Operations Research

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

Multiagent Mechanism Design Without Money

Santiago R. Balseiro, Huseyin Gurkan, Peng Sun

To cite this article:
Santiago R. Balseiro, Huseyin Gurkan, Peng Sun (2019) Multiagent Mechanism Design Without Money. Operations Research 67(5):1417-1436. https://doi.org/10.1287/opre.2018.1820

http://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Vol. 67, No. 5, September–October 2019, pp. 1417–1436
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Methods

Multiagent Mechanism Design Without Money

Santiago R. Balseiro,[a] Huseyin Gurkan,[b] Peng Sun[b]

[a] Graduate School of Business, Columbia University, New York, New York 10027; [b] Fuqua School of Business, Duke University, Durham, North Carolina 27708

Contact: srb2155@columbia.edu, http://orcid.org/0000-0002-0012-3292 (SRB); hg67@duke.edu (HG); psun@duke.edu, http://orcid.org/0000-0002-3892-5936 (PS)

Received: December 20, 2017
Revised: June 29, 2018
Accepted: September 4, 2018
Published Online in Articles in Advance: September 10, 2019

Subject Classifications: games/group decisions: stochastic; probability: dynamic programming/stochastic model applications
Area of Review: Stochastic Models

https://doi.org/10.1287/opre.2018.1820

Abstract. We consider a principal repeatedly allocating a single resource in each period to one of multiple agents, whose values are private, without relying on monetary payments over an infinite horizon with discounting. We design a dynamic mechanism that induces agents to report their values truthfully in each period via promises/threats of future favorable/unfavorable allocations. We show that our mechanism asymptotically achieves the first-best efficient allocation (the welfare-maximizing allocation as if values are public) as agents become more patient and provide sharp characterizations of convergence rates to first best as a function of the discount factor. In particular, in the case of two agents we prove that the convergence rate of our mechanism is optimal—that is, no other mechanism can converge faster to first best.

Supplemental Material: The online appendices are available at https://doi.org/10.1287/opre.2018.1820.

Keywords: dynamic mechanism design • social efficiency • multiagent games • resource allocation without money

1. Introduction

Mechanism design for resource allocation with asymmetric information has been extensively studied in economics and, more recently, in computer science (see, for example, Nisan et al. 2007) and operations research (see, for example, Vohra 2011; Li et al. 2012; and Zhang 2012a, b). Most of these studies allow, and often rely on, monetary transfers as part of the mechanism. In certain problem settings, especially with repeated interactions between agents, monetary transfers may not be practical. For example, monetary transfers may be inconvenient when allocating CPU or memory resources in shared computing environments; using money to manage incentives may sound awkward when an organization is deciding on the allocation of an internal resource, such as scheduling a conference room; in some medical resource allocation settings, monetary transfer may be a source of controversy. In the examples above, resource allocation occurs repeatedly, and agents' values for the resource might change over time.

In this paper, we study the problem of a socially maximizing planner repeatedly allocating a single resource without relying on monetary transfers. Specifically, we consider a discrete time infinite horizon setting where agents' private valuations for the resource are independent. The planner is able to commit to a long-term allocation mechanism but is not able to collect monetary transfers from agents or transfer money between agents. Both the planner and agents share the same time discount factor. The objective of the planner is to maximize allocation efficiency—that is, the expected total discounted utilities from the resource in all periods.

If agents can pay for the resource with money, repeatedly implementing the standard Vickrey–Clarke–Groves (VCG) mechanism achieves the "first-best" allocation (also referred to as the "efficient allocation"). That is, the resource is allocated to the agent with the highest realized value in every period. Obviously, first best can also be achieved in settings where valuations are publicly observable. Without monetary payment, however, agents have the natural tendency of claiming that their values for the resource are the highest possible. In this case, if the resource is allocated only once, the planner can do no better than allocating the resource to the agent with the highest expected value. Repeated interactions, however, allow the planner to leverage future allocations when eliciting current period values, which may improve efficiency.

In this paper, we design mechanisms without monetary transfers that induce agents to truthfully reveal private information via promises/threats of future favorable/unfavorable allocations. Moreover, we show that our mechanism asymptotically achieves the first-best allocation as agents become more patient. The fact that the first-best allocation can be approximated may not be surprising, given the Folk theorem established in Fudenberg et al. (1994). In comparison, however, our paper takes an operational focus. In particular, while Fudenberg et al. (1994) implies the existence of an approximately efficient mechanism as the discount factor is close enough to 1, we present a specific, easy to implement, mechanism for given time discount factors. Furthermore, our construction and analysis yields the convergence rate for the approximation. The equilibrium strategy for each agent in the game under our mechanism is also quite simple: each agent truthfully reports the valuation in each period.

1.1. An Overview of Our Approach

Invoking the revelation principle, we focus on direct dynamic mechanisms in which allocations in each period depend on reported private values over time. A direct dynamic mechanism induces a game between agents. Our solution concept for this game is perfect Bayesian equilibrium (PBE). Without loss of generality, we restrict attention to so-called *incentive compatible* mechanisms, under which all agents reporting truthfully regardless of past history is a PBE.

We consider the set of all achievable utilities—that is, the set of vectors representing all agents' total discounted expected utilities that can be attained by incentive compatible dynamic mechanisms. Using the set of achievable utilities, one could readily optimize any objective involving the total expected utility of each agent and, in particular, identify the most efficient mechanism. Characterizing the set of achievable utilities by analyzing all dynamic mechanisms directly appears impossible because the dimensionality of the history grows exponentially with time.

Therefore, we provide an alternative characterization of the mechanism and set of achievable utilities using the so-called "promised utility" framework, which allows us to represent long-term contracts recursively (Spear and Srivastava 1987, Thomas and Worrall 1990). In this framework, agents' total discounted utilities, also referred to as promised utilities, are state variables. In each time period, the planner selects a "stage mechanism" consisting of an allocation function as well as a future promise function, both depending on the current promised utility state and reported values. These functions map the current time period's reports to an allocation and promised utilities for the next time period, respectively. A stage mechanism is incentive compatible if each agent's total expected utility from the current period's allocation and the discounted future promise is maximized by reporting truthfully. Furthermore, the stage mechanism needs to satisfy "promise keeping" constraints, which impose that the total expected utility delivered by the mechanism is equal to the current promised utility. Therefore, implementing an incentive compatible stage mechanism recursively delivers the promised utility for each agent.

Following Abreu et al. (1990), we provide a recursive formulation in the spirit of dynamic programming to characterize the set of achievable utilities. Specifically, we define a Bellman-like operator that maps a target set of future promised utilities to a set of current period promised utilities. The mapping specifies that there exists an incentive compatible stage mechanism that achieves every current promised utility with future promises lying in the target set. The set of achievable utilities is, therefore, a fixed-point of this Bellman-like operator for sets.

Our main contribution is the construction of an incentive compatible mechanism that can attain, as the discount factor approaches 1, the "perfect information" (PI) achievable set—that is, the set of utilities attainable when values are publicly observable. It is clear that the vector of first-best utilities following efficient allocation is in the PI achievable set. Our approach, therefore, provides a constructive proof that first best is asymptotically achievable in repeated settings without monetary transfers. Although our mechanism is not necessarily optimal for a fixed discount factor, it is relatively simple to implement, in the sense that one does not need to solve for a fixed point of the aforementioned Bellman-like operator. Moreover, in the case of two agents, we show that the average social welfare of our mechanism converges to first best at rate $1-\beta$, where $\beta \in (0, 1)$ is the discount factor, approaching 1. Notably, using the linear programming approach to approximate dynamic programming, we also show that this rate is tight—that is, no other mechanism can converge faster to first best.

More specifically, our mechanism allocates the resource to the agent with the largest *weighted* value, where weights are dynamically adjusted based on promised utilities. This allocation function is inspired by the PI achievable set: because every point in the efficient frontier of the PI achievable set can be attained by a weighted allocation, our mechanism seeks to maximize efficiency by setting the weights according to the "closest" point in the efficient frontier of the PI achievable set.

The design of future promise functions resembles the d'Aspremont–Gérard-Varet–Arrow mechanism (Arrow 1979, d'Aspremont and Gérard-Varet 1979) and the risk-free transfers of Esö and Futo (1999). Fixing the allocation rule, the interim future promises are uniquely determined from the incentive compatibility constraints. To implement ex-post future promises, we try to minimize the risk of the next time period's promised utilities lying outside the achievable set. The proof that the proposed mechanism is incentive compatible is geometric in nature and relies on ideas from convex optimization, which is quite involved and resembles the equilibrium construction for the Folk theorem in Fudenberg et al. (1994). The essential intuition of our mechanism, however, is clear. Reporting a higher value increases an agent's chance of receiving the resource in the current period, while lowering the agent's future promised utility. A lower promised utility, in turn, leads to a lower weight in the allocation rule and, hence, a lower chance of allocating the resource to this agent in the future.

1.2. Related Literature

There is an extensive literature on dynamic mechanism design problems. Most of the literature focuses on settings in which monetary transfers are allowed. Settings under consideration include dynamically changing populations with fixed information and fixed populations with dynamically changing information. Under these environments, the efficient outcome can be implemented using natural generalizations of the static VCG mechanism to dynamic setting (see, e.g., Parkes and Singh 2004, Bergemann and Välimäki 2010, Gershkov and Moldovanu 2010). These mechanisms maintain efficient allocation of resources while incentivizing truthful reporting by choosing transfers that are equal to the externality that an agent imposes on others. In our setting, however, incentives cannot be readily aligned with transfers, and the efficient outcome is not implementable in general. We refer the reader to the survey by Bergemann and Said (2011) for a more in-depth discussion on dynamic mechanism design problems with monetary transfers.

Jackson and Sonnenschein (2007) study a general framework for resource allocation in a finite horizon model without discounting in which agents learn all private information at time 0. They consider a budget-based mechanism in which each agent can report each type a limited number of times and prove that, as the number of time periods increases, the inefficiency due to asymmetric information diminishes to 0. In this paper, we extend their budget-based mechanism to our discounted infinite horizon setting in which agents sequentially learn their values. We show that the best possible rate of convergence of a budget-based mechanism to first best is at most $(1-\beta)^{1/2}$; lower than the convergence rate of our mechanism. Even though explicitly characterizing the equilibria of a budget-based mechanism is challenging, a remarkable feature of budget-based mechanisms is that all equilibria are asymptotically efficient. In comparison, under our incentive compatible mechanism, one simple equilibrium that achieves efficiency asymptotically is reporting truthfully (which does not necessarily constitute an equilibrium under budget-based mechanisms).

There is a recent stream of studies considering dynamic mechanism design without money. Guo and Hörner (2015) consider the problem of repeatedly allocating a costly resource to a single agent whose values evolve according to a two-state Markov chain and characterize the optimal allocation rule. In this paper we study settings with multiple agents and continuous values. In our setting where the marginal cost of resource is 0, the problem becomes trivial with a single agent. This is because it is optimal to always allocate the resource to the agent in each period. Guo et al. (2009) study the design of dynamic mechanisms with multiple agents and provide a mechanism that achieves at least 75% of the efficient allocation. While their mechanism is guaranteed to attain a fixed proportion of the efficient allocation for all discount factors larger than a threshold, it does not necessarily achieve the first-best allocation as the discount rate converges to 1. Johnson (2014) studies a similar problem with multiple agents and discrete private values, and provides some numerical evidence that the optimal mechanism achieves higher social welfare as the discount factor increases. In our paper, we provide a relatively simple mechanism in quasi-closed form and analytically prove that it achieves first best asymptotically. In addition, agents' values are continuous in our model, which requires tackling some technical challenges in characterizing the optimal mechanism, in exchange for simpler mechanisms without complicated tie-breaking randomization for discrete value settings. Gorokh et al. (2016) also study a similar setting with a finite number of periods and discrete values. They provide a mechanism that can be implemented via artificial currencies and show that the performance of their mechanism approximately achieves first best. Different from ours, the mechanism in Gorokh et al. (2016) satisfies incentive compatibility constraints approximately. In particular, truthful reporting does not constitute an equilibrium when the horizon is finite. In comparison, our mechanism is guaranteed to be incentive compatible.

Our model and analysis relies on the promised utility framework (Spear and Srivastava 1987, Abreu et al. 1990, Thomas and Worrall 1990). In settings with monetary transfers, any nonnegative promised utility can be achieved by having the planner transfer money to the agents. Thus, constructing feasible incentive compatible mechanism is relatively straightforward, and the problem of the planner reduces to that of optimizing certain objective. When monetary transfers are not allowed, however, the planner can no longer subsidize agents. Constructing a feasible incentive compatible mechanism is challenging in this case because the planner needs to guarantee that future promises can be delivered exclusively via allocations. Abreu et al. (1990) introduce a recursive approach to study pure strategy sequential equilibria of repeated games with imperfect monitoring. In the paper, they characterize a *self-generating* set of sequential equilibria payoffs. We extend their recursive approach to characterize a self-generating set of utilities that can be achieved by incentive compatible mechanisms. Although our setting, which is focused on adverse selection (private signal) issues in mechanism design, is different from that of Abreu et al. (1990), we adopt some of their proof techniques related to self-generating sets.

Fudenberg et al. (1994) build on the dynamic programming framework of Abreu et al. (1990) to establish Folk theorems for finite repeated games under imperfect information. In particular, theorem 8.1 of Fudenberg et al. (1994) implies that in a setting similar to ours (except that the valuation set is finite), agents' payoff under efficient allocation can be approximated as long as the time discount factor is close enough to 1. Fudenberg et al. (1994), however, do not provide an explicit description of such a direct mechanism. For potentially indirect mechanisms (for example, "allocating to the agent with the highest report"), describing agents' equilibrium reporting strategies to sustain the Folk theorem approximation appears nontrivial. As a matter of fact, computing equilibrium strategies in repeated games is a challenging problem (see, e.g., Judd et al. 2003, Abreu et al. 2017, and Yeltekin et al. 2017). In comparison, our mechanism is well specified. Another advantage of our mechanism is that the equilibrium strategy for agents is straightforward: reporting the true value in each period. Furthermore, our construction and analysis explicitly characterize the rate of convergence to first best of the mechanism's social welfare in terms of the time discount factor. It turns out that some steps in our construction resemble the steps used in Fudenberg et al. (1994) to prove existence of an asymptotically efficient equilibrium, and we will point them out in our paper.

Another stream of literature that is related to our work is the study of "scrip systems," which are nonmonetary trade economies for the exchange of resources (Friedman et al. 2006; Kash et al. 2007, 2012, 2015; Johnson et al. 2014). In these systems, scrips are used in place of government issued money, and the resource is priced at a fixed amount of scrips whenever trade occurs. The promised utilities in our model can be perceived as scrips. According to our mechanism, the agent who receives the resource in a period sees his promised utility decrease while others' increases. The exchange of promised utilities according to our mechanism, however, is not fixed. In fact, it depends on the current promised utilities of all agents. From this perspective, our mechanism is more general than the ones considered in the existing studies of scrip systems.

The remaining of the paper is organized as the following. We first introduce the model, its recursive formulation, and various concepts related to self-generating sets in Section 2. In Section 3, we present the main phase of our mechanism. We then focus on the two-agent case and complete the description of the mechanism for the boundary region in Section 4. We compare the performances of our mechanism and the mechanism inspired by Jackson and Sonnenschein (2007) in Section 5. In Section 6, we describe the intuition behind the truthfulness and asymptotic efficiency of our mechanism. The general case of more than two agents is discussed in Section 7. Finally, Section 8 concludes the paper with comments on potential future directions. Proofs for results are presented in the online appendix.

2. Model and Problem Formulation

We consider a discrete time infinite horizon setting where a social planner repeatedly allocates a single resource to one of multiple agents in each period without relying on monetary transfers. We index agents by $i \in \{1, \ldots, n\}$ and denote by $\mathbf{v}_t = (v_{i,t})_{i=1}^n$ the random vector of agents' (private) values for the resource in period $t \geq 1$. Agent i's values in each period are independent and identically distributed with cumulative distribution function $F_i(\cdot)$ and density function $f_i(\cdot)$. Values are supported in the bounded set $[0, \bar{v}]$ and densities are bounded in their domain—that is, $0 < \underline{f} \leq f_i(v_i) \leq \bar{f} < \infty$ for all $v_i \in [0, \bar{v}]$. Moreover, we denote the minimum and maximum of the first moment of agents' values by $\underline{m} = \min_{i \in \{1,\ldots,n\}} \mathbb{E}[v_i]$ and $\overline{m} = \max_{i \in \{1,\ldots,n\}} \mathbb{E}[v_i]$; and assume that $0 < \underline{m} \leq \overline{m} < \infty$. The planner and agents share the same discount factor $\beta \in (0,1)$.[1] An agent's overall utility is given by the discounted sum of the valuations generated by the allocations of the resource across the horizon. The objective of the planner is to maximize the expected discounted sum of total valuations in all periods.

Notation. For a sequence of vectors $\mathbf{a} = ((a_{i,t})_{i=1}^n)_{t=1}^\infty \in \mathbb{R}^{n \times \infty}$, we denote by $a_{i,1:t} = (a_{i,\ell})_{\ell=1}^t \in \mathbb{R}^t$ the i^{th} components of the 1st to the tth vectors, and by $\mathbf{a}_{1:t} = ((a_{i,\ell})_{i=1}^n)_{\ell=1}^t \in \mathbb{R}^{n \times t}$ the entire 1st to the tth vectors. For a given vector $\boldsymbol{x}$, we denote by $\boldsymbol{x}_{-i}$ the vector obtained by removing x_i from $\boldsymbol{x}$, and $\boldsymbol{x}^\top$ its transpose. For any two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ in $\mathbb{R}^n$, the inequality $\boldsymbol{x} \leq (\geq) \boldsymbol{y}$ represents that $x_i \leq (\geq) y_i$ for each component i. For any function $g : \mathbb{R}^n \to \mathbb{R}$, we use ∇g to represent its gradient. We use $\mathbf{1}\{\cdot\}$ to represent the indicator function.

2.1. Dynamic Mechanisms and Achievable Utilities

We assume the planner commits to a direct dynamic mechanism. That is, in each period, the agents learn their valuations of the resource, and each reports a value to the planner. The planner, in turn, determines the allocation of the resource for the period based on the entire history of reports and allocations, and publicly announces all agents' reports and allocations in the end of the period.

Nonanticipating Mechanisms and Strategies. Formally, a dynamic mechanism $\boldsymbol{\pi}$ is a sequence of allocation rules $\boldsymbol{\pi} = ((\pi_{i,t})_{i=1}^n)_{t=1}^\infty$, where $\pi_{i,t}$ is the probability that the resource is allocated to agent i in period t. We denote by $\mathcal{P} \subseteq \mathbb{R}^n$ the set of n-dimensional feasible allocations—that is, $\mathcal{P} \triangleq \{\boldsymbol{\pi} \in [0,1]^n : \boldsymbol{\pi} \geq \mathbf{0}, \sum_{i=1}^n \pi_i \leq 1\}$—and restrict that $\boldsymbol{\pi}_t \in \mathcal{P}$ for all t. Any mechanism $\boldsymbol{\pi}$ induces a dynamic game among agents, in which each agent i submits a report $\hat{v}_{i,t} \in [0, \bar{v}]$ in each period t and receives an allocation $\pi_{i,t}$. We define the history available at time t, $\mathbf{h}_t = (h_{i,t})_{i=1}^n$, as all reports and previous allocations up to time t, where $h_{i,t} = (\hat{v}_{i,1:t-1}, \pi_{i,1:t-1})$.[2] We define the set of all possible histories that can be observed in periods $t \geq 2$ as $\mathcal{H}_t = \prod_{i=1}^n \mathcal{H}_{i,t}$ where $\mathcal{H}_{i,t} = [0,\bar{v}]^{t-1} \times [0,1]^{t-1}$, and $\mathcal{H}_1 = \{\emptyset\}$. We assume that the planner discloses the reports and the allocations after each round, so that the history is publicly observed. A nonanticipating strategy profile $\boldsymbol{\sigma} = ((\sigma_{i,t})_{t=1}^\infty)_{i=1}^n$ for agents consists of reporting functions for each period, $\sigma_{i,t} : [0,\bar{v}] \times \mathcal{H}_t \to [0,\bar{v}]$, that depend only on the value $v_{i,t}$ of agent i in period t, and the public history $\mathbf{h}_t$ up to that period—i.e., $\sigma_{i,t}(v_{i,t}, \mathbf{h}_t) = \hat{v}_{i,t}$. We say a dynamic mechanism $\boldsymbol{\pi}$ is nonanticipating if $\pi_{i,t}$ depends only on the current period reports $\hat{\mathbf{v}}_t$ and the history $\mathbf{h}_t$—that is, $\boldsymbol{\pi}_t : [0,\bar{v}]^n \times \mathcal{H}_t \to \mathcal{P}$. Because agents' values in each period are independent, we restrict attention to nonanticipating mechanisms and reporting strategies that depend on past reports and allocations, but not on past values. Because agents' actions and planner's allocations are only conditioned on previous reports and allocations, and this information is publicly observed, agents do not need to form beliefs about the past actions of competitors.

Direct Mechanisms and Truthful Reporting. Following the Revelation Principle, without loss of generality, we can focus on direct mechanisms in which agents report their values truthfully to the planner. In particular, for the game induced by a mechanism, we consider perfect Bayesian equilibria (PBE) in truthful reporting strategies, with beliefs that assign probability 1 to the event that other agents report truthfully. Therefore, we enforce *interim incentive compatibility* constraints, which ensure that an agent is better off adopting the truthful reporting strategy than any other reporting strategy when other agents report their values truthfully under mechanism $\boldsymbol{\pi}$.

To elaborate, we introduce some notations. We use the notation $V_{i,t}$ to represent agent i's utility-to-go in period t when the planner implements mechanism $\boldsymbol{\pi}$, all agents employ strategy profile $\boldsymbol{\sigma}$, agent i's value for the resource is $v_{i,t}$, and the history is $\mathbf{h}_t$. That is,

$$V_{i,t}(\boldsymbol{\pi}, \boldsymbol{\sigma} | v_{i,t}, \mathbf{h}_t) \triangleq (1-\beta)\mathbb{E}^{\boldsymbol{\pi},\boldsymbol{\sigma}} \left[v_{i,t}\pi_{i,t}(\hat{\mathbf{v}}_t, \mathbf{h}_t) + \sum_{\ell=t+1}^{\infty} \beta^{\ell-t} v_{i,\ell}\pi_{i,\ell}(\hat{\mathbf{v}}_\ell, \tilde{\mathbf{h}}_\ell) \mid v_{i,t}, \mathbf{h}_t \right],$$

where $\mathbb{E}^{\boldsymbol{\pi},\boldsymbol{\sigma}}[-|v_{i,t}, \mathbf{h}_t]$ represents the expectation with respect to histories $(\tilde{\mathbf{h}}_\ell)_{\ell > t}$ induced by the mechanism $\boldsymbol{\pi}$ and the strategy profile $\boldsymbol{\sigma}$, given that the value of agent i at time t is $v_{i,t}$ and an initial history $\mathbf{h}_t$. For $\ell \geq t$, we denote by $\hat{v}_{i,\ell} = \sigma_{i,\ell}(v_{i,\ell}, \tilde{\mathbf{h}}_\ell)$ agent i's reported value in period ℓ, in which history $\tilde{\mathbf{h}}_\ell$ is recursively defined as $\tilde{\mathbf{h}}_\ell = (\tilde{\mathbf{h}}_{\ell-1}, (\hat{\mathbf{v}}_{\ell-1}, \boldsymbol{\pi}_{\ell-1}(\hat{\mathbf{v}}_{\ell-1}, \tilde{\mathbf{h}}_{\ell-1})))$ starting from $\tilde{\mathbf{h}}_t = \mathbf{h}_t$. To facilitate comparisons across different discount factors, in the expression for $V_{i,t}$ we multiply by $1-\beta$ to obtain an "average" discounted utility-to-go.

Using this notation, the interim incentive compatibility constraints on mechanism $\boldsymbol{\pi}$ are given as follows:

$$V_{i,t}\left(\boldsymbol{\pi}, \boldsymbol{\iota} | v_{i,t}, \mathbf{h}_t\right) \geq V_{i,t}\left(\boldsymbol{\pi}, (\sigma_i, \boldsymbol{\iota}_{-i}) | v_{i,t}, \mathbf{h}_t\right), \quad \forall v_{i,t}, \sigma_i, \mathbf{h}_t, i, t, \tag{1}$$

where ι_i represents the truthful reporting strategy for agent i—that is, $\iota_{i,t}(v_{i,t}, \mathbf{h}_t) = v_{i,t}$. Hereafter, we say a nonanticipating mechanism $\boldsymbol{\pi}$ is perfect Bayesian incentive compatible (PBIC) if $\boldsymbol{\pi}$ satisfies (1). These constraints enforce that, at every point in time, each agent is better off reporting truthfully when other agents report truthfully, regardless of past reports and allocations.

We next define the set of achievable utilities, $\mathcal{U}_\beta$ as the following:

$$\mathcal{U}_\beta \triangleq \{\mathbf{u} \in \mathbb{R}^n \mid u_i = V_i(\boldsymbol{\pi}, \boldsymbol{\iota}), \text{ for a PBIC mechanism } \boldsymbol{\pi}\}, \tag{2}$$

where $V_i(\boldsymbol{\pi}, \boldsymbol{\sigma}) \triangleq \mathbb{E}_{v_{i,1}}[V_{i,1}(\boldsymbol{\pi}, \boldsymbol{\sigma} | v_{i,1}, \emptyset)]$ is the total expected utility of agent i when the planner implements mechanism $\boldsymbol{\pi}$, and agents employ strategy profile $\boldsymbol{\sigma}$. For any state $\mathbf{u} = (u_i)_{i=1}^n$ in $\mathcal{U}_\beta$, there must exist a nonanticipating direct (dynamic) PBIC mechanism $\boldsymbol{\pi}$ that achieves utility u_i for all agents $i = 1, \ldots, n$. Specifically, when the planner implements mechanism $\boldsymbol{\pi}$, every agent truthfully reporting each period's value while believing with probability 1 that others report truthfully is a PBE. Moreover, the corresponding expected discounted sum of the valuations generated by $\boldsymbol{\pi}$ for agent i is equal to u_i.

Because social welfare is the expected discounted sum of total valuations, the component sum of $\mathbf{u}$ corresponds to the social welfare obtained by $\boldsymbol{\pi}$. Therefore, given the set of achievable utilities $\mathcal{U}_\beta$, the maximum social welfare that can be obtained by a PBIC mechanism is readily given by $J_\beta^* = \max_{\mathbf{u} \in \mathcal{U}_\beta} \sum_{i=1}^n u_i$. We use the vector $\mathbf{u}_\beta^* \in \arg\max_{\mathbf{u} \in \mathcal{U}_\beta} \sum_{i=1}^n u_i$ to represent the utilities of the agents under an optimal mechanism. (Although in this paper we only focus on the maximum achievable social welfare, other objectives involving the total expected utility of each agent can be easily accommodated accordingly from the feasible set $\mathcal{U}_\beta$.) Unfortunately, characterizing $\mathcal{U}_\beta$ directly by analyzing all nonanticipating mechanisms is not possible in general, because the dimensionality of the history grows exponentially with time. Therefore, in the following section, we provide an equivalent, recursive definition of $\mathcal{U}_\beta$ using the promised utility framework.

2.2. Promised Utility Framework

Because the planner has commitment power and values are independent, we can employ the *promised utility* framework to recursively formulate the set of achievable utilities, $\mathcal{U}_\beta$. We first present the framework and then show its equivalence with the definition (2) in the end of this subsection.

Note that to determine if an allocation for the current time period is incentive compatible, the planner needs to understand the impact of today's actions on the continuation game induced by the mechanism. The promised utility framework builds on the observation that, because agents are expected value maximizers, the next period's continuation utilities constitute a sufficient statistic for the problem of determining if an allocation for the current time period is incentive compatible. Loosely speaking, we use $\mathbf{u}_t = (u_{i,t})_{i=1}^n$ to represent the vector of expected discounted total future values starting from period t. That is,

$$u_{i,t} = (1-\beta)\mathbb{E}\left[\sum_{\ell=t}^{\infty} \beta^{\ell-t} v_{i,\ell}\pi_{i,\ell}\right].$$

In the beginning of period t, the planner needs to fulfill the current state $\mathbf{u}_t$ as a promise to agents through future allocations. We enforce this recursively by having the planner determine the current period's allocation of the resource, as well as next period's promised utilities $\mathbf{u}_{t+1}$. For each agent, the expected value of the current period allocation plus the next period's promised utility has to equal the current period promised utility.

We refer to the mechanism associated with the allocation of a single resource in a time period as a *stage mechanism*. More formally, for any given state $\mathbf{u} \in \mathbb{R}^n$, a stage mechanism is given by an allocation function $\mathbf{p}(\cdot|\mathbf{u}) : [0, \bar{v}]^n \to \mathcal{P}$ and a *future promise function* $\mathbf{w}(\cdot|\mathbf{u}) : [0, \bar{v}]^n \to \mathbb{R}^n$, which map the vector of agents' reports in the current period to an allocation vector $\mathbf{p}$ and a next state $\mathbf{w}$, respectively. We drop the dependence on $\mathbf{u}$ when referring to a fixed state.

A stage mechanism $(\mathbf{p}, \mathbf{w})$ should satisfy the following constraints. First, the allocation function should be feasible. That is,

$$\sum_{i=1}^{n} p_i(\mathbf{v}) \le 1 \text{ and } \mathbf{p}(\mathbf{v}) \ge \mathbf{0}, \quad \forall \mathbf{v}. \tag{FA}$$

Additionally, the mechanism should satisfy the following *promise keeping* constraint,

$$u_i = \mathbb{E}[(1-\beta)v_i p_i(\mathbf{v}) + \beta w_i(\mathbf{v})], \quad \forall i, \tag{PK(u)}$$

which guarantees that the promised utility $\mathbf{u}$ is fulfilled by the mechanism. Finally, an agent should not have an incentive to misreport the true type. The following interim *incentive compatibility* constraint imposes that, for each agent, reporting the value truthfully yields an expected utility at least as large as any other strategy, when other agents report truthfully. Denote by $P_i(v) \triangleq \mathbb{E}_{\mathbf{v}_{-i}}[p_i(v, \mathbf{v}_{-i})]$ the interim allocation, and by $W_i(v) \triangleq \mathbb{E}_{\mathbf{v}_{-i}}[w_i(v, \mathbf{v}_{-i})]$ the interim future promise of agent i. The incentive compatibility constraints are given by

$$(1-\beta)vP_i(v) + \beta W_i(v) \ge (1-\beta)vP_i(v') + \beta W_i(v'), \ \forall i, v, v'. \tag{IC}$$

Using the constraints defined above, we next define the set operator $B_\beta : 2^{\mathbb{R}_+^n} \to 2^{\mathbb{R}_+^n}$ for a given set $\mathcal{A} \subset \mathbb{R}_+^n$ as follows:

$$B_\beta(\mathcal{A}) = \{\mathbf{u} \in \mathbb{R}_+^n \mid \exists(\mathbf{p}, \mathbf{w}) \text{ satisfying (IC), (FA), (PK}(\mathbf{u})\text{), and } \mathbf{w}(\mathbf{v}) \in \mathcal{A} \text{ for all } \mathbf{v}\}. \tag{3}$$

Essentially, set $B_\beta(\mathcal{A})$ contains all promised utilities $\mathbf{u}$ that can be supported by future promise functions $\mathbf{w}$ in set $\mathcal{A}$, while satisfying feasibility and incentive compatibility constraints. Although the operator B_β is analogous to operator B in Abreu et al. (1990, p. 1047), the specific constraints used in our definition are different. Abreu et al. (1990) study sequential equilibria of repeated games with imperfect monitoring, in which each player has a finite action space. In our setting, however, the stage game itself is induced by the stage mechanism selected by the planner. In this game, each agent has an uncountable action space because agents' private values are continuous.

The operator B_β provides a certificate to check whether a given set is a subset of $\mathcal{U}_\beta$, according to the following result.

Proposition 2.1. *If a set $\mathcal{A}$ satisfies $\mathcal{A} \subseteq B_\beta(\mathcal{A})$, then we have $B_\beta(\mathcal{A}) \subseteq \mathcal{U}_\beta$.*

Following Abreu et al. (1990), we refer to sets that satisfy $\mathcal{A} \subseteq B_\beta(\mathcal{A})$ as *self-generating*. That is, all of the promised utilities in such a set $\mathcal{A}$ can be fulfilled with future promises from within the same set. Proposition 2.1 implies that any state in a self-generating set can be achieved by a PBIC mechanism, because this state is in the set $\mathcal{U}_\beta$. Furthermore, if there exists a specific mechanism $(\mathbf{p}, \mathbf{w})$ that satisfies (IC), (FA), (PK($\mathbf{u}$)) for all $\mathbf{u} \in \mathcal{A}$, and $\mathbf{w}(\mathbf{v}|\mathbf{u}) \in \mathcal{A}$ for all $\mathbf{v}$ and $\mathbf{u} \in \mathcal{A}$, then we call the set $\mathcal{A}$ to be *self-generating with respect to mechanism* $(\mathbf{p}, \mathbf{w})$.

Remark 2.1. It is worth noting that the "lower triangle" set $L \triangleq \left\{\mathbf{u} \in \mathbb{R}_+^n \,\middle|\, \sum_{i=1}^n u_i/\mathbb{E}[v_i] \le 1\right\}$ is self-generating. In fact, for any state $\mathbf{u} \in L$, consider the random allocation rule $p_i^L(\mathbf{v}|\mathbf{u}) \triangleq u_i/\mathbb{E}[v_i]$ regardless of the value $\mathbf{v}$. Such an allocation rule $\mathbf{p}^L$, together with the future promise functions $\mathbf{w}^L(\mathbf{v}|\mathbf{u}) \triangleq \mathbf{u}$, achieves utilities $\mathbf{u}$—i.e., satisfy (PK($\mathbf{u}$)). Therefore, the lower triangle set L is self-generating with respect to mechanism $(\mathbf{p}^L, \mathbf{w}^L)$ and, therefore, is a subset of $\mathcal{U}_\beta$ for any discount factor β.

Proposition 2.1 states that any self-generating set is a subset of the set of achievable utilities, $\mathcal{U}_\beta$. The following result further demonstrates that set $\mathcal{U}_\beta$ itself is self-generating, and, therefore, a fixed point of the operator B_β. The result implies that the set of achievable utilities, $\mathcal{U}_\beta$, defined according to summations of allocations over an infinite horizon, can be equivalently represented through stage mechanisms. In particular, the stage mechanisms satisfy constraints (PK(u)) and (IC) in a recursive manner, and with the future promise functions $\mathbf{w}$ lying in the same set $\mathcal{U}_\beta$.

Proposition 2.2. *The set of achievable utilities, $\mathcal{U}_\beta$, satisfies $\mathcal{U}_\beta \subseteq B_\beta(\mathcal{U}_\beta)$. Therefore, $\mathcal{U}_\beta = B_\beta(\mathcal{U}_\beta)$.*

In Online Appendix G, we provide a procedure to numerically compute the set $\mathcal{U}_\beta$ using value iteration in the space of support functions.

2.3. Perfect Information Achievable Set

We next propose a "perfect information achievable set" that provides an upper bound on the set of achievable utilities. Specifically, we define the perfect information (PI) achievable set, $\mathcal{U}$, as the set of utilities attainable when values are publicly observable by the planner. This set is given by

$$\mathcal{U} \triangleq \{\mathbf{u} \in \mathbb{R}^n_+ \mid u_i = \mathbb{E}[v_i p_i(\mathbf{v})] \text{ for all } i, \text{ for some } \mathbf{p} \text{ satisfying (FA)}\}. \tag{4}$$

Clearly, for any $\beta \in [0,1)$ we have that $\mathcal{U}_\beta \subseteq \mathcal{U}$.

Following Luenberger (1969, p. 44), any convex set can be represented by its support functions. In particular, for any fixed $\boldsymbol{\alpha} \in \mathbb{R}^n$ such that $\|\boldsymbol{\alpha}\|_1 = 1$, the support function of set $\mathcal{U}$ is given by

$$\begin{aligned}\phi(\boldsymbol{\alpha}) \triangleq \sup_{\mathbf{u}\in\mathcal{U}} \boldsymbol{\alpha}^\top \mathbf{u} &= \sup_{\mathbf{p} \text{ s.t. (FA)}} \sum_{i=1}^n \mathbb{E}_{\mathbf{v}}\left[\alpha_i v_i p_i(\mathbf{v})\right] \\ &= \mathbb{E}_{\mathbf{v}}\left[\max_{i=1,\dots,n} \alpha_i v_i\right],\end{aligned} \tag{5}$$

where the second equation follows from the definition of the PI achievable set, and the third from optimizing pointwise over values. The support function ϕ satisfies the following properties.

Proposition 2.3. *The support function $\phi(\boldsymbol{\alpha})$ given in (5) is convex, differentiable for $\boldsymbol{\alpha} \in \mathbb{R}^n_+$, and twice differentiable for $\boldsymbol{\alpha} \in \mathbb{R}^n_+$ such that $\boldsymbol{\alpha} > \mathbf{0}$. Moreover, the partial derivatives for all i are given by*

$$\frac{\partial \phi}{\partial \alpha_i}(\boldsymbol{\alpha}) = \mathbb{E}_{\mathbf{v}}\left[v_i \mathbf{1}\left\{\alpha_i v_i \ge \max_{j\neq i} \alpha_j v_j\right\}\right].$$

Differentiability of the support function follows because values are absolutely continuous. The gradient $\nabla\phi$ of the support function for any $\boldsymbol{\alpha}$ corresponds to a point on the efficient frontier of the set $\mathcal{U}$. Specifically, for any convex set $\mathcal{A} \subset \mathbb{R}^n$, define $\mathcal{E}(\mathcal{A})$ to be its efficient frontier:

$$\mathcal{E}(\mathcal{A}) \triangleq \left\{\mathbf{u} \in \mathcal{A} \mid \nexists \mathbf{u}' \in \mathcal{A} \text{ with } \mathbf{u}' \neq \mathbf{u} \text{ and } \mathbf{u}' \ge \mathbf{u}\right\}.$$

Because the support function is differentiable and the set $\mathcal{U}$ is convex and closed, for every state $\mathbf{u}$ on the efficient frontier of $\mathcal{U}$ there exists some $\boldsymbol{\alpha}$ such that $\nabla\phi(\boldsymbol{\alpha}) = \mathbf{u}$ (see, e.g., Schneider 2013, corollary 1.7.3, p. 47).

Furthermore, Proposition 2.3 implies that all points on the efficient frontier of the PI set are achievable by allocations of the form $p_i(\mathbf{v}) = \mathbf{1}\{\alpha_i v_i \ge \max_{j\neq i} \alpha_j v_j\}$. That is, the resource is allocated to the agent with the highest $\boldsymbol{\alpha}$-weighted value. More generally, for any point $\mathbf{u} \in \mathcal{U}$, not necessarily on the efficient frontier, we can define $\boldsymbol{\alpha}^*(\mathbf{u})$ with some abuse of notation as

$$\boldsymbol{\alpha}^*(\mathbf{u}) \in \underset{x:\|x\|_1=1, x\ge \mathbf{0}}{\arg\max} \left\{x^\top \mathbf{u} - \phi(x)\right\}. \tag{6}$$

That is, $\boldsymbol{\alpha}^*(\mathbf{u})$ is the normal vector of the point "closest" to $\mathbf{u}$ on the efficient frontier. It is easy to verify that for any $\mathbf{u} \in \mathcal{E}(\mathcal{U})$, we have $\nabla\phi(\boldsymbol{\alpha}^*(\mathbf{u})) = \mathbf{u}$.

The *first-best* total utility, J^{FB}, is achieved by allocating the resource to the agent that values it most in each period—that is, $J^{\text{FB}} = \mathbb{E}_{\mathbf{v}}[\max_{i=1,\dots,n} v_i]$. Because the first-best utility is attained by the allocation rule $p_i(\mathbf{v}) = \mathbf{1}\left\{v_i \ge \max_{j\neq i} v_j\right\}$,[3] and $\mathcal{U}_\beta \subseteq \mathcal{U}$, we must have

$$J^{\text{FB}} = \max_{\mathbf{u}\in\mathcal{U}} \sum_{i=1}^n u_i \ge \max_{\mathbf{u}\in\mathcal{U}_\beta} \sum_{i=1}^n u_i = J^*_\beta. \tag{7}$$

We denote by $\mathbf{u}^*$ the agents' utilities under the efficient allocation—that is, $u^*_i = \mathbb{E}_{\mathbf{v}}[v_i \mathbf{1}\{v_i \ge \max_{j\neq i} v_j\}]$. Now we provide a high-level summary our approach.

2.4. An Overview of Our Approach

In general, it is not possible to fully characterize the *optimal* mechanism in closed form. Thus, to implement the optimal mechanism, the designer would need to numerically compute the set $\mathcal{U}_\beta$, which might be challenging in some settings. Instead, in this paper we introduce an approximation mechanism that guarantees incentive compatibility, is easy to implement, and asymptotically achieves first best.

Our approach involves designing a factor, $\mathsf{F}_\beta \in [0,1]$ and mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$, such that the *scaled perfect information set* $\mathsf{F}_\beta\mathcal{U} \triangleq \{\mathsf{F}_\beta\mathbf{u} \mid \mathbf{u} \in \mathcal{U}\}$ is self-generating with respect to the mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$. Consequently, every state in $\mathsf{F}_\beta\mathcal{U}$ is achievable following this incentive compatible mechanism. Because the set $\mathsf{F}_\beta\mathcal{U}$ is self-generating, or $\mathsf{F}_\beta\mathcal{U} \subseteq B_\beta(\mathsf{F}_\beta\mathcal{U})$, Proposition 2.1 implies that $\mathsf{F}_\beta\mathcal{U} \subseteq \mathcal{U}_\beta$. Therefore, we obtain the following "sandwich" condition for set $\mathcal{U}_\beta$:

$$\mathsf{F}_\beta\mathcal{U} \subseteq \mathcal{U}_\beta \subseteq \mathcal{U}.$$

In particular, the sequence of factors F_β is designed to converge to 1 as β converges to 1, which implies that the mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$ achieves the efficient allocation asymptotically as the discount factor approaches 1.

3. Central Region and Main Phase Mechanism

Generally speaking, our mechanism specifies allocation and future promise functions that depend on the current promised utilities and reported values. To ensure that the future promises are within the self-generating set, they need to be nonnegative and beneath the efficient frontier of the self-generating set. This motivates different designs when the promised utility is close to 0 and close to the efficient frontier. Therefore, we partition the self-generating set under study into a *central region* and a *boundary region*. In the central region, we use the *main phase* of the mechanism, which is designed to guarantee that promised utilities lie beneath the efficient frontier. In the main phase of the mechanism, the designer allocates the resource by weighing the reported valuations differently. The weights are determined by projecting the current promised utility state to the efficient frontier of the (scaled) perfect information set, and taking the normal vector of the projection as weights. Future promises (the next period's states) are chosen so that they lie as far as possible from the efficient frontier. When the promised utility eventually drifts out of the central region into the *boundary region*, we use the simple randomization mechanism described in Remark 2.1 to deliver the promised utility. This guarantees that, when the promised utility is close to 0, future promises do not to jump too far from the current promised utility and remain nonnegative. The crucial design issue, which is the focus on this section, is the main phase mechanism in the central region. (See Figure 1 for an illustration of the two regions.)

Formally, the central region $\hat{\mathcal{U}}_\beta$ is defined through a pair of scalars $\underline{u}_\beta \in [0, \underline{m}]$ and $\mathsf{F}_\beta \in [0,1]$. The specific values of those scalars depend on the number of agents involved and the distribution of values, and are deferred to the following sections. For this section, it is sufficient to keep in mind that $\underline{u}_\beta$ approaches 0 and F_β approaches 1 as β approaches 1. Define the *central region* $\hat{\mathcal{U}}_\beta$ as

$$\hat{\mathcal{U}}_\beta = \mathsf{F}_\beta \mathcal{U} \cap \{\mathbf{u} \in \mathbb{R}^n \,|\, u_i \geq \underline{u}_\beta,\ \forall i\}, \tag{8}$$

which is a subset of $\mathsf{F}_\beta\mathcal{U}$ that includes states that are component-wise higher than the threshold $\underline{u}_\beta$. Figure 1(a) illustrates such a central region, where we implement the main phase mechanism, to be described next.

Figure 1. The Central Region and Boundary Region in a Two-Agent Case

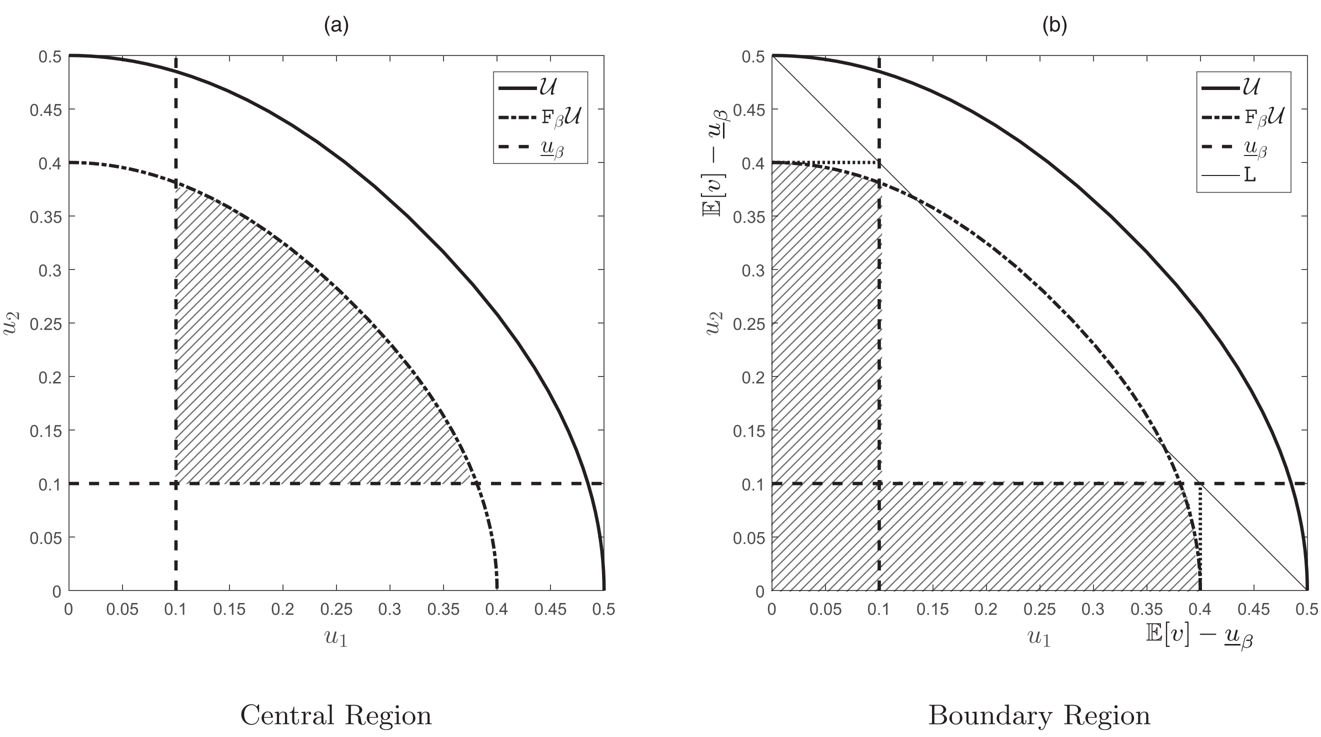

Notes. Values of $\mathbf{v}$ are uniformly distributed in $[0,1]$, and $\mathsf{F}_\beta = 0.8$ and $\underline{u}_\beta = 0.1$. The solid curve is the efficient frontier of the perfect information achievable set $\mathcal{U}$. The dashed lines represent the threshold $\underline{u}_\beta$, and the dotted-dashed curve represents the scaled set, $\mathsf{F}_\beta\mathcal{U}$. In panel (a) the shaded area represents the central region, $\hat{\mathcal{U}}_\beta$. In panel (b), the shaded area represents the boundary region within the lower triangle set L. In this case, $\underline{m} = \overline{m} = 0.5$ and $\mathbb{E}[v] - \mathbb{E}[v]\underline{u}_\beta/\underline{m} = 0.4$.

The main phase mechanism consists of allocation functions $\hat{\mathbf{p}}(\cdot|\mathbf{u}) : [0,\bar{v}]^n \to \mathcal{P}$ and future promise functions $\hat{\mathbf{w}}(\cdot|\mathbf{u}) : [0,\bar{v}]^n \to \mathsf{F}_\beta\mathcal{U}$ for all $\mathbf{u} \in \hat{\mathcal{U}}_\beta$, which satisfy (IC), (PK(u)), and $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}) \in \mathsf{F}_\beta\mathcal{U}$ for all $\mathbf{v}$. Next we explain them separately for any state $\mathbf{u}$ in the central region.

3.1. Allocation $\hat{\mathbf{p}}$

Note that for any point $\mathbf{u}$ on the efficient frontier $\mathcal{E}(\mathsf{F}_\beta\mathcal{U})$ of the set $\mathsf{F}_\beta\mathcal{U}$, the state $\mathbf{u}/\mathsf{F}_\beta$ is on the efficient frontier $\mathcal{E}(\mathcal{U})$. Recall from the discussion in the last section, with perfect information, promised utilities $\mathbf{u}/\mathsf{F}_\beta$ can be achieved with allocations that weigh values according to $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$ defined in (6). Motivated from this, for any state $\mathbf{u}$ in the central region $\hat{\mathcal{U}}_\beta$ defined in (8), we define the allocation function to be

$$\hat{p}_i(\mathbf{v}|\mathbf{u}) = \mathbf{1}\left\{\alpha_i^*(\mathbf{u}/\mathsf{F}_\beta)v_i \ge \max_{j\ne i} \alpha_j^*(\mathbf{u}/\mathsf{F}_\beta)v_j\right\}, \tag{9}$$

that is, the allocation function $\hat{\mathbf{p}}$ corresponds to allocating the resource to the agent with the largest weighted value, where the weights are the normal vector of the point closest to $\mathbf{u}$ in the efficient frontier of the scaled perfect information achievable set. Following Proposition 2.3, the expected utility of this allocation satisfies $\mathbb{E}_{\mathbf{v}}[v_i\hat{p}_i(\mathbf{v}|\mathbf{u})] = \frac{\partial\phi}{\partial\alpha_i}(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta))$—that is, the allocation delivers the utility of the point on the efficient frontier of the set $\mathcal{U}$ with normal $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$.

Remark 3.1. It is worth illustrating geometrically how the weights $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$ are determined. According to the optimality conditions for (6), for any state $\mathbf{u} \in \mathsf{F}_\beta\mathcal{U}$, we must have $\mathbf{u}/\mathsf{F}_\beta = \nabla\phi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)) - \mu\boldsymbol{e}$ for some nonnegative scalar μ, where $\boldsymbol{e}$ is the vector of 1s in $\mathbb{R}^n$. (Here, μ corresponds to the dual variable of the $\|x\|_1 = 1$ constraint.) Furthermore, as explained before, the vector $\mathsf{F}_\beta\nabla\phi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta))$ is on the efficient frontier of set $\mathsf{F}_\beta\mathcal{U}$. Therefore, as illustrated in Figure 2(a), the weights $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$ are determined by first projecting the state $\mathbf{u}$ along the "45° line" up to the point $\mathbf{u} + \mathsf{F}_\beta\mu\boldsymbol{e}$ on the efficient frontier of $\mathsf{F}_\beta\mathcal{U}$. Then, we scale this point by $1/\mathsf{F}_\beta$ to obtain a point $\mathbf{u}/\mathsf{F}_\beta + \mu\boldsymbol{e}$ on the efficient frontier of $\mathcal{U}$. Weights $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$ correspond to the normal vector of the hyperplane supporting the point $\mathbf{u}/\mathsf{F}_\beta + \mu\boldsymbol{e}$ on the efficient frontier of $\mathcal{U}$.

Note that the allocation (9) is not efficient, except when the weights $\alpha_i^*(\mathbf{u}/\mathsf{F}_\beta)$ are the same across agents or, equivalently, if the current state lies on the 45° line passing through the (scaled) welfare maximizing state $\mathsf{F}_\beta\mathbf{u}^*$. (As a reminder, $\mathbf{u}^*$ is the vector of agents' utilities under the efficient allocation.) Even if the initial state is set to $\mathsf{F}_\beta\mathbf{u}^*$, over time, state trajectories tend to drift

Figure 2. Examples of Weights and Future Promises of the Two-Agent Case

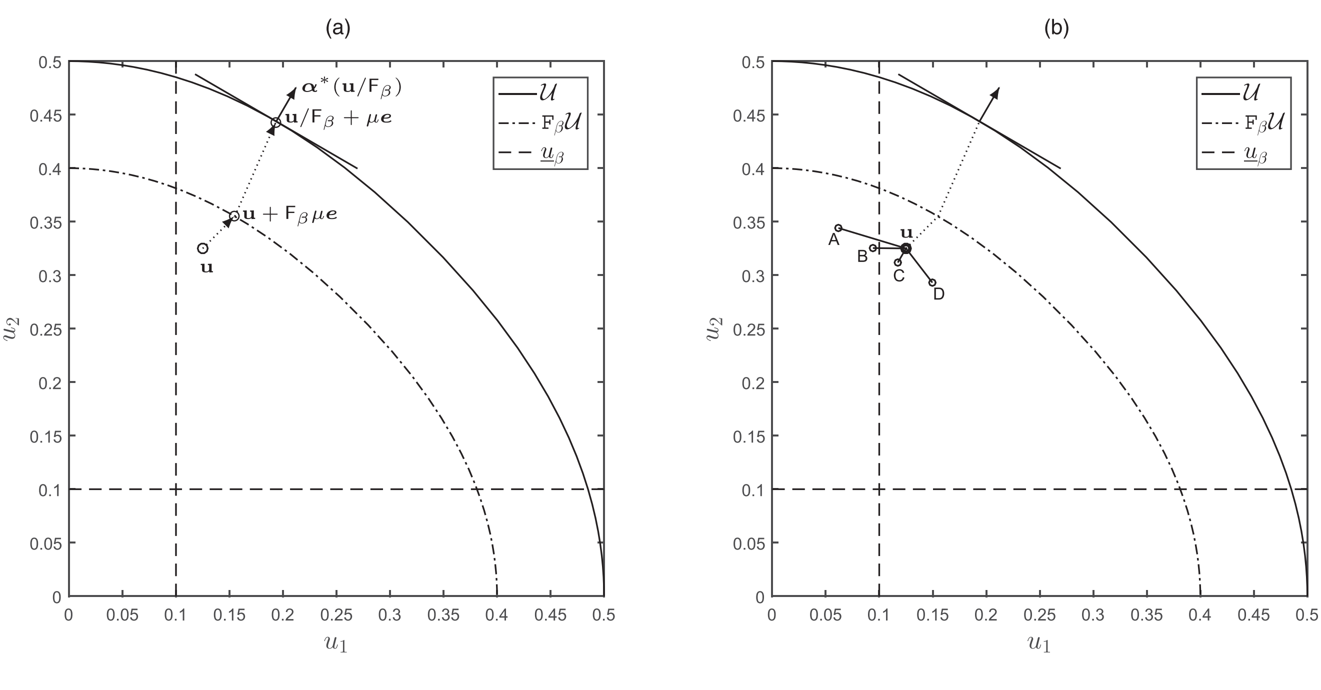

Illustration of the weight $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$ Illustration of the future promises $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$

Notes. Panel (a) illustrates the weight $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$ in a two-agent case with $\mathsf{F}_\beta = 0.8$, $\underline{u}_\beta = 0.1$, and $\mathbf{u} = (0.125, 0.325)$. In this case, $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta) = (0.3675, 0.6325)$, and $\mu = 0.03$. Panel (b) illustrates the future promises $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ given in (11) for some choices of $\mathbf{v}$ in a two-agent case (A, B, C, and D) with $\beta = 0.9$, $\mathsf{F}_\beta = 0.8$, $\underline{u}_\beta = 0.1$, $\mathbf{u} = (0.125, 0.325)$. In particular, the coordinates of the points A, B, C, and D correspond to $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ given in (11) for $\mathbf{v}$ taking values (1, 0), (0, 0), (1, 1), and (0, 1), respectively. (For the purpose of illustration, the value of F_β here is not calculated according to the formula presented in Section 4.)

away from this 45° line, because of the stochastic nature of private valuations. This introduces inefficiency, which can be interpreted as the information rent the designer has to pay to induce truthful revelation. Following our design of the future promise functions, as the time discount factor gets closer to 1, however, promised utility trajectories tend to concentrate around the 45° line for longer periods of time, resulting in more efficient allocations.

3.2. Future Promise Functions $\hat{\mathbf{w}}$

First, it is straightforward to determine the interim future promise function $\hat{W}_i(v_i|\mathbf{u}) \triangleq \mathbb{E}_{\mathbf{v}_{-i}}[\hat{w}_i(\mathbf{v}|\mathbf{u})]$ for each agent i from the incentive compatibility constraints. Following the standard argument of Myerson (1981), the (IC) constraints uniquely determine an expression for $\hat{W}_i(v_i|\mathbf{u})$ in terms of the interim allocation $\hat{P}_i(v_i|\mathbf{u}) \triangleq \mathbb{E}_{\mathbf{v}_{-i}}[\hat{p}_i(\mathbf{v}|\mathbf{u})]$ and the interim promise for the lowest type $\hat{W}_i(0|\mathbf{u})$. Further removing $\hat{W}_i(0|\mathbf{u})$ using the (PK(u)) constraints yields the following interim future promise function:

$$\hat{W}_i(v_i|\mathbf{u}) = \frac{1}{\beta}\left(u_i + (1-\beta)\left(\int_0^{v_i} \hat{P}_i(y|\mathbf{u})\mathrm{d}y - \hat{P}_i(v_i|\mathbf{u})v_i - \int_0^{\bar{v}} \bar{F}_i(y)\hat{P}_i(y|\mathbf{u})\mathrm{d}y\right)\right), \forall i. \tag{10}$$

Because the interim allocation is increasing in an agent's report, Equation (10) is necessary and sufficient to guarantee incentive compatibility.

Although the interim future promise functions $\hat{\mathbf{W}}$ in (10) are uniquely determined by $\hat{\mathbf{p}}$, the ex-post future promise functions $\hat{\mathbf{w}}$ may not be uniquely determined given $\hat{\mathbf{p}}$. This is because multiple ex-post future promise functions may correspond to the same interim future promise function. For example, we can set the ex-post future promise function $\hat{w}_i(v_i, \mathbf{v}_{-i}|\mathbf{u})$ to the interim future promise function $\hat{W}_i(v_i|\mathbf{u})$ for all $\mathbf{v}_{-i}$. This choice, however, does not guarantee that starting from an initial state $\mathbf{u} \in \mathsf{F}_\beta\mathcal{U}$, the future promises remain in $\mathsf{F}_\beta\mathcal{U}$ for all possible reports. In fact, if the state $\mathbf{u}$ is sufficiently close to the efficient frontier $\mathcal{E}(\mathsf{F}_\beta\mathcal{U})$, there exist values of $\mathbf{v}$ such that the future promise $\hat{\mathbf{W}}$ falls outside the set $\mathsf{F}_\beta\mathcal{U}$. If so, state $\mathbf{u}$ cannot be achieved using such a mechanism.

The key consideration, therefore, is that for any value $\mathbf{v} \in [0, \bar{v}]^n$ and state $\mathbf{u} \in \mathsf{F}_\beta\mathcal{U}$, we must ensure that the future promise function $\hat{\mathbf{w}}$ is feasible—that is, $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}) \in \mathsf{F}_\beta\mathcal{U}$, so as the set $\mathsf{F}_\beta\mathcal{U}$ is self-generating with respect to mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$. Specifically, the ex-post future promises need to be (i) nonnegative and (ii) within the efficient frontier of $\mathsf{F}_\beta\mathcal{U}$. In this section, we address the design issue (ii), and leave the issue (i) to Sections 4 and 7. In fact, the threshold $\underline{u}_\beta$ is designed to resolve issue (i).

To guarantee condition (ii), we implement interim future promises so as to minimize the risk of the next time period's promise utilities lying outside the achievable set $\mathsf{F}_\beta\mathcal{U}$. The design of the future promise functions resembles the d'Aspremont–Gérard-Varet–Arrow mechanism (Arrow 1979, d'Aspremont and Gérard-Varet 1979) and the risk-free transfers of Esö and Futo (1999). Specifically, the ex-post future promise function for any state $\mathbf{u}$ in the central region $\hat{\mathcal{U}}_\beta$ and value $\mathbf{v} \in [0,1]^n$ is given by

$$\hat{w}_i(\mathbf{v}|\mathbf{u}) = \hat{W}_i(v_i|\mathbf{u}) - \frac{1}{n-1}\sum_{j\neq i}\frac{\alpha_j^*(\mathbf{u}/\mathsf{F}_\beta)}{\alpha_i^*(\mathbf{u}/\mathsf{F}_\beta)} \cdot\left[\hat{W}_j(v_j|\mathbf{u}) - \mathbb{E}_{\tilde{v}_j}[\hat{W}_j(\tilde{v}_j|\mathbf{u})]\right]. \tag{11}$$

It is clear from expression (11) that the summation term has an expectation of 0, which is consistent with the first term, $\hat{W}_i(v_i|\mathbf{u})$, being the interim future promise. The following two propositions in the remainder of this section characterize important properties of the future promise functions as defined in (11).

In the context of a risk-averse seller in a static setting, Esö and Futo (1999) propose a similar transfer rule that gives rise to a constant ex-post revenue. Analogously, our choice of future promise functions guarantees that the weighted sum of the ex-post future promises is a constant. The following result shows that the future promises lie in a plane with normal $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$.

Proposition 3.1. *For any given state $\mathbf{u} \in \hat{\mathcal{U}}_\beta$, the future promise vector $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ lies within a plane in $\mathbb{R}^n$ for all $\mathbf{v}$. Specifically, the plane is described by the following equation,*

$$\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)^\top(\mathbf{u} - \hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})) = \frac{1-\beta}{\beta}\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)^\top (\nabla\phi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)) - \mathbf{u}) \geq 0.$$

The fact that our future promises lie on a plane resembles the concept of "enforceability with respect to hyperplanes" in Fudenberg et al. (1994). Furthermore, our plane being parallel to a support function of the (PI) achievable set echoes the construction of future promises given in the proof of lemma 6.1 in Fudenberg et al. (1994).

The properties stated in Proposition 3.1 are useful toward establishing that future promises $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ fall in the set $\mathsf{F}_\beta\mathcal{U}$. Recall that $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$ gives the normal of the supporting hyperplane of a point $\mathbf{u}$ on the efficient frontier of $\mathsf{F}_\beta\mathcal{U}$. Interpreting $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)^\top(\mathbf{u}-\mathbf{x})$ as the "signed directional distance" of a point $\mathbf{x} \in \mathbb{R}^n$ to the hyperplane supporting $\mathbf{u}$, we obtain that $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)^\top(\mathbf{u} - \hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}))$ measures how close the future promises are to the hyperplane supporting $\mathbf{u}$. Among all ex-post future promise functions that satisfy the interim future promise functions $\hat{W}_i(v_i|\mathbf{u})$, the one according to expression (11) is "pushed" the farthest from the efficient frontier $\mathcal{E}(\mathsf{F}_\beta\mathcal{U})$. This is formalized in the next result.

Proposition 3.2. *Fix a state* $\mathbf{u} \in \mathcal{E}(\hat{\mathcal{U}}_\beta)$*, and let* $\hat{W}_i(v_i|\mathbf{u})$ *be the interim future promise given in* (10)*. Then, the ex-post future promise* $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ *defined in* (11) *is an optimal solution of the following optimization problem:*

$$\begin{aligned} \max_{\tilde{\mathbf{w}}(\cdot)} \quad & \min_{\mathbf{v}} \boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)^\top(\mathbf{u} - \tilde{\mathbf{w}}(\mathbf{v})) \\ \text{s.t.} \quad & \mathbb{E}_{\mathbf{v}_{-i}}[\tilde{w}_i(v_i, \mathbf{v}_{-i})] = \hat{W}_i(v_i|\mathbf{u}) \quad \forall v_i, i, \end{aligned} \tag{12}$$

where $\boldsymbol{\alpha}^*(\cdot)$ *is defined in* (6).

Intuitively, setting future promises away from the efficient frontier is desirable because we need to ensure, for every report $\mathbf{v}$, that the ex-post future promise stays beneath the efficient frontier. Therefore, this choice of ex-post future promise allows us to satisfy the self-generating constraint in constructing the set $\mathsf{F}_\beta\mathcal{U}$.

Proposition 3.1 shows that the signed directional distance between $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ and the hyperplane supporting the projection of $\mathbf{u}$ onto $\mathcal{E}(\mathsf{F}_\beta\mathcal{U})$ is independent of $\mathbf{v}$ and strictly positive. This provides some indication that $\hat{\mathbf{w}}$ lies within the efficient frontier of $\mathsf{F}_\beta\mathcal{U}$. As in Fudenberg et al. (1994), we show that this is the case by carefully balancing the "step size" $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}) - \mathbf{u}$ against the curvature of the efficient frontier at the projection of $\mathbf{u}$ onto $\mathcal{E}(\mathsf{F}_\beta\mathcal{U})$. We formally establish this for the two-agent case in the next section.

Figure 2(b) demonstrates, in a two-agent case, the future promised function $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ starting from a state $\mathbf{u} \in \hat{\mathcal{U}}_\beta$. The figure shows that the realized future promise $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ ranges on a line when we vary values of $\mathbf{v}$. Moreover, we also observe that for some values of $\mathbf{v}$, the ex-post future promise $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ may fall out of the central region $\hat{\mathcal{U}}_\beta$ but remains in the set $\mathsf{F}_\beta\mathcal{U}$.

4. Boundary Region and Mechanism for the Two-Agent Case

In the previous section, we described the mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$ for the central region. To complete the description of the mechanism, we need to specify the mechanism for the "boundary region," or when any of the initial state u_i is below the threshold $\underline{u}_\beta$. As it turns out, the two-agent case is much simpler to describe compared with the general $n > 2$ case. In this section, we focus on the two-agent case. We generalize the analysis to larger n in Section 7.

We start by specifying the scalars F_β and $\underline{u}_\beta$ as follows:

$$\mathsf{F}_\beta = 1 - \frac{\underline{u}_\beta}{\underline{m}} \quad \text{and} \quad \underline{u}_\beta = \xi(1-\beta), \tag{13}$$

where ξ is a constant scalar independent of β and is given in Equation (C.8) of the online appendix. In fact, the threshold $\underline{u}_\beta$ is determined so that the efficient frontier $\mathcal{E}(\mathsf{F}_\beta\mathcal{U})$ intersects with axis i at the point $\mathsf{F}_\beta\mathbb{E}[v_i] = \mathbb{E}[v_i] - \mathbb{E}[v_i]\underline{u}_\beta/\underline{m}$. Therefore, for any $\mathbf{u}$ in the boundary region of $\mathsf{F}_\beta\mathcal{U}$ (that is, when either u_1 or u_2 is below the threshold $\underline{u}_\beta$), the state $\mathbf{u}$ must be within the lower triangle set L. See Figure 1(b) for an illustration.

We have already described the mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$ for the central region in the previous section. For a state $\mathbf{u}$ in the boundary region, because it is also in the lower triangle region, we define the mechanism to be, simply, random allocation. That is,

$$\begin{aligned} &\hat{p}_i(\mathbf{v}|\mathbf{u}) = p_i^{\mathrm{L}}(\mathbf{v}|\mathbf{u}) = u_i/\mathbb{E}[v_i], \quad \text{and} \\ &\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}) = \mathbf{w}^{\mathrm{L}}(\mathbf{v}|\mathbf{u}) = \mathbf{u}, \quad \text{if } \exists u_i < \underline{u}_\beta, \ i \in \{1,2\}. \end{aligned}$$

Following Remark 2.1, the mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$ satisfies (IC), (FA), (PK(u)) for $\mathbf{u}$ in the boundary region. Furthermore, it is obvious that future promises remain in the set $\mathsf{F}_\beta\mathcal{U}$ in this region. Therefore, the boundary region is self-generating with respect to mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$.

Now we have completed the description of mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$ for all initial state $\mathbf{u} \in \mathsf{F}_\beta\mathcal{U}$ and are ready to proceed with the claim that the set $\mathsf{F}_\beta\mathcal{U}$ is self-generating with respect to the mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$. What remains to be shown is that following the definitions of F_β and $\underline{u}_\beta$ in (13), the future promises indeed remain in set $\mathsf{F}_\beta\mathcal{U}$ according to the main phase mechanism. This is formally established in the following result.

Proposition 4.1. *Let* $\underline{\beta} \in (0,1)$ *be such that* $\mathsf{F}_{\underline{\beta}} \geq 0.5$*. For any* $\beta \geq \underline{\beta}$ *and initial state* $\mathbf{u} \in \mathsf{F}_\beta\mathcal{U}$*, the ex post future promise function* $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}) \in \mathsf{F}_\beta\mathcal{U}$*.*

The above key result of our paper relies crucially on the design of the constant ξ in (13), which determines the scaling factor F_β and the lower bound $\underline{u}_\beta$. First of all, the lower bound $\underline{u}_\beta$ has to be high enough so that the future promises $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$, defined in (11), remain positive. To achieve this we first provide an upper bound on $\|\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}) - \mathbf{u}\|_2$, which measures the distance between current and future promised utilities, in terms of β and other model parameters. Using this bound we show that future promises remain positive for all $\mathbf{u}$ with $u_i \geq \underline{u}_\beta$, when ξ is suitably chosen.

Second, and perhaps more importantly, we need to make sure that the boundary of the convex central region $\mathsf{F}_\beta\mathcal{U}$ (its efficient frontier) is "flat" enough so that starting from a point $\mathbf{u}$ very close to the efficient frontier, the next point $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ does not fall outside of it. Bounding the curvature of the efficient frontier $\mathcal{E}(\mathsf{F}_\beta\mathcal{U})$ is not easy. We instead work with the *signed distance function* between a point $\mathbf{u}$ and the convex set $\mathcal{U}$, which is defined as

$$I_{\mathcal{U}}(\mathbf{u}) = \max_{x:\|x\|_1=1, x\geq 0} \left\{x^\top \mathbf{u} - \phi(x)\right\}. \tag{14}$$

The signed distance $I_{\mathcal{U}}(\mathbf{u})$ measures the distance of a point to the efficient frontier of the PI set $\mathcal{U}$. It is positive if the point lies outside the set, 0 if on the efficient frontier of the set, and negative otherwise. Following (6), we obtain that $I_{\mathcal{U}}(\mathbf{u}) = \boldsymbol{\alpha}^*(\mathbf{u})^\top \mathbf{u} - \phi(\boldsymbol{\alpha}^*(\mathbf{u})) = \boldsymbol{\alpha}^*(\mathbf{u})^\top(\mathbf{u} - \nabla\phi(\boldsymbol{\alpha}^*(\mathbf{u})))$, because $\phi(\boldsymbol{\alpha}) = \nabla\phi(\boldsymbol{\alpha})^\top \boldsymbol{\alpha}$. Recall that $\nabla\phi(\boldsymbol{\alpha})$ corresponds to a point on the efficient frontier of the PI set with normal $\boldsymbol{\alpha}$. Therefore, the signed distance function measures the "directional distance" of a point $\mathbf{u}$ to the "closest" point on the efficient frontier.

The signed distance between any point $\mathbf{u}$ and the scaled set $\mathsf{F}_\beta\mathcal{U}$ is given by $I_{\mathsf{F}_\beta\mathcal{U}}(\mathbf{u}) = \mathsf{F}_\beta I_{\mathcal{U}}(\mathbf{u}/\mathsf{F}_\beta)$. We design the constant F_β so that $I_{\mathcal{U}}(\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})/\mathsf{F}_\beta) \le 0$ (the scalar F_β can be dropped because it is positive), which guarantees that future promises $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ lie in the scaled set $\mathsf{F}_\beta\mathcal{U}$. Here is some intuition using the shorthand notation $\hat{\mathbf{w}} = \hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$. Consider the following quadratic approximation of the signed distance function at the current state $\mathbf{u}/\mathsf{F}_\beta$,

$$I_{\mathcal{U}}\left(\frac{\hat{\mathbf{w}}}{\mathsf{F}_\beta}\right) \approx I_{\mathcal{U}}\left(\frac{\mathbf{u}}{\mathsf{F}_\beta}\right) + \nabla I_{\mathcal{U}}\left(\frac{\mathbf{u}}{\mathsf{F}_\beta}\right)^\top \left(\frac{\hat{\mathbf{w}} - \mathbf{u}}{\mathsf{F}_\beta}\right) + \frac{1}{2}\left(\frac{\hat{\mathbf{w}} - \mathbf{u}}{\mathsf{F}_\beta}\right)^\top \operatorname{Hess} I_{\mathcal{U}}\left(\frac{\mathbf{u}}{\mathsf{F}_\beta}\right)\left(\frac{\hat{\mathbf{w}} - \mathbf{u}}{\mathsf{F}_\beta}\right),$$

where $\operatorname{Hess} I_{\mathcal{U}}(\mathbf{u}) \in \mathbb{R}^{n\times n}$ represents the Hessian of function $I_{\mathcal{U}}$ evaluated at $\mathbf{u}$. The 0th-order term is nonpositive because the current state $\mathbf{u}$ lies in $\mathsf{F}_\beta\mathcal{U}$. The envelope theorem implies that $\nabla I_{\mathcal{U}}(\mathbf{u}) = \boldsymbol{\alpha}^*(\mathbf{u})$. Thus, the first-order term is negative and *independent* of $\hat{\mathbf{w}}$, according to Proposition 3.1. Because the signed distance function $I_{\mathcal{U}}$ is convex, the second-order term is nonnegative. We control the contribution of the second-order term by bounding the maximum eigenvalue of the Hessian matrix in terms of model parameters and using the aforementioned bound on $\|\hat{\mathbf{w}} - \mathbf{u}\|_2$. As it turns out, our design of the constant ξ allows us to ensure that the signed distance of $\hat{\mathbf{w}}$ is at most 0. Therefore, future promises $\hat{\mathbf{w}}$ fall below the efficient frontier $\mathscr{E}(\mathsf{F}_\beta\mathcal{U})$ starting from any point $\mathbf{u} \in \mathsf{F}_\beta\mathcal{U}$.

Proposition 4.1 implies that starting from any state $\mathbf{u} \in \mathsf{F}_\beta\mathcal{U}$, we can construct a sequence of allocation and future promises that delivers that initial state. That is, every state in $\mathsf{F}_\beta\mathcal{U}$ is achievable following this mechanism. In particular, consider the state that maximizes $u_1 + u_2$ in $\mathsf{F}_\beta\mathcal{U}$. The mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$ is able to achieve the social welfare given by

$$J_\beta \triangleq \max_{\mathbf{u}\in\mathsf{F}_\beta\mathcal{U}} (u_1 + u_2) = \mathsf{F}_\beta \max_{\mathbf{u}\in\mathcal{U}} (u_1 + u_2) = \mathsf{F}_\beta J^{\mathrm{FB}},$$

because $\mathbf{u} \in \mathsf{F}_\beta\mathcal{U}$ if and only if $\mathbf{u}/\mathsf{F}_\beta \in \mathcal{U}$. Because the set $\mathsf{F}_\beta\mathcal{U}$ is self-generating, it is a subset of $\mathcal{U}_\beta$, following Proposition 2.1. Therefore, the total social welfare satisfies $J_\beta \le J_\beta^* = \max_{\mathbf{u}\in\mathcal{U}_\beta}(u_1 + u_2)$. Overall, we have the following main theorem.[4]

Theorem 4.1. *Let $\underline{\beta} \in (0,1)$ be such that $\mathsf{F}_{\underline{\beta}} \ge 0.5$. For any $\beta \ge \underline{\beta}$, we have*

$$\mathsf{F}_\beta J^{\mathrm{FB}} = J_\beta \le J_\beta^* \le J^{\mathrm{FB}}.$$

Therefore, as β approaches 1, the relative gap in social welfare between our mechanism and first best converges to 0 at rate $(J^{\mathrm{FB}} - J_\beta)/J^{\mathrm{FB}} = 1 - \mathsf{F}_\beta = O(1-\beta)$.

It is clear that F_β as defined in (13) approaches 1 as β approaches 1. Therefore, the achievable set $\mathsf{F}_\beta\mathcal{U}$ approaches the perfect information set $\mathcal{U}$, and, correspondingly, the optimally achievable social welfare J_β^* approaches the first-best social welfare J^{FB}. In particular, the convergence rate of our mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$ to efficiency can be measured by their relative difference $1 - \mathsf{F}_\beta = O(1-\beta)$. Additionally, because $\mathsf{F}_\beta\mathcal{U}$ converges to the PI achievable set $\mathcal{U}$ as β approaches 1, every point in the PI achievable set is asymptotically achievable according to our mechanism.

Our mechanism provides a lower bound on the optimally achievable social welfare J_β^*, which, in theory, could converge to J^{FB} faster than $O(1-\beta)$. The following result indicates that the relative social welfare gap between any mechanism and first best cannot be smaller than $\Omega(1-\beta)$.

Theorem 4.2. *Suppose agents' values are i.i.d., then*

$$(J^{\mathrm{FB}} - J_\beta^*)/J^{\mathrm{FB}} \ge \Omega(1-\beta).$$

A typical approach used to provide upper bounds in dynamic mechanism design involves relaxing the IC constraint of all time periods except the first. While this approach works when values are discrete (see, e.g., Guo and Hörner 2015), it does not lead to an optimal convergence rate when values are continuous. Instead, we prove this result by considering a relaxation where the valuations of one agent are publicly observable for all periods (the valuations of the other agent are private). We formulate a dynamic programming problem to find the optimal social welfare under this relaxation and upper bound its objective value using the linear programming approach to approximate dynamic programming (see De Farias and Van Roy 2003). In particular, we impose a quadratic parametric form for the value function and optimize over one coefficient of the quadratic form. Denote by $V(\beta)$ the optimal objective value of the resulting optimization problem for any given β. Using the envelope theorem, we show that $V(\beta)$ is left differentiable around $\beta = 1$ and its left derivative $V'(1_-)$ is strictly positive. Using that $V(1) = J^{\mathrm{FB}}$, we obtain that $V(\beta) = J^{\mathrm{FB}} - (1-\beta)V'(1_-) + o(1-\beta)$, and the result follows because $V(\beta)$ is an the upper bound for J_β^*. We are hopeful that this method can be applied more broadly in the analysis of other dynamic mechanism design problems.

Now we have completed the description of our mechanism for the two-agent setting; in the next section, we provide a comparison with a mechanism motivated by the paper of Jackson and Sonnenschein (2007).

5. A Comparison with Jackson and Sonnenschein (2007)

Jackson and Sonnenschein (2007) (referred to as JS07 hereafter) propose a budget-based mechanism that allocates the resource to the agent with the highest report (ties are broken randomly) and, at the same time, restricts the number of times each agent can report a given type. JS07 is mainly focused on a finite-horizon static setting in which values are discrete and all private information is revealed to each agent at time 0. In comparison, we consider an infinite-horizon dynamic setting in which values are continuous and sequentially revealed over time. Despite some fundamental differences between our setting and that of JS07, we formally extend their simple design to our dynamic setting and analytically compare the performance of the resulting mechanism to first best. In short, we establish that while budget-based mechanisms are asymptotically efficient, the best relative inefficiency achievable with a budget-based mechanism is $\Omega((1-\beta)^{1/2})$.

It is worth explaining the setting and mechanism of JS07. The mechanism of JS07 critically depends on the number of time periods, τ_β, being finite and private information taking a finite number of possible types. In this setting, the JS07 mechanism sets a budget on the total number of times an agent can report a certain type. The budget of each type is set to be the probability of this type times τ_β, or the expected number of times that this type is realized over τ_β periods. Using this simple mechanism, an agent following an "approximately truthful strategy" (JS07, p. 252) receives an expected utility that converges to the utility obtained from the efficient allocation as τ_β goes to infinity. A remarkable feature of the mechanism from JS07 is that reporting as truthful as possible secures this level of utility regardless of the strategies followed by the other agents. Although JS07 does not explicitly characterize the equilibria of their mechanism, the latter security result implies that all equilibria are asymptotically efficient.

To extend the JS07 mechanism to our dynamic setting, we need to also assume a finite set of types for the private valuation. To reconcile the finite horizon in the budget-based mechanism with the infinite horizon nature of our setting, we divide the infinite time horizon into an infinite number of τ_β-period cycles, where the cycle length τ_β is a design parameter that depends on the discount factor. The budget-based mechanism is then used within each τ_β period cycle. We say that a mechanism $\boldsymbol{\pi}$ *secures efficient utility levels* if for each agent i we have $\lim_{\beta\uparrow 1}\sup_{\sigma_i}\inf_{\sigma_{-i}} V_i(\boldsymbol{\pi},(\sigma_i,\boldsymbol{\sigma}_{-i})) = u_i^*$. This maximin property implies that every agent has a sequence of strategies that guarantee convergence to the efficient utility levels regardless of the competitors' strategies. For the budget-based mechanism to secure efficient utility levels in an infinite-horizon setting, we need to carefully balance how the length of the cycle τ_β grows as β approaches 1.

Proposition 5.1. *The budget-based mechanism with cycles of length τ_β secures efficient utility levels if and only if*

$$\lim_{\beta\uparrow 1}\tau_\beta = \infty \quad \text{and} \quad \lim_{\beta\uparrow 1}\beta^{\tau_\beta} = 1. \tag{15}$$

Proposition 5.1 implies that the number of time periods in each cycle should be $\tau_\beta = o(1/(1-\beta))$ for the security result to hold. We prove the result in two steps. For the "if" part, we consider an approximately truthful reporting strategy under which an agent reports truthfully while the budget of the current type is still available, and randomly chooses another type to report if the budget of the current type has been exhausted. We show that any agent using these strategies can secure utility levels that converge to first best regardless of the strategy of the competitors, if the cycle length τ_β satisfies condition (15). By design, the mechanism guarantees that the empirical distribution of the competitor's reports over a cycle coincides with the distribution 1 would expect if the reports are truthful. When β^{τ_β} is close to 1, the net present values of all periods in a cycle are similar, and the order in which the competitor's reports appear is immaterial. Thus, by reporting truthfully whenever possible, the agent can achieve the efficient utility level in the limit. In this last step, we need to control the number of lies an agent is forced to make because budget constraints are hard and values are random. For the "only if" part, we consider a symmetric setting in which agents' values can be either high (= 1) or low (= 0) with equal probabilities. If a competitor bids high in the first $\tau_\beta/2$ periods, and low in the remaining periods, we show that no strategy secures an efficient utility level unless condition (15) holds. Intuitively, because of discounting, earlier periods have higher net present value and the agent is hurt the most when the competitor's report is high in the first periods. In the proof, we show that if β^{τ_β} does not converge to 1, the loss introduced in the first periods does not vanish, and the security result does not hold.

Finally, we establish that when cycle length τ_β satisfies condition (15), the relative gap from the budget-based mechanism's expected social welfare to first best is at least $\Omega((1-\beta)^{1/2})$ for every strategy profile. Formally, let $J^{\mathrm{JS}}_{\beta,\tau_\beta}(\sigma)$ be the expected social welfare of the budget-based mechanism with cycle length τ_β when agents employ strategy profile $\boldsymbol{\sigma}$.

Theorem 5.1. *If the length of each cycle τ_β satisfies condition (15) and agents have two types, then $\big(J^{\mathrm{FB}} - \sup_\sigma J^{\mathrm{JS}}_{\beta,\tau_\beta}(\sigma)\big)/J^{\mathrm{FB}} \ge \Omega((1-\beta)^{1/2})$.*

Inefficiencies are introduced whenever agents do not report their values. Agents are forced to lie whenever some type runs out of budget, which occurs with positive probability, because budgets are hard and values are random. Using concentration inequalities, we can show that an agent needs to lie at least $\Omega(\tau_\beta^{1/2})$ times in each cycle. Because of discounting, the best possible strategy is to lie near the end of the cycle, which introduces inefficiencies of order $\beta^{\tau_\beta}\tau_\beta^{1/2}$ per cycle. Summing over all possible cycles, we obtain that the total inefficiencies are at least $\Omega((1-\beta)^{1/2})$ under condition (15).

Proposition 5.1 and Theorem 5.1 together imply that the best possible rate of convergence to first best of a budget-based mechanism that secures efficient utility levels is $\Omega((1-\beta)^{1/2})$. This separates our mechanism with the budget-based mechanism extended from JS07. That is, in light of Theorem 4.1, our mechanism converges to first best at a faster rate than the budget-based mechanism as the discount factor β approaches 1 in the two-agent setting. (It is possible to show that if agents follow an approximately truthful reporting strategy, then the convergence rate to first best of a budget-based mechanism satisfying (15) is $O((1-\beta)^{1/2})$).

We continue with economic insights that can be derived from our mechanism in the next section, before presenting the analysis of our mechanism for the case of $n > 2$ in Section 7.

6. Economic Insights

In this section, we shed light on the economic insights derived from our mechanism. In particular, we explain, intuitively, why our main phase mechanism is dynamically incentive compatible (inducing agents to report truthfully) and approximately efficient (approaching the first-best social welfare as the discount factor converges to 1).

6.1. Incentive Compatibility

The main phase mechanism achieves incentive compatibility by introducing intertemporal substitution of consumption. That is, according to our mechanism, reporting a high value today reduces the chance of receiving the resource in the future. Consequently, if today's value is not as high, agents may forego today's allocation in view of more valuable future opportunities. This effect stems from the following results.

Proposition 6.1. *Agent i's allocation $\hat{p}_i(\mathbf{v}|\mathbf{u})$ is nondecreasing in the agent's report v_i (for fixed $\mathbf{v}_{-i}$ and $\mathbf{u}$), and nondecreasing in the agent's promised utility u_i (for fixed $\mathbf{v}$ and $\mathbf{u}_{-i}$).*

According to (9), in each period the allocation is determined by the comparison of weighted values among agents. The weights are determined by the promised utilities $\mathbf{u}$. Following Remark 3.1, it is clear that for any $\mathbf{u}$ in the central region, if agent i's promised utility is larger than agent j's (i.e., $u_i > u_j$), then agent i collects a higher utility via the allocation of the resource (i.e., $\mathbb{E}[v_i\hat{p}_i(\mathbf{v}|\mathbf{u})] > \mathbb{E}[v_j\hat{p}_j(\mathbf{v}|\mathbf{u})]$). Furthermore, (IC) implies that a higher valuation v_i increases agent i's chance of receiving the resource. Proposition 6.1 formalizes these intuitive features of our allocation rule.

Incentive compatibility implies that the interim future promise function of each agent is nonincreasing in his own report. That is, an agent's future promised utility tends to be lower if the current period report is higher. The following result characterizes the ex-post future promise function.

Proposition 6.2. *The future promise function $\hat{w}_i(\mathbf{v}|\mathbf{u})$ is nonincreasing in v_i and nondecreasing in v_j for $j \neq i$ (for fixed $\mathbf{u}$).*

As a result, reporting a higher value entails a higher chance of receiving the resource in the current period, at the expense of a lower future promised utility. This, in turn, translates to a lower chance of receiving the resource in the future. This provides an intuitive explanation on how the mechanism ensures that agents do not report higher than their true values.

Furthermore, it is interesting to see how other agents' reports affect a focal agent's future promise. Figure 2(b) illustrates such a point. Compare, for example, points A and C, or B and D. As agent 2's value increases from 0 to 1, agent 1's future promised utility w_1 increases. Intuitively, if another agent other than i reports a higher value, it decreases agent i's chance of receiving the resource in the current period. Agent i is then compensated with a higher future utility, to be fulfilled by future allocations.

6.2. Efficiency

We next discuss convergence to first best using an example in which agents' values are identically distributed. Recall that the main phase mechanism starts at the state $\mathsf{F}_\beta\mathbf{u}^*$, where $\mathbf{u}^*$ is the vector of agents' utilities under the efficient allocation. Therefore, the mechanism starts at the state with the largest component sum in the scaled set $\mathsf{F}_\beta\mathcal{U}$. At this state, the allocation is efficient. In fact, if a state $\mathbf{u}$ has equal components (i.e., $u_i = u_j$ for all i, j), then the allocation $\hat{\mathbf{p}}(\mathbf{v}|\mathbf{u})$ is always efficient, because the weights $\alpha_i^*(\mathbf{u}/\mathsf{F}_\beta)$ are the same for all i.

In Figure 3, we plot sample trajectories of promised utilities starting at state $\mathsf{F}_\beta\mathbf{u}^*$ following our mechanism. As we can see, the sample trajectories concentrate around the 45° line. Correspondingly, the weights $\alpha_i^*(\mathbf{u}/\mathsf{F}_\beta)$ of all agents tend to be close to each other along these trajectories. As a result, the resource is allocated almost efficiently, until it reaches the boundary region. In Figure 3(a), all trajectories reach the boundary region within 450 time periods.

As agents become more patient and the discount factor β increases, the step size between current promised utility $\mathbf{u}$ and future promises $\hat{\mathbf{w}}$ decreases (see

Figure 3. Sample Trajectories of Promised Utilities

(a) (b)

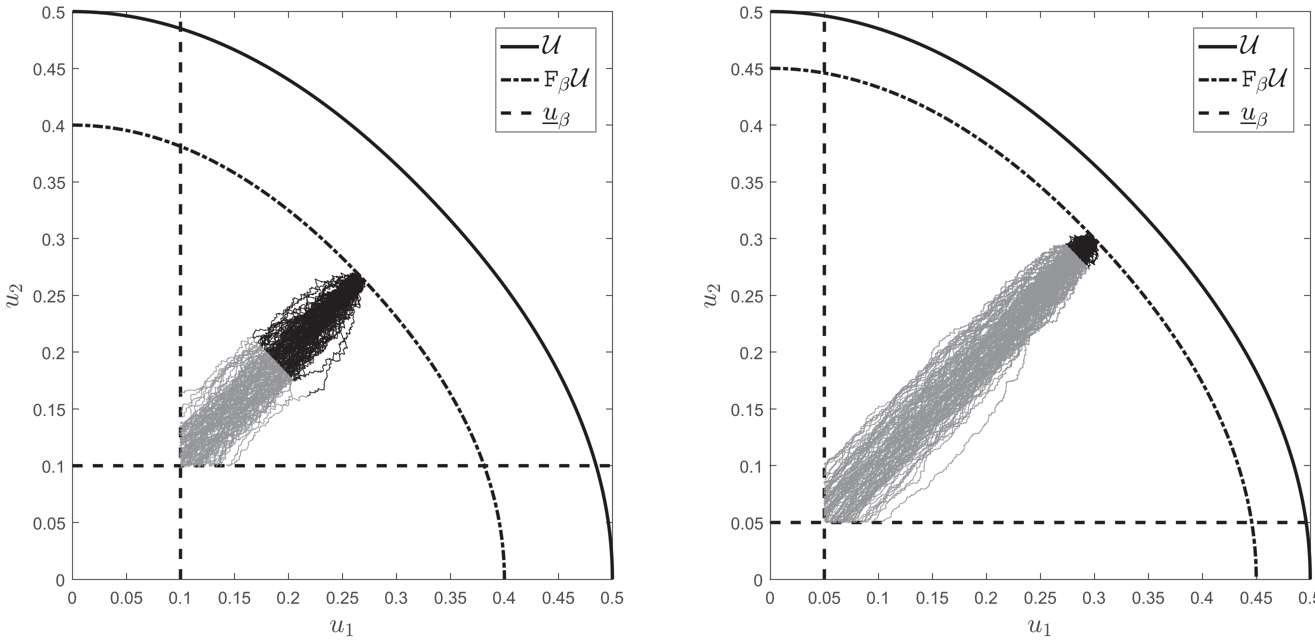

$\mathsf{F}_\beta = 0.8$, $\underline{u}_\beta = 0.1$ and $\beta = 0.9969$ $\mathsf{F}_\beta = 0.9$, $\underline{u}_\beta = 0.05$ $\beta = 0.9985$

Note. Each panel demonstrates 100 sample trajectories.

Lemma C.2 of the online appendix). In Figure 3, (a) and (b), the first 250 steps of each trajectory are marked black, while later steps are colored grey. As we can see from Figure 3(b), after 250 time periods, the promised utilities are still concentrated around the initial state, and it takes longer to reach the boundary region.

Figure 4 illustrates the expected social welfare per round generated by our mechanism, which is given by $\mathbb{E}_{\mathbf{u}_t,\mathbf{v}_t}[\sum_{i=1}^n v_{i,t}\hat{p}_i(\mathbf{v}_t|\mathbf{u}_t)]$, as a function of time. During the first time periods, the mechanism is in the main phase and the expected social welfare per round is very close to that of the efficient allocation. Over time, trajectories

Figure 4. Average Welfare of the Mechanism as a Function of Time

(a) (b)

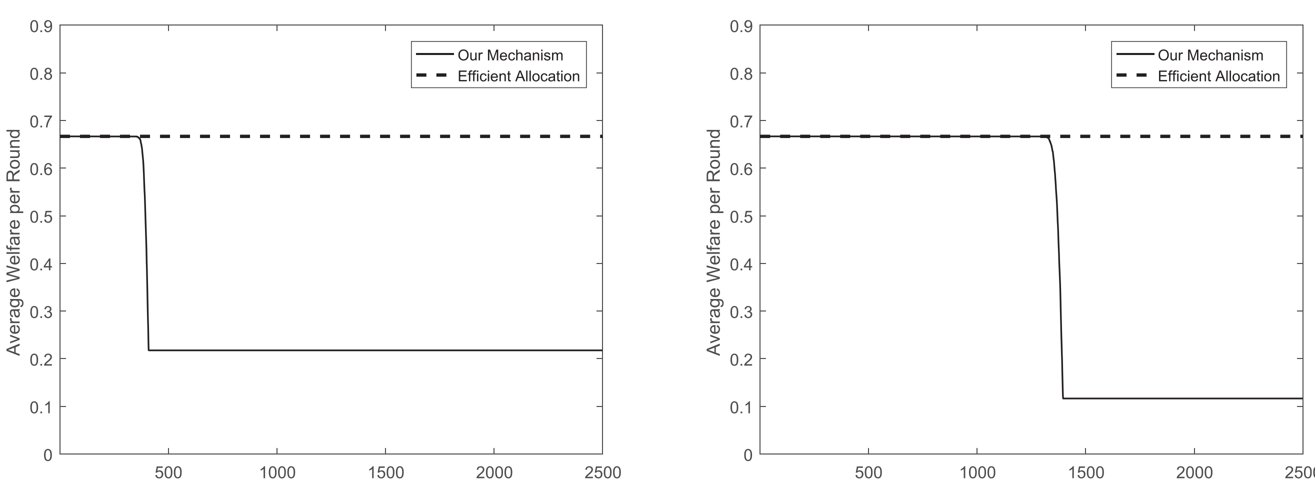

$\mathsf{F}_\beta = 0.8$, $\underline{u}_\beta = 0.1$ and $\beta = 0.9969$ $\mathsf{F}_\beta = 0.9$, $\underline{u}_\beta = 0.05$ and $\beta = 0.9985$

Notes. Main phase mechanism allocates efficiently until the state reaches the boundary region. In panel (a), the average time to reach the boundary region is 409, and in panel (b) it is 1,396.

drift to the boundary region where the expected social welfare per round drops significantly because there the allocation is highly inefficient. The boundary mechanism, albeit inefficient, is necessary to ensure that the incentive compatibility and promise-keeping constraints are sustained. Furthermore, as the discount factor increases, the mechanism remains in the main phase for more time periods. As we can see from Figure 4(a), it takes between 400 and 500 periods for all trajectories to reach the boundary region. In comparison, for a higher discount factor, Figure 4(b) shows that it takes between 1,000 and 1,500 time periods to reach the boundary region. As the time discount factor approaches 1, the boundary mechanism is pushed further into the future, and the mechanism allocates approximately efficiently for longer periods of time.

7. Generalization to *n* Agents

In this section, we generalize the analysis to settings with n agents, where $n > 2$. In this case, we need to carefully design the scaling factor and lower bound, and the boundary mechanism depending on the number n of agents involved. We first describe the difficulties that arise in the general case, and then provide our mechanism and analysis. To simplify the exposition, we assume that agents' values are identically distributed. Our results easily extend to the case of nonidentical distributions. We denote the p.d.f., the c.d.f., and the expected value of the agents' value distribution by $f(\cdot)$, $F(\cdot)$, and $\mathbb{E}[v]$, respectively.

First of all, the boundary mechanism for the two-agent setting no longer works for the general setting. When there are only two agents, in the boundary region we use the randomization mechanism $(\hat{\mathbf{p}}^{\mathrm{L}}, \hat{\mathbf{w}}^{\mathrm{L}})$. If $n > 2$, however, the boundary region is not always contained within the lower triangle set L anymore, as long as the lower bound for the n agent case, $\underline{u}_\beta^{(n)}$, approaches 0 with β approaching 1. As a result, the randomized allocation is no longer feasible for the boundary region.

Specifically, consider, for example, a three-agent setting as illustrated in Figure 5. Here, we focus on a situation where one agent's promised utility is below the threshold $\underline{u}_\beta^{(3)}$. In Figure 5(a), we plot the efficient frontier of the PI set $\mathscr{E}(\mathcal{U})$, its scaled version $\mathscr{E}(\mathsf{F}_\beta\mathcal{U})$, and the plane corresponding to $u_1 = 0.01 < \underline{u}_\beta^{(3)} = 0.0167$. Figure 5(b) demonstrates the intersections between the plane with efficient frontiers of sets $\mathcal{U}$, $\mathsf{F}_\beta\mathcal{U}$, and L, respectively. All points on the intersection of the plane and $\mathsf{F}_\beta\mathcal{U}$ lie in the boundary region. The state $\mathbf{u}$ represented by the circle, however, is outside the efficient frontier of the lower triangle set. At this state,

Figure 5. The Efficient Frontier and the Level Set in the Three-Agent Case

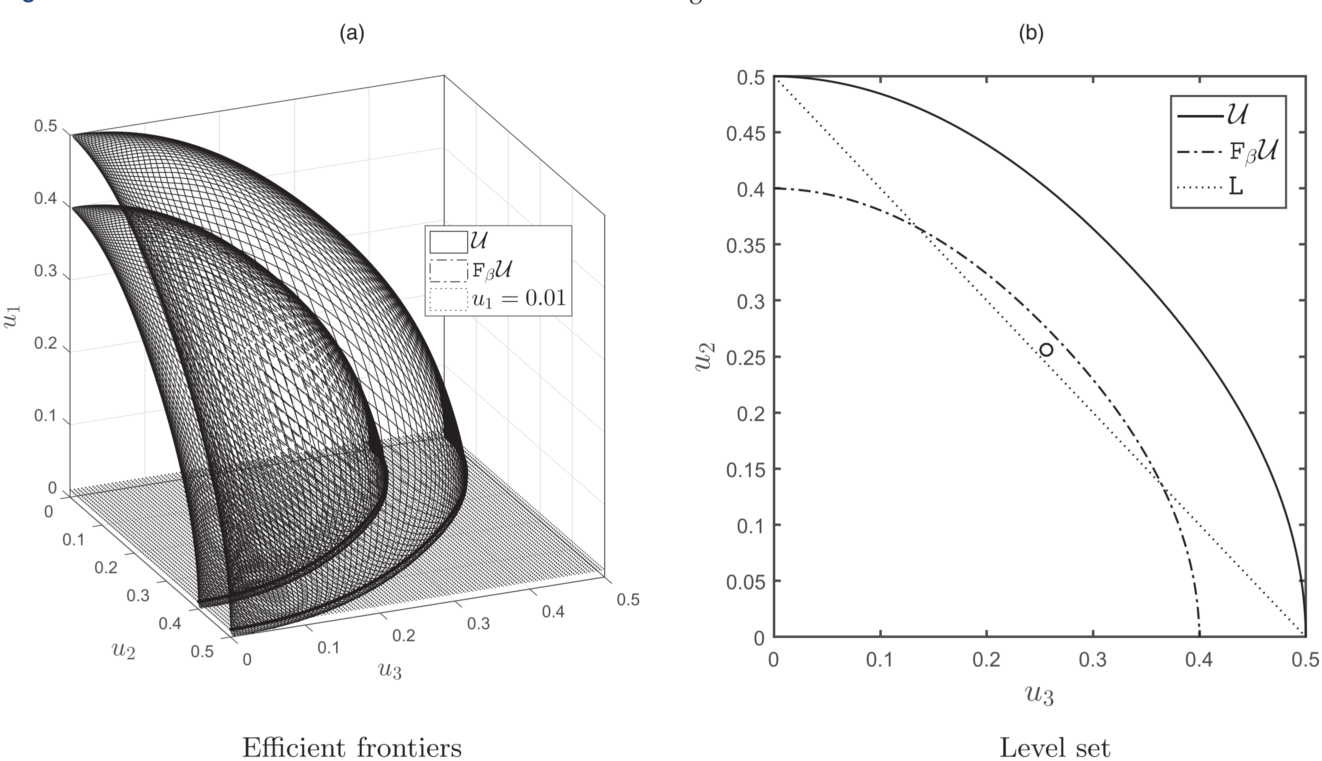

Notes. In panel (a), the outer surface with a solid mesh is $\mathscr{E}(\mathcal{U})$; the inner surface with a dotted-dashed mesh is the scaled efficient frontier $\mathscr{E}(\mathsf{F}_\beta\mathcal{U})$; and the horizontal plane represents the slice corresponding to $u_1 = 0.01$, which is lower than the threshold $\underline{u}_\beta^{(3)} = 0.0167$. Panel (b) plots the intersection of the surfaces in panel (a) with level set $u_1 = 0.01$ in the (u_2, u_3) space. In both panels, the value distributions are assumed to be uniform $[0, 1]$, $\mathsf{F}_\beta = 0.8$, and $\underline{u}_\beta^{(3)} = 0.0167$.

$u_1 + u_2 + u_3 > \mathbb{E}[v]$, which implies that the randomized allocation $\hat{\mathbf{p}}^{\mathrm{L}}$ is no longer feasible.

While the two-agent mechanism allocates the resource randomly when some agent's promise utility lies in the boundary region, the three-agent mechanism shall allocate randomly only to the agent with low promised utility and implement the two-agent mechanism for the other agents. Because the two-agent mechanism has been shown to attain every point in the two-agent PI achievable set as the discount factor increases, this construction allows us to attain points otherwise not achievable with random allocations.

More generally, suppose the main phase mechanism defined in Section 3 carries the promised utilities into a state in the boundary region, which means that at least some promised utility, say u_i, is below the threshold $\underline{u}_\beta^{(n)}$. By allocating the resource from this period on to agent i with probability $u_i/\mathbb{E}[v]$, we can guarantee that agent i's future promise $\hat{w}_i(\mathbf{u}|\mathbf{v})$ remains at u_i. As such, we convert the problem into one at a lower dimensional space. Therefore, when there are more than two agents, the boundary region mechanism can be defined recursively, depending on how many agents are still involved.

7.1. Mechanism

We refer to the agents with promise utilities above the threshold $\underline{u}_\beta^{(n)}$ as the *active agents*, in which $\underline{u}_\beta^{(n)}$ is defined as

$$\underline{u}_\beta^{(n)} = \xi^{(n)}(1-\beta)^{\frac{1}{n+4}}.$$

Here, $\xi^{(n)}$ is a constant scalar independent of β, provided in Definition F.2 of the online appendix.[5]

At any point in time, the allocation and future promise for an inactive agent i are, simply, $\hat{p}_i(\mathbf{v}|\mathbf{u}) = u_i/\mathbb{E}[v]$ and $\hat{w}_i(\mathbf{v}|\mathbf{u}) = u_i$, respectively. The mechanism for the active agents resembles the main phase mechanism described in Section 3.

Consider a case with k active agents. That is, k out of the n components in the state vector $\mathbf{u} \in \mathbb{R}^n$ are above the threshold $\underline{u}_\beta^{(n)}$. Denote by $\mathbf{u}^{(k)} \in \mathbb{R}^k$ the subvector of promised utilities for the active agents. Further define the total probability that the resource is allocated to an active agent to be

$$s(\mathbf{u}) = 1 - \sum_{i=1}^{n} \frac{u_i \mathbf{1}\{u_i < \underline{u}_\beta^{(n)}\}}{\mathbb{E}[v]}.$$

Consider the k dimensional PI set $\mathcal{U}^{(k)}$, the corresponding support function $\phi^{(k)}$, and weights $\boldsymbol{\alpha}^{(k)}$, as defined in (4), (5), and (6), respectively. (In these definitions, the state vector $\mathbf{u}$ corresponds to a k dimensional vector.) Similar to Section 3, we scale the PI set $\mathcal{U}^{(k)}$ with a factor $s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}$, in which $\mathsf{F}_\beta^{(k,n)}$ is defined as the following:

$$\mathsf{F}_\beta^{(k,n)} \triangleq 1 - \frac{n(k-1)\underline{u}_\beta^{(n)}}{\mathbb{E}[v]}.$$

Note that the scaling of the PI set needs to involve the factor $s(\mathbf{u})$, because with probability $1 - s(\mathbf{u})$ the resource is allocated to the inactive agents.

For the n agent case, the constants $\underline{u}_\beta$ and F_β in Section 3 correspond to $\underline{u}_\beta^{(n)}$ and $\mathsf{F}_\beta^{(n,n)}$, respectively. In the boundary region with k active agents, the allocation of our mechanism for an active agent i is defined similarly to the main phase mechanism (9) for the k agent case. That is,

$$p_i^{(k)}(\mathbf{v}|\mathbf{u}) = s(\mathbf{u})\mathbf{1}\left\{\alpha_i^{(k)}\left(\frac{\mathbf{u}^{(k)}}{s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}}\right)v_i \geq \max_{j\neq i} \alpha_j^{(k)}\left(\frac{\mathbf{u}^{(k)}}{s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}}\right)v_j\right\}. \tag{16}$$

The corresponding future utility is defined similar to (11), as

$$w_i^{(k)}(\mathbf{v}|\mathbf{u}) = W_i^{(k)}(v_i|\mathbf{u}) - \frac{1}{k-1}\sum_{j\neq i} \frac{\alpha_j^{(k)}\left(\frac{\mathbf{u}^{(k)}}{s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}}\right)}{\alpha_i^{(k)}\left(\frac{\mathbf{u}^{(k)}}{s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}}\right)} \cdot \left\{W_j^{(k)}(v_j|\mathbf{u}) - \mathbb{E}_{\tilde{v}_j}\left[W_j^{(k)}(\tilde{v}_j|\mathbf{u})\right]\right\}. \tag{17}$$

Here, the interim future promise function $W_i^{(k)}$ is as defined in (10) of Section 3, with the interim allocation function defined accordingly.

To summarize, at the beginning of each time period, the number of active agents k is updated to reflect the remaining number of active agents. Then, our mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$ is defined as the following. The allocation is given by

$$\hat{p}_i(\mathbf{v}|\mathbf{u}) = \begin{cases} p_i^{(k)}(\mathbf{v}|\mathbf{u}), & \text{if } u_i \geq \underline{u}_\beta^{(n)}, \\ u_i/\mathbb{E}[v], & \text{otherwise.} \end{cases} \tag{18}$$

and the future promise function is given by

$$\hat{w}_i(\mathbf{v}|\mathbf{u}) = \begin{cases} w_i^{(k)}(\mathbf{v}|\mathbf{u}), & \text{if } u_i \geq \underline{u}_\beta^{(n)}, \\ u_i, & \text{otherwise.} \end{cases} \tag{19}$$

7.2. Self-Generating Set

Note that when the number of agents n is larger than 2, the set $\mathsf{F}_\beta^{(n,n)}\mathcal{U}^{(n)}$ is not self-generating with respect to our mechanism. This is because as the number of active agents decreases to $k < n$, the PI achievable set $\mathcal{U}^{(k)}$ is *not* a scaled version of the intersection between the n-dimensional PI set and a subspace $\mathbb{R}^k$. In Figure 5(b), for example, the solid curve marks the boundary of the intersection between the three-dimensional PI achievable set $\mathcal{U}^{(3)}$ and the subspace $u_1 = 0.01$. However, this set is different from the efficient frontier of the two-dimensional PI achievable set $\mathcal{U}^{(2)}$, even with scaling. Consequently, when the number of active agents drops to $k<n$, there is a priori no guarantee that the promise utility lies in the set $s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathcal{U}^{(k)}$.

Therefore, we define the following set Ω_β, which we show to be self-generating with respect to our mechanism $(\hat{\mathbf{p}},\hat{\mathbf{w}})$. Some additional notations are in order. For a vector $\mathbf{u}\in\mathbb{R}^n$ and set $\kappa\subset\{1,\ldots,N\}$, we denote by $\mathbf{u}^{(\kappa)}$ the subvector corresponding to components of $\mathbf{u}$ in the index set κ. We denote the complement of κ by $\bar{\kappa}$. An inequality between a vector and a scalar means that each component of the vector satisfies this inequality. The set Ω_β is given by

$$\Omega_\beta=\bigcup_{\kappa\subseteq\{1,\ldots,N\}}\Big\{\mathbf{u}\in\mathbb{R}^n\,|\,\mathbf{u}^{(\kappa)}\geq \underline{u}_\beta^{(n)},\mathbf{u}^{(\bar{\kappa})}<\underline{u}_\beta^{(n)},\quad\text{and}\quad \mathbf{u}^{(\kappa)}\in s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathcal{U}^{(k)}\Big\}.$$

That is, for any state $\mathbf{u}$ in Ω_β with k active agents, the scaled PI achievable set $s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathcal{U}^{(k)}$ is contained in the Ω_β set. The following proposition confirms that our construction of the lower bounds ensures that Ω_β is indeed self-generating.

Proposition 7.1. *The set Ω_β is self-generating with respect to the mechanism $(\hat{\mathbf{p}},\hat{\mathbf{w}})$ defined in (18) and (19).*

To show that Ω_β set is self-generating with respect to our mechanism $(\hat{\mathbf{p}},\hat{\mathbf{w}})$, we need to argue that for all $\mathbf{u}$ in Ω_β, the mechanism given in (18) and (19) satisfies the conditions given in (3). The (IC), (FA), and (PK(u)) constraints follow by construction. The main step of the proof involves showing that future promises satisfy $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})\in\Omega_\beta$ for every report $\mathbf{v}$.

For any state $\mathbf{u}$ in Ω_β with k active agents, it is clear that future promised utilities of inactive agents remain in Ω_β. In the proof of Proposition 7.1, we first show that the subvector of future promises satisfies

$$\mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u})\in s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathcal{U}^{(k)}.\tag{20}$$

This step extends the geometric approach of Proposition 4.1 to higher dimensions by using the fact that the mechanism for active agents also satisfies the properties in Section 3. (In particular, in Online Appendix H, we prove extensions of Propositions 3.1 and 3.2 that account for the scaling factor.)

Condition (20) alone is not sufficient to establish our result because $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ could involve fewer active agents than $\mathbf{u}$. That is, in the next step the number of active agents may decrease to $k'<k$. We need to show the stronger result that the subvector $\mathbf{w}^{(k')}(\mathbf{v}|\mathbf{u})$, consisting of the components of $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ above the threshold $\underline{u}_\beta^{(n)}$, lies in $s(\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}))\mathsf{F}_\beta^{(k',n)}\mathcal{U}^{(k')}$. Therefore, we design the constant $\xi^{(n)}$ accordingly such that the intersection between $s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathcal{U}^{(k)}$ and the k' dimensional space of active agents is contained in $s(\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}))\mathsf{F}_\beta^{(k',n)}\mathcal{U}^{(k')}$. With this argument, (20) is sufficient to guarantee that all future promises lie in Ω_β.

Now we are ready to state the next theorem, which is the main result of this section.

Theorem 7.1. *There exists $\underline{\beta}\in(0,1)$ such that for any $\beta\geq\underline{\beta}$ the mechanism $(\hat{\mathbf{p}},\hat{\mathbf{w}})$ for n agents satisfies*

$$\mathsf{F}_\beta^{(n,n)}J^{\text{FB}}\leq J_\beta\leq J_\beta^*\leq J^{\text{FB}}.$$

Because the scaling factor $\mathsf{F}_\beta^{(n,n)}$ converges to 1 as β approaches 1, Theorem 7.1 implies that the maximum expected discounted social welfare achieved by the mechanism $(\hat{\mathbf{p}},\hat{\mathbf{w}})$ approaches first best. In particular the convergence rate is $1-\mathsf{F}_\beta^{(n,n)}=O\Big((1-\beta)^{\frac{1}{n+4}}\Big)$. Similar to the two-agent case, because $\underline{u}_\beta^{(n)}$ converges to 0 as β converges to 1, set Ω_β converges to the PI achievable set $\mathcal{U}$ as β approaches 1. Thus, every point in the PI achievable set is asymptotically achievable following our mechanism.

Note that the convergence rate to first best given in Theorem 7.1 for $n=2$ is slower than the one in Theorem 4.1. The proof for the two-agent case given in Section 4 leverages some special structure of the problem, which is not present in the general case. One may be able to provide better convergence rates by tightening the analysis. We leave this to future research.

In fact, the budget-based mechanism of JS07 for $n=2$ cannot be directly generalized to $n>2$ agents. The budget-based mechanism in the case of $n>2$ agents needs to be modified, and the analysis of the impact of this modification on the convergence rate appears to be nontrivial. Specifically, the modification needs to ensure that the approximately truthful reporting strategy secures efficient utility levels, to guarantee that the empirical distribution of reports is close to the true distribution for every set of $n-1$ players. Without this modification, approximately truthful reporting strategies are not guaranteed to be an equilibrium, because the security result does not hold. Therefore, we leave the analysis of the budget-based mechanism for the case of $n>2$ agents as a future research direction.

8. Conclusion and Extensions

In this paper, we study resource allocation with asymmetric information and no monetary transfer in a dynamic setting. The marginal cost for the resource is 0 in each period. Therefore, the mechanism designer focuses on allocation efficiency. We propose a mechanism that achieves asymptotic efficiency as the time discount factor approaches 1. From an algorithmic perspective, our mechanism is readily implementable. We provide the explicit suboptimality of our mechanism—that is, the rate of convergence to first best is $O(1-\beta)$ and $O\Big((1-\beta)^{\frac{1}{n+4}}\Big)$ for the two-agent case and the n-agent case, respectively.

Our analytical framework is focused on self-generating sets of agents' promised utilities. The essence of our approach is to establish a self-generating set with respect to our mechanism that expands as the discount factor increases, and eventually approaches the set of utilities achievable when values are publicly observable.

For challenging dynamic mechanism design problems, there has been previous work on constructing mechanisms that satisfy incentive compatibility constraints approximately. Our approach, on the other hand, constructs approximately optimal mechanisms that satisfy incentive compatibility constraints exactly, working closely with self-generating sets of future promises. We believe this approach can be applied in other dynamic mechanism/contract design settings, with or without monetary transfers.

There are a number of potential extensions to our work. For example, in our current mechanism, future promises are determined only through the interim (but not the ex-post) allocation. If we perceive promised utilities as money, our mechanism requires the planner to introduce lotteries, which may not be appealing in practice. It is, therefore, interesting to explore whether it is possible to asymptotically achieve efficiency with mechanisms in which future promises depend on the ex-post allocations. Such a mechanism, if it exists, establishes an indirect implementation without lotteries. That is, the agent who receives the resource in a period pays with future promises, while other agents are potentially compensated by higher future promises.

Along the line of thinking about ex-post versus interim promised utilities, we can also consider varying the incentive compatibility constraint. Currently, we enforce incentive compatibility at the "interim" level. That is, truthful reporting is each agent's best strategy, taking expectations with respect to other agents' values and assuming that competitors report truthfully. Alternatively, one can enforce certain "ex-post" incentive compatibility. That is, truthful reporting could be weakly dominant for every agent in each period regardless of other agents' reports (and assuming all agents report truthfully in the future). The optimal social welfare under ex-post incentive compatibility is less than or equal to that of our setting for any fixed discount factor, because there are more constraints in the definition of the self-generating set. It remains to be seen whether one can still establish asymptotic efficiency in this case.

Another extension is to consider a positive marginal production cost for the resource in each period. If the cost is positive, the mechanism designer needs to trade-off efficiency with production cost. In this case, the planner may have to withhold the resource if agents' valuations are low. In such a setting, merely studying the achievable set of promised utilities is no longer sufficient.

Acknowledgments

The authors thank Ozan Candogan, Xing Hu, Jacob Leshno, Giuseppe Lopomo, and Saša Pekeč for valuable comments. Furthermore, the authors thank the area editors, an anonymous associate editor, and three anonymous referees for providing a number of helpful comments that greatly improved the paper.

Endnotes

[1] We do not allow each agent's valuations to be correlated across periods, because the corresponding model would involve not only promised utility, but also threat utility, which may dramatically increase the dimensionality of the model (Fernandes and Phelan 2000). We need the support of the probability distribution to be bounded to ensure that the ex-post future promise function is inside the self-generating set. We keep the same discount factor for different agents for ease of exposition.

[2] Note that if some components of the allocation $\boldsymbol{\pi}_t$ are strictly between 0 and 1, we assume here that $\boldsymbol{\pi}_t$, and not the actual resource allocation in the end of the period, is publicly observable. More generally, the planner can decide what information to reveal to the agents in each period, and base the resource allocation rule on previous realized allocations. Given the type of results that we provide in this paper, it is clear that asymptotically the planner can do no better through such manipulations.

[3] In our setting, the probability of having a tie is 0.

[4] We say $f(\beta) = O(g(\beta))$ if and only if there exists $C > 0$ and $\beta_0 \in (0,1)$ such that $|f(\beta)| \leq C|g(\beta)|$ for all $\beta_0 \leq \beta < 1$. We say $f(\beta) = \Omega(g(\beta))$ if and only if there exists $C > 0$ and $\beta_0 \in (0,1)$ such that $|f(\beta)| \geq C|g(\beta)|$ for all $\beta_0 \leq \beta < 1$.

[5] Note that the values of the threshold $\underline{u}_\beta$ and the constant ξ for the two-agent mechanism provided in Section 4 are different from $\underline{u}_\beta^{(n)}$ and $\xi^{(n)}$—that is, $\underline{u}_\beta^{(n)} \neq \underline{u}_\beta$ and $\xi^{(n)} \neq \xi$—because we can provide stronger guarantees in the two-agent case.

References

Abreu D, Brooks B, Sannikov Y (2017) A "pencil-sharpening" algorithm for two player stochastic games with perfect monitoring. Working paper, Princeton University, Princeton, NJ.

Abreu D, Pearce D, Stacchetti E (1990) Toward a theory of discounted repeated games with imperfect monitoring, *Econometrica* 58(5): 1041–1063.

Arrow K (1979) The property rights doctrine and demand revelation under incomplete information, Boskin MJ, ed. *Economics and Human Welfare* (Academic Press, New York), 23–39.

Bergemann D, Said M (2011) *Dynamic Auctions*, *Wiley Encyclopedia of Operations Research and Management Science* (John Wiley & Sons, New York).

Bergemann D, Välimäki J (2010) The dynamic pivot mechanism. *Econometrica* 78(2):771–789.

d'Aspremont C, Gérard-Varet L-A (1979) Incentives and incomplete information. *J. Public Econom.* 11(1):25–45.

De Farias DP, Van Roy B (2003) The linear programming approach to approximate dynamic programming. *Oper. Res.* 51(6):850–865.

Esö P, Futo G (1999) Auction design with a risk averse seller. *Econom. Lett.* 65(1):71–74.

Fernandes A, Phelan C (2000) A recursive formulation for repeated agency with history dependence. *J. Econom. Theory* 2(91): 223–247.

Friedman EJ, Halpern JY, Kash I (2006) Efficiency and nash equilibria in a scrip system for p2p networks. *Proc. 7th ACM Conf. Electronic Commerce* (ACM, New York), 140–149.

Fudenberg D, Levine D, Maskin E (1994) The folk theorem with imperfect public information. *Econometrica* 62(5):997–1039.

Gershkov A, Moldovanu B (2010) Efficient sequential assignment with incomplete information. *Games Econom. Behav.* 68(1):144–154.

Gorokh A, Banerjee S, Iyer K (2016) Near-efficient allocation using artificial currency in repeated settings *Proc. 12th Conf. Web Internet Econom., Montreal*.

Guo M, Conitzer V, Reeves DM (2009) Competitive repeated allocation without payments. Leonardi S, ed. *Internet and Network Economics. WINE 2009*, Lecture Notes in Computer Science, vol. 5929 (Springer, Berlin), 244–255.

Guo Y, Hörner J (2015) Dynamic mechanisms without money. Technical report, Cowles Foundation for Research in Economics, Yale University, New Haven, CT.

Jackson MO, Sonnenschein HF (2007) Overcoming incentive constraints by linking decisions. *Econometrica* 75(1):241–257.

Johnson K, Simchi-Levi D, Sun P (2014) Analyzing scrip systems. *Oper. Res.* 62(3):524–534.

Johnson TR (2014) Dynamic mechanism design without transfers: Promises and confidentiality. Working paper, University of Notre Dame, Notre Dame, IN.

Judd KL, Yeltekin S, Conklin J (2003) Computing supergame equilibria. *Econometrica* 71(4):1239–1254.

Kash IA, Friedman EJ, Halpern JY (2007) Optimizing scrip systems: Efficiency, crashes, hoarders, and altruists. *Proc. 8th ACM Conf. Electronic Commerce* (ACM, New York), 305–315.

Kash IA, Friedman EJ, Halpern JY (2012) Optimizing scrip systems: Crashes, altruists, hoarders, sybils and collusion. *Distributed Comput.* 25(5):335–357.

Kash IA, Friedman EJ, Halpern JY (2015) An equilibrium analysis of scrip systems. *ACM Trans. Econom. Comput.* 3(3):13.

Li H, Zhang H, Fine C (2012) Dynamic business share allocation in a supply chain with competing suppliers. *Oper. Res.* 61(2):280–297.

Luenberger DG (1969) *Optimization by Vector Space Methods*, 1st ed. (John Wiley & Sons, Inc., New York).

Myerson R (1981) Optimal auction design. *Math. Oper. Res.* 6(1):58–73.

Nisan N, Roughgarden T, Tardos E, Vazirani VV (2007) *Algorithmic Game Theory*, vol. 1 (Cambridge University Press, Cambridge, UK).

Parkes DC, Singh S (2004) An MDP-based approach to online mechanism design. Thrun S, Saul LK, Schölkopf B, eds. *Advances in Neural Information Processing Systems*, vol. 17 (MIT Press, Cambridge, MA), 791–798.

Schneider R (2013) *Convex Bodies: The Brunn–Minkowski Theory*, vol. 151 (Cambridge University Press, Cambridge, UK).

Spear S, Srivastava S (1987) On repeated moral hazard with discounting. *Rev. Econom. Stud.* 54(4):599–617.

Thomas J, Worrall T (1990) Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *J. Econom. Theory* 51(2):367–390.

Vohra RV (2011) *Mechanism Design: A Linear Programming Approach* (Cambridge University Press, New York).

Yeltekin S, Cai Y, Judd KL (2017) Computing equilibria of dynamic games. *Oper. Res.* 65(2):337–356.

Zhang H (2012a) Analysis of a dynamic adverse selection model with asymptotic efficiency. *Math. Oper. Res.* 37(3):450–474.

Zhang H (2012b) Solving a dynamic adverse selection model through finite policy graphs. *Oper. Res.* 60(4):850–864.

Santiago R. Balseiro is an associate professor in the Decision, Risk, and Operations Division at the Graduate School of Business, Columbia University. His primary research interests are in the area of dynamic optimization, stochastic systems, and game theory with applications in revenue management and internet advertising.

Huseyin Gurkan is a PhD candidate in the Decision Sciences Area at the Fuqua School of Business, Duke University. He is interested in applying mathematical techniques to operational problems where system objectives are not aligned with incentives of self-interested, economic agents.

Peng Sun is a professor of decision sciences at the Fuqua School of Business, Duke University. His research interests include optimization under uncertainty and optimal design of contracts and mechanisms in dynamic environments.