$. $t_i(b) = 0$ for all $i > m$, and $s_i(b) = 0$ for all $i \in N$.

The 2BPB mechanism first asks agents to bid on the maximum penalties they are willing to accept for the option to use the resource for free, and assigns the resources to the highest bidders. The mechanism then asks each assigned agent to bid a penalty that is weakly higher than the $m + 1^{\text{th}}$ bid, representing the amount she would like to be charged in case of a no-show.[8]

To establish the dominant strategy equilibrium under the mechanism, we first prove some useful properties of $\tilde{u}_i(z)$.

Lemma 3.2. *Given an agent with any type θ_i that satisfies (A1)-(A3), the agent's subjective expected utility $\tilde{u}_i(z)$ as a function of the penalty z satisfies:*

(i) $\tilde{u}_i(0) \geq 0$, and $\lim_{z\to+\infty} \tilde{u}_i(z) \leq \mathbb{E}[V_i]$.
(ii) $\tilde{u}_i(z)$ is right continuous and upper semi-continuous.

Proof. We first prove part (i). $\tilde{u}_i(0) \geq 0$ holds given (4) and the fact that $\tilde{\beta}_i \leq 1$ and $w_i \geq 0$. For the limit as $z \to +\infty$, observe that:

$$\begin{aligned}
&\tilde{u}_i(z) \\
=&\mathbb{E}\left[(V_i^{(1)} + \tilde{\beta}_i v_i^{(2)})\mathbb{1}\left\{(V_i^{(1)} + \tilde{\beta}_i v_i^{(2)}) \geq -z\right\}\right] + \\
&(1-\tilde{\beta}_i)v_i^{(2)}\mathbb{P}\left[(V_i^{(1)} + \tilde{\beta}_i v_i^{(2)}) \geq -z\right] - z\mathbb{P}\left[(V_i^{(1)} + \tilde{\beta}_i v_i^{(2)}) < -z\right] \\
=&\mathbb{E}\left[\max\{V_i^{(1)} + \tilde{\beta}_i v_i^{(2)}, -z\}\right] + (1-\tilde{\beta}_i)v_i^{(2)}\mathbb{P}\left[(V_i^{(1)} + \tilde{\beta}_i v_i^{(2)}) \geq -z\right].
\end{aligned} \tag{5}$$

By the monotone convergence theorem, as $z \to +\infty$, the first term of (5) converges to $\mathbb{E}\left[V_i^{(1)} + \tilde{\beta}_i v_i^{(2)}\right] = \mathbb{E}\left[V_i^{(1)}\right] + \tilde{\beta}_i v_i^{(2)}$. The second term is upper bounded by $(1-\tilde{\beta}_i)v_i^{(2)}$, therefore we get $\lim_{z\to+\infty} \tilde{u}_i(z) \leq \mathbb{E}\left[V_i^{(1)}\right] + v_i^{(2)} = \mathbb{E}[V_i]$.

For (ii), $\max\{V_i^{(1)} + \tilde{\beta}_i v_i^{(2)}, -z\}$ is continuous in z, therefore its expectation $\mathbb{E}\left[\max\{V_i^{(1)} + \tilde{\beta}_i v_i^{(2)}, -z\}\right]$ is also continuous in z. $\mathbb{P}\left[(V_i^{(1)} + \tilde{\beta}_i v_i^{(2)}) \geq -z\right]$ is right continuous, implying the right continuity of $\tilde{u}_i(z)$. The upper semi-continuity (i.e. $\lim_{z\uparrow z^*} \tilde{u}_i(z) \leq \tilde{u}_i(z^*)$ for all $z^* \geq 0$) holds because of the fact that $(1-\tilde{\beta}_i)v_i^{(2)} \geq 0$, and that $\mathbb{P}\left[(V_i^{(1)} + \tilde{\beta}_i v_i^{(2)}) \geq -z\right]$ is upper semi-continuous. □

[8]Instead of using two rounds of bidding, we may also consider a direct revelation mechanism, where agents report their private information $\tilde{\theta}_i$, with which the mechanism determines the assignment and the contingent payments.

Ma et al. [21] had earlier proved that for a rational agent without present bias, her expected utility as a function of the penalty is continuous, convex, and monotonically decreasing. These properties no longer hold for present-biased agents, since a higher penalty may incentivize an agent to use the resource more optimally, resulting in a higher expected utility.

For any penalty z, we define $\tilde{U}_i(z)$ as agent i's highest subjective expected utility for the best choice of penalty, assuming this penalty must be at least z:

$$\tilde{U}_i(z) = \sup_{z' \geq z} \tilde{u}_i(z'). \tag{6}$$

The following lemma proves the continuity and monotonicity of $\tilde{U}_i(z)$, together with the existence of a zero-crossing for $\tilde{U}_i(z)$. This zero-crossing point is the maximum penalty an agent will accept, in the case that this agent can choose to be charged any penalty weakly larger than this penalty.

Lemma 3.3. *Given any agent with type θ_i that satisfies (A1)-(A3), the agent's subjective expected utility $\tilde{U}_i(z)$ as a function of the minimum penalty z satisfies:*

(i) $\tilde{U}_i(z)$ is continuous and monotonically decreasing in z.
(ii) There exists a zero-crossing z_i^0 s.t. $\tilde{U}_i(z_i^0) = 0$ and $\tilde{U}_i(z) < 0$ for all $z > z_i^0$.

Proof. For part (i), the monotonicity of $\tilde{U}_i(z)$ is obvious, and the continuity of $\tilde{U}_i(z)$ is implied by the right continuity of $\tilde{u}_i(z)$ as shown in Lemma 3.2.

For part (ii), Lemma 3.2 and assumption (A3) imply $\lim_{z\to\infty} \tilde{u}_i(z) \leq \mathbb{E}[V_i] < 0$. Therefore, there exists $Z \in \mathbb{R}$ s.t. $\tilde{u}_i(z) < 0$ for all $z \geq Z$. As a result, $\tilde{U}_i(z) < 0$ holds for all $z \geq Z$, and the monotonicity and continuity of $\tilde{U}_i(z)$ then imply that the following supreme exists:

$$z_i^0 \triangleq \sup\{z \in \mathbb{R} \mid \tilde{U}_i(z) \geq 0\},$$

and that we must have $\tilde{U}_i(z_i^0) = 0$ and $\tilde{U}_i(z) < 0$ for all $z > z_i^0$. □

The following example illustrates the expected utility functions of an agent with (c_i, p_i) type (see Example 2.1), and shows that there may not exist a DSE under the CSP mechanism.

Example 3.4 (Expected utility of (c_i, p_i) type agents). Consider a sophisticated agent whose type follow the (c_i, p_i) model, who is assigned a resource and charged no-show penalty z. With probability $1 - p_i$, the agent is not able to use the resource, and has to pay the penalty. With probability p_i, the agent is able to use the resource at a cost of c_i, but will use the resource if and only if $\beta_i w_i - c_i \geq -z \Leftrightarrow z \geq c_i - \beta_i w_i$. Therefore, the agent's expected utility as a function of the no-show penalty is of the form:

$$u_i(z) = \begin{cases} -z, & \text{if } z < c_i - \beta_i w_i, \\ (w_i - c_i)p_i - (1-p_i)z, & \text{if } z \geq c_i - \beta_i w_i, \end{cases}$$

and $\tilde{u}_i(z) = u_i(z)$ holds since the agent is sophisticated. Figure 2a illustrates $\tilde{u}_i(z)$ for an agent with $c_i - \beta_i w_i > 0$. Intuitively, $c_i - \beta_i w_i$ is the minimum penalty the agent needs to be charged so that she will use the resource when she is able to. When $z < c_i - \beta_i w_i$, the agent ends up always paying the penalty, which is too small to incentivize utilization. $\tilde{U}_i(z)$ of this agent is as shown in Figure 2b. The maximum penalty the agent is willing to accept is $z_i^0 = (w_i - c_i)p_i/(1-p_i)$.

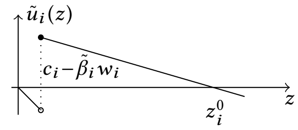

(a) $\tilde{u}_i(z)$.

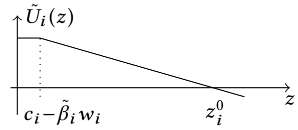

(b) $\tilde{U}_i(z)$.

Figure 2: Subjective expected utility functions of a sophisticated agent with (c_i, p_i) type, with $c_i - \tilde{\beta}_i w_i > 0$.

There is no dominant strategy for this agent under the CSP mechanism. Consider the assignment of a single resource. If the highest bid among the rest of the agents satisfies $\max_{i' \neq i} b_{i'} \in [c_i - \tilde{\beta}_i w_i, z_i^0)$, the agent gets positive utility from bidding $b_i = z_i^0$, getting allocated and charged penalty $\max_{i' \neq i} b_{i'}$. However, if $\max_{i' \neq i} b_{i'} < c_i - \tilde{\beta}_i w_i$, bidding $b_i = z_i^0$ results in negative utility—the agent will be allocated, but charged a penalty that is too small to overcome her present bias. In this case, the agent is better off bidding $b_i = 0$ and get zero utility. □

We now state and prove the main theorem of this paper.

Theorem 3.5 (Dominant strategy equilibrium). *Assuming (A1)-(A3), under the two-bid penalty bidding mechanism, it is a dominant strategy for each agent $i \in N$ to bid $\bar{b}_i^* = z_i^0$. If agent i is assigned a resource and given a minimum penalty $\underline{z}$, it is then a dominant strategy to bid $\underline{b}_i^* = \arg\max_{z \geq \underline{z}} \tilde{u}_i(z)$. Moreover, the mechanism satisfies voluntary participation and no deficit.*

Proof. We first consider an agent who is assigned a resource and asked by the mechanism to bid an amount that is at least $\underline{z}$. The right continuity of $\tilde{u}_i(z)$ (see Lemma 3.2) implies that the highest utility $\tilde{U}_i(\underline{z})$ when the agent can choose any penalty weakly higher than $\underline{z}$ is achieved at $\arg\max_{z \geq \underline{z}} \tilde{u}_i(z)$. Since whichever amount an agent bids as $\underline{b}_i$ will be the penalty she is charged by the mechanism, it is a dominant strategy to bid $\underline{b}_i^* = \arg\max_{z \geq \underline{z}} \tilde{u}_i(z)$.

Given that an assigned agent will get expected utility $\tilde{U}_i(\underline{z})$ when she is asked to choose a penalty $\underline{b}_i$ that is weakly above $\underline{z}$, $\tilde{U}_i(z)$ is effectively her expected utility function in the first round of bidding. With the monotonicity of $\tilde{U}_i(z)$ and the fact that the minimum penalty is determined by the $m + 1^{\text{th}}$ highest bid, it is standard that an agent bids in DSE the highest "minimum penalty to choose from" that she is willing to accept, which is z_i^0. □

The following example shows that the 2BPB mechanism can achieve better social welfare and utilization than the $m + 1^{\text{th}}$ price auction by assigning to a "better" agent and charging a proper penalty as the commitment device.

Example 3.6. Consider the allocation of one resource to two sophisticated agents with (c_i, p_i) types, where $c_1 = 10$, $p_1 = 0.8$, $w_1 = 16$, $\beta_1 = \tilde{\beta}_1 = 0.5$ and $c_2 = 6$, $p_2 = 0.5$, $w_2 = 10$, $\beta_2 = \tilde{\beta}_2 = 0.8$.

When $z < c_1 - \beta_1 w_1 = 2$, agent 1 never uses the resource. On the other hand, $c_2 - \beta_2 w_2 < 0$ means that agent 2 uses the resource with probability p_2 while facing any non-negative penalty. $u_i(z) = \tilde{u}_i(z)$ for $i = 1, 2$ since both agents are sophisticated, and the subjective expected utility functions of the two agents are as shown in Figure 3.

Under the second price auction, agents bid in DSE $b_{1,\text{SP}}^* = \tilde{u}_1(0) = 0$ and $b_{2,\text{SP}}^* = \tilde{u}_2(0) = (w_2 - c_2)p_2 = 2$, the values of the option to use the resource without any penalty (the free option to use the resource has no value to agent 1 since she knows that she will never show up). Agent 2 gets assigned the resource and charged no penalty, achieving social welfare $(w_2 - c_2)p_2 = 2$ and utilization $p_2 = 0.5$. Under the 2BPB mechanism, the agents bid in DSE $\bar{b}_1^* = z_1^0 = (w_1 - c_1)p_1/(1 - p_1) = 24$, and $\bar{b}_2^* = z_2^0 = (w_2 - c_2)p_2/(1 - p_2) = 4$. Agent 1 is therefore assigned and will bid $\underline{b}_1^* = 4$ when asked to choose a penalty weakly above $\bar{b}_2^* = 4$, since $\tilde{u}_1(z)$ is decreasing in z for $z \geq c_1 - \tilde{\beta}_1 w_1 = 2$. The 2BPB mechanism achieves social welfare $(w_1 - c_1)p_1 = 4.8$ and utilization $p_1 = 0.8$, both are higher than those under the second price auction. □

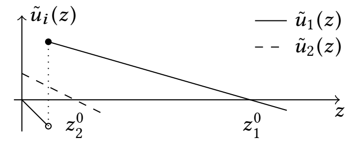

Figure 3: Expected utility of two agents in Example 3.6.

3.1 Discussion

For fully rational agents with $\tilde{\beta}_i = \beta_i = 1$, $\tilde{u}_i(z)$ is monotonically decreasing in z, and $\tilde{u}_i(z) = \tilde{U}_i(z)$ for all $z \in \mathbb{R}$. In this case, the equilibrium outcome under the 2BPB mechanism coincides with that under the $m + 1^{\text{th}}$-price generalization of the CSP mechanism.

Since $\tilde{u}_i(z)$ is what an agent considers while bidding, in period 0 a naive agent will bid as if she was rational. In period 1, however, present bias will take effect, and the naive agent may make suboptimal decisions. The actual expected utility a naive agent gets from participating in the $m+1^{\text{th}}$ price CSP or the 2BPB mechanisms, therefore, may be negative, despite the fact that she is willing to participate and believes she will get non-negative expected utility.

For two agents i and i' who are identical except that $\tilde{\beta}_i > \tilde{\beta}_{i'}$, we can prove that $\tilde{u}_i(z) \geq \tilde{u}_{i'}(z)$ holds for all $z \geq 0$. As a result, $z_i^0 \geq z_{i'}^0$. This implies that an agent who believes that she is less present-biased (i.e. with higher $\tilde{\beta}_i$) will bid higher under both the 2BPB mechanism and the $m + 1^{\text{th}}$ price auction. See Proposition B.1 in Appendix B.1 of the full version of this paper.

For rational agents without present bias, the CSP mechanism optimizes utilization among a large family of mechanisms with a set of desirable properties [21]. The 2BPB mechanism, however, does not provably optimize utilization for present-biased agents. The reason is that the actual present bias factor does not affect a naive agent's bid, and it is still possible for a very biased naive agent to be assigned but rarely show up (see Examples B.2 and B.3 in Appendix B of the full version of this paper).

The 2BPB mechanism can also be generalized for assigning multiple heterogeneous resources, where each agent $i \in N$ has a random value $V_{i,a} = V_{i,a}^{(1)} + v_{i,a}^{(2)}$ for using each resource $a \in M$. $\tilde{u}_{i,a}(z)$ and $\tilde{U}_{i,a}(z)$ can be defined similarly to (4) and (6). The 2BPB mechanism can be generalized through the use of a minimum Walrasian equilibrium price mechanism, which computes the assignment and the minimum penalty each agent faces using $\{\tilde{U}_{i,a}(z)\}_{i \in N, a \in M}$ [1, 11, 21]. As a second step, each assigned agent is asked to report a weakly higher penalty that she wants to be charged by the mechanism.

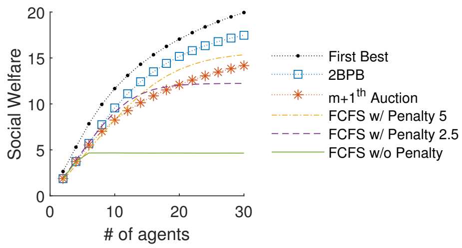

(a) Social welfare.

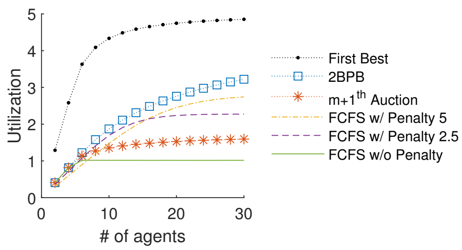

(b) Utilization.

Figure 4: Average social welfare and utilization for naive agents with exponential types.

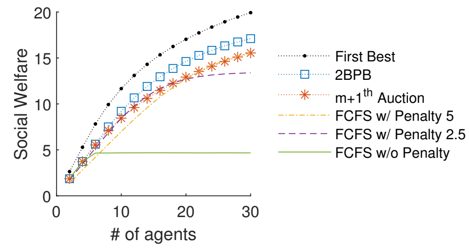

(a) Social welfare.

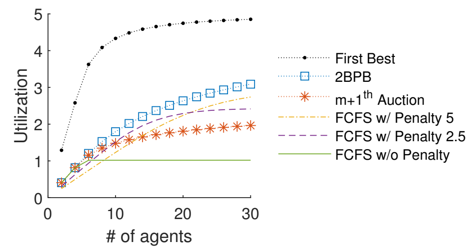

(b) Utilization.

Figure 5: Average social welfare and utilization for sophisticated agents with exponential types.

4 SIMULATION RESULTS

In this section, we adopt the exponential model (Example 2.2) and compare in simulation the social welfare and utilization achieved by different mechanisms and benchmarks. Additional simulation results assuming other type models are presented in Appendix A of the full version of this paper. For the exponential model, $\mathbb{E}[V_i] = -\lambda_i^{-1} + w_i$, where λ_i^{-1} is the expected period 1 opportunity cost for using the resource. We consider a type distribution with w_i and $\lambda_{i,a}^{-1}$ uniformly distributed as $\lambda_i^{-1} \sim \mathrm{U}[0,\ L]$ and $w_i \sim \mathrm{U}[0,\ \lambda_i^{-1}]$. $w_i < \lambda_i^{-1}$ holds almost surely thus (A1)-(A3) are satisfied. The results are not sensitive to the choices of L in defining this type distribution, and we fix $L = 20$ for the rest of this section.

4.1 Varying Resource Scarcity

Fixing the number of resources at five, we study the impact of varying the scarcity of the resource, by varying the number of agents from 2 to 30. We define the *first best* as the highest achievable welfare (or utilization) assuming full knowledge of agent types, and without violating voluntary participation or no deficit. The *first-come-first-serve with fixed penalty mechanism* (FCFS) assumes a random order of arrival, with the effect of assigning to a random subset of at most m agents who are willing to accept the penalty. We consider three levels of penalties for FCFS: 5, 2.5 and 0, where 5 is equal to the expectation of the future value w_i.

Naive Agents. For naive agents with $\tilde{\beta}_i = 1$ and present bias factor distribution $\beta_i \sim \mathrm{U}[0, 1]$, the average social welfare and utilization over 100,000 randomly generated profiles are as shown in Figure 4. When the number of agents is small, the outcomes under 2BPB, the $m + 1^{\text{th}}$ price auction, and the FCFS without penalty are similar, since all three effectively assign the resources to all agents, without charging any penalty. As the number of agents increases, the 2BPB mechanism achieves higher welfare and substantially higher utilization than the $m + 1^{\text{th}}$ price auction (which optimizes welfare for rational agents without present bias), and does this without charging any payments from agents who do show up.

The 2BPB mechanism achieves higher welfare and utilization for economies of any size, and does not require any prior knowledge about the number of agents or their bias level or value distributions. The FCFS mechanism, by comparison, requires careful adjustments of the fixed penalty level. A smaller penalty works fine when the number of agents is small but fails to keep up as the economy becomes more competitive. A larger penalty outperforms the $m+1^{\text{th}}$ price auction for larger economies, but deters participation and leaves resources unallocated when the number of agents is small.

Sophisticated Agents. Consider sophisticated agents with $\tilde{\beta}_i = \beta_i$ and present bias factor distribution $\beta_i \sim \mathrm{U}[0, 1]$. As the number of agents vary from 2 to 30, the average welfare and utilization over 100,000 randomly generated economies are as shown in Figure 5.

As with the setting with naive agents, we can see that the 2BPB mechanism achieves higher welfare and utilization than the $m + 1^{\text{th}}$ price auction, and that the performance of FCFS is very sensitive to the fixed penalty and the competitiveness of the economy. The $m + 1^{\text{th}}$ price auction achieves higher welfare and utilization for sophisticates, in comparison to the setting with naive agents. This is because sophisticated agents are able to adjust their bids depending on their present bias level, and avoid the situation where a naive agent bids too much, gets assigned, but rarely show up.

In Appendix A.1 of the full paper, we present additional simulation results assuming all agents are fully rational or partially naive. The outcome for partially naive agents is between the outcome for fully naive agents and fully sophisticated agents. For rational agents, the 2BPB mechanism achieves slightly worse welfare than the $m + 1^{\text{th}}$ price auction, which is provably optimal for this setting. The 2BPB mechanism, however, still achieves higher utilization and also a significantly better outcome than the FCFS benchmarks.

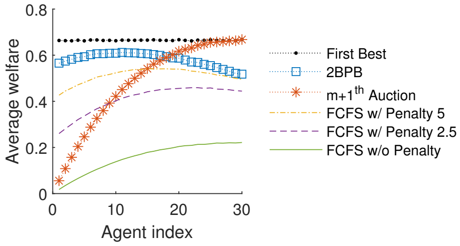

(a) Average welfare.

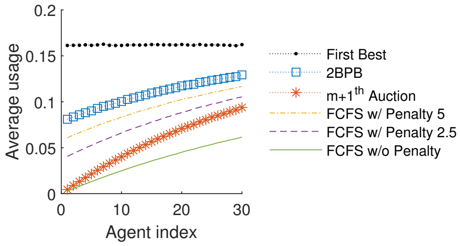

(b) Average usage.

Figure 6: Average welfare and usage for naive agents with exponential types, fixing bias factor $\beta_i = i/n$.

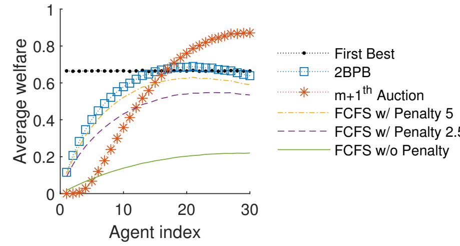

(a) Average welfare.

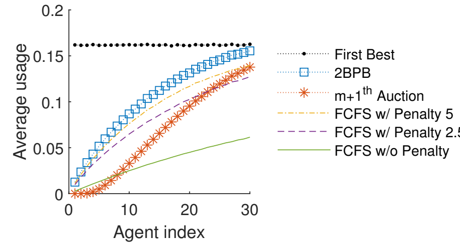

(b) Average usage.

Figure 7: Average welfare and usage for sophisticated agents with exponential types, fixing bias factor $\beta_i = i/n$.

4.2 Agents with Different Degrees of Bias

In this section, we study the different outcomes for agents with different degrees of present bias. We assume the same type distribution as in the previous section, but fix the present-bias factor of each agent i at $\beta_i = i/n$, where n is the total number of agents—the smaller an agent's index, the more present-biased an agent.

Naive Agents. We first consider the scenario where all agents are naive. Fixing $n = 30$, for 1 million randomly generated economies, the average per economy welfare and usage (i.e. the probability of being assigned and showing up) of *each agent* is as shown in Figure 6. Note that naive agents behave in period 0 as if they were rational, thus all agents bid in the same way despite their different degrees of bias, and therefore are assigned with the same probability.

Figure 6a shows that the less biased agents get substantially higher welfare than the more biased agents under the $m + 1^{\text{th}}$ price auction. By contrast, the 2BPB mechanism helps agents who are more biased to achieve substantially higher welfare than the outcome under the $m + 1^{\text{th}}$ price auction, and at the same time slightly reducing the welfare for the least biased agents. This is because the least biased agents are able to make close to optimal decisions in period 1 by themselves, and charging a penalty sometimes leads to suboptimal utilization decisions.

From Figure 6b, we see that all agents have higher average usage under the 2BPB mechanism, and agents who are more biased achieve a higher gain in comparison with the $m + 1^{\text{th}}$ price auction. Overall, the outcome under the 2BPB mechanism is subtantially more equitable for agents with all levels of bias. It is also worth noting that while naive agents do not see the value of commitment and do not take any commitment device when offered [5, 6], the 2BPB mechanism is still able to help, since a commitment device is designed through the mechanism, and it is not an option to not accept a commitment.

Sophisticated Agents. With fully sophisticated agents, the average welfare and usage are as shown in Figure 7. The first observation is that under the $m + 1^{\text{th}}$ price auction, the welfare and usage for the most biased agents are effectively zero, while the least biased agents achieve better welfare and utilization than the first-best outcome. This is because the more biased agents bid lower and get assigned with lower probability. The 2BPB mechanism is not able to achieve the same level of welfare for all agents, but achieves large improvements for the more biased population compared to the $m + 1^{\text{th}}$ price auction, and also higher welfare and better equity than the FCFS benchmarks.

5 CONCLUSION

We propose the two-bid penalty-bidding mechanism for resource allocation in the presence of uncertain future values and present bias. We prove the existence of a simple dominant strategy equilibrium, regardless of an agent's value distribution, level of present bias, or degree of sophistication. Simulation results show that the mechanism improves utilization and achieves higher welfare and better equity in comparison with mechanisms broadly used in practice, as well as mechanisms that are welfare-optimal for settings without present bias.

In future work, it will be interesting to conduct empirical studies to better understand behavior in settings such as exercise studios and events, with the goal of separating the effect on utilization of uncertainty from that of present bias. Another interesting direction is to generalize the model to allow for more than two time periods, where agents may arrive asynchronously, when uncertainty unfolds over time, and where resources can be re-allocated.

ACKNOWLEDGMENTS

The authors would like to thank Ido Erev, Matt Juszczak, Scott Kominers, and Kyle Pasake for helpful comments and discussions.

REFERENCES

[1] S. Alaei, K. Jain, and A. Malekian. 2016. Competitive Equilibria in Two-Sided Matching Markets with General Utility Functions. *Operations Research* 64, 3 (2016), 638–645.

[2] Alp E Atakan and Mehmet Ekmekci. 2014. Auctions, actions, and the failure of information aggregation. *American Economic Review* 104, 7 (2014).

[3] Ned Augenblick and Matthew Rabin. 2015. An experiment on time preference and misprediction in unpleasant tasks. *Review of Economic Studies. Forthcoming* (2015).

[4] John Beshears, James J Choi, Christopher Harris, David Laibson, Brigitte C Madrian, and Jung Sakong. 2015. *Self Control and Commitment: Can Decreasing the Liquidity of a Savings Account Increase Deposits?* Working Paper 21474. National Bureau of Economic Research. https://doi.org/10.3386/w21474

[5] John Beshears, James J Choi, David Laibson, Brigitte C Madrian, and Jung Sakong. 2011. *Self control and liquidity: How to design a commitment contract.* Technical Report WR-895-SSA. RAND Working Paper Series.

[6] Gharad Bryan, Dean Karlan, and Scott Nelson. 2010. Commitment devices. *Annu. Rev. Econ.* 2, 1 (2010), 671–698.

[7] Richard E Caves. 2003. Contracts between art and commerce. *Journal of economic Perspectives* (2003), 73–84.

[8] Jonathan D Cohen, Keith Marzilli Ericson, David Laibson, and John Myles White. 2016. *Measuring Time Preferences.* Working Paper 22455. National Bureau of Economic Research. https://doi.org/10.3386/w22455

[9] Pascal Courty and Hao Li. 2000. Sequential screening. *The Review of Economic Studies* 67, 4 (2000), 697–717.

[10] Rahul Deb and Debasis Mishra. 2014. Implementation with contingent contracts. *Econometrica* 82, 6 (2014), 2371–2393.

[11] G. Demange and D. Gale. 1985. The strategy structure of two-sided matching markets. *Econometrica* 53 (1985), 873–888.

[12] Mehmet Ekmekci, Nenad Kos, and Rakesh Vohra. 2016. Just Enough or All: Selling a Firm. *American Economic Journal: Microeconomics* 8, 3 (2016), 223–56.

[13] Keith Marzilli Ericson and David Laibson. 2018. *Intertemporal Choice.* Working Paper 25358. National Bureau of Economic Research. https://doi.org/10.3386/w25358

[14] Xavier Giné, Dean Karlan, and Jonathan Zinman. 2010. Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation. *American Economic Journal: Applied Economics* 2, 4 (October 2010), 213–35. https://doi.org/10.1257/app.2.4.213

[15] Nick Gravin, Nicole Immorlica, Brendan Lucier, and Emmanouil Pountourakis. 2016. Procrastination with Variable Present Bias. In *Proceedings of the 2016 ACM Conference on Economics and Computation.* ACM, 361–361.

[16] Kenneth Hendricks and Robert H Porter. 1988. An empirical study of an auction with asymmetric information. *The American Economic Review* (1988), 865–883.

[17] Jon Kleinberg, Sigal Oren, and Manish Raghavan. 2017. Planning with multiple biases. In *Proceedings of the 2017 ACM Conference on Economics and Computation.* ACM, 567–584.

[18] Jon M. Kleinberg and Sigal Oren. 2014. Time-Inconsistent Planning: A Computational Problem in Behavioral Economics. *CoRR* abs/1405.1254 (2014). arXiv:1405.1254 http://arxiv.org/abs/1405.1254

[19] Jon M. Kleinberg, Sigal Oren, and Manish Raghavan. 2016. Planning Problems for Sophisticated Agents with Present Bias. *CoRR* abs/1603.08177 (2016). arXiv:1603.08177 http://arxiv.org/abs/1603.08177

[20] David Laibson. 1997. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics* 112, 2 (1997), 443–478.

[21] Hongyao Ma, Reshef Meir, David C. Parkes, and James Zou. 2019. Contingent Payment Mechanisms for Resource Utilization. In *Proceedings of the 18th Conference on Autonomous Agents and MultiAgent Systems (AAMAS'19). arXiv preprint arXiv:1607.06511.* IFAAMAS.

[22] Hongyao Ma, Valentin Robu, Na Li, and David C. Parkes. 2016. Incentivizing Reliability in Demand-Side Response. In *Proceedings of The 25th International Joint Conference on Artificial Intelligence (IJCAI'16).* 352–358.

[23] Hongyao Ma, Valentin Robu, and David C. Parkes. 2017. Generalizing Demand Response Through Reward Bidding. In *Procedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'17).* 60–68.

[24] Ashwin Mehra, Claire J Hoogendoorn, Greg Haggerty, Jessica Engelthaler, Stephen Gooden, Michelle Joseph, Shannon Carroll, and Peter A Guiney. 2018. Reducing Patient No-Shows: An Initiative at an Integrated Care Teaching Health Center. *J Am Osteopath Assoc* 118, 2 (2018), 77–84.

[25] Katherine L Milkman, Julia A Minson, and Kevin GM Volpp. 2013. Holding the Hunger Games hostage at the gym: An evaluation of temptation bundling. *Management science* 60, 2 (2013), 283–299.

[26] Ted O'Donoghue and Matthew Rabin. 1999. Doing it now or later. *American Economic Review* 89, 1 (1999), 103–124.

[27] Ted O'Donoghue and Matthew Rabin. 2001. Choice and procrastination. *The Quarterly Journal of Economics* 116, 1 (2001), 121–160.

[28] Andrzej Skrzypacz. 2013. Auctions with contingent paymentsâĂŤan overview. *International Journal of Industrial Organization* 31, 5 (2013), 666–675.

[29] Hal R Varian. 2007. Position auctions. *international Journal of industrial Organization* 25, 6 (2007), 1163–1178.