Dynamic Mechanism Design without Money

by

Huseyin Gurkan

Business Administration
Duke University

Date: ________________________
Approved:

Peng Sun, Supervisor

Santiago R. Balseiro

Ozan Candogan

Ali Makhdoumi

Aleksandar Pekeč

Dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Business Administration
in the Graduate School of
Duke University

2019

ABSTRACT

Dynamic Mechanism Design without Money

by

Huseyin Gurkan

Business Administration
Duke University

Date: ______________________
Approved:

____
Peng Sun, Supervisor

____
Santiago R. Balseiro

____
Ozan Candogan

____
Ali Makhdoumi

____
Aleksandar Pekeč

An abstract of a dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Business Administration
in the Graduate School of
Duke University

2019

Abstract

In this dissertation, we study settings where a principal repeatedly determines the allocation of a single resource to i) a single agent, ii) one of two agents, and iii) one of n agents without monetary transfers over an infinite horizon with discounting. In all settings, the value of each agent for the resource in each period is private and the value distribution is common knowledge. For these settings, we design dynamic mechanisms that induce agents to report their values truthfully in each period via promises/threats of future favorable/unfavorable allocations. We show that our mechanisms asymptotically achieve the first-best efficient allocation (the welfare-maximizing allocation as if values are public) as the discount factor increases. Our results provide sharp characterizations of convergence rates to first best as a function of the discount factor.

In the single-agent setting, the principal incurs a positive cost from allocating the resource to the agent. We first consider the case in which the allocation cost is random in each period with a known distribution. Next, we extend the model such that the allocation cost follows one of two possible probability distributions. The principal and the agent share the same belief about the true cost distribution and update their beliefs in each period using Bayes' rule. In both cases, we provide mechanisms whose convergence rates are optimal, i.e., no other mechanism can converge faster to first best. In the settings with two or more agents, we do not consider allocation cost. We study the two-agent case before extending it to n agents. For two agents, we prove that the convergence rate of our mechanism is optimal. For n agents, we provide the convergence rate of our mechanism as a function of n.

Keywords: dynamic mechanism design, social efficiency, multi-agent games, resource allocation without money

Acknowledgements

I am very fortunate to have the support of many people in this journey.

I would like to express my deepest gratitude to my advisors, Santiago R. Balseiro and Peng Sun, for their excellent guidance, caring, patience, and providing me a perfect atmosphere for doing research. They both helped me at every moment when I needed. Right after being admitted to the Ph.D. program, I started working with Santiago. He taught me how to work tenaciously on challenging problems. During our discussions, I learned a lot from him and I am grateful for all that he has done. In the second year of the Ph.D. program, I started working with Peng. Since then, he became a great advisor to me. I learned many lessons from him that are not only related to research but also work discipline, dedication and the philosophy of life. Both Santiago and Peng always encouraged me to do more and gave inspiration to me in my academic career.

Moreover, I am greatly indebted to Ozan Candogan. His good advice and support, for which I am extremely grateful, have been invaluable for my academic and personal growth. Since the first year of the Ph.D. program, he never hesitated to guide me towards producing coherent and rigorous analysis. It is a privilege for me to work with him. I would also like to thank Sasa Pekec and Ali Makhdoumi, my committee members, for their insightful comments and recommendations. Besides, I am also grateful to Pino Lopomo for his comments that greatly improved the exposition of this study.

During my time at the Fuqua School of Business, the Decision Sciences area provided me the best conditions that a Ph.D. student may ask for. Therefore, I would like to thank Sasa one more time for playing a significant role at recruiting me as the graduate coordinator of the Decision Sciences area. Furthermore, I am grateful to the faculty in the Fuqua School of Business: Alessandro Arlotto, Alex Belloni, David Brown, Bora Keskin, Bob Nau, Jim Smith, Yehua Wei, Bob Winkler and Jiaming Xu–for their help at numerous occasions.

Before starting the Ph.D. program, I would not imagine that this journey will bring me the best friendship. We shared a lot and faced the hurdles of the Ph.D. study together. I would like to thank my friends: Levi, Xinchang, Soudipta, Anyi, Mingliu, Yunke, Chen, Ahmet, Asa, Janko for their support which always kept me motivated in really challenging conditions of the graduate study. Furthermore, I would like to specially thank to DeValve's and Murphy's for inviting me to their houses for the special meals as a member of their own families. I am very lucky to know you all.

My special thanks goes to Safak, Ezgi, Nadi; and my roommates Hasan and Oguz for supporting me to maintain a work/life balance in Durham. Their presence made me feel like at home. Moreover, I will always be grateful to my friends Ali, Bahadir, Busra, Gokhan, Merve, Murat, Nil, Nihal, Onursal, Yunus Emre, Zahid for their invaluable friendship. Despite the geographical distance, they always cheered me up.

Finally, my mother Aynur, my father Nail, and my little brother Onur have given me eternal support, love and encouragement; they were always there and stood by me through the good and bad times.

Contents

List of Figures

Chapter 1

Introduction

1.1 Dynamic Mechanism Design without Money

Mechanism design for resource allocation with asymmetric information have 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; 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 dissertation, 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. In Chapter 2, we study a setting with a single agent and the allocation of the resource is costly. In Chapters 3 and 4, there is no allocation cost, however, we study two-agent and n-agent settings, respectively. The objective of the planner is to maximize allocation efficiency, that is, the expected total discounted utilities from the resource minus the allocation cost, if there is any, 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, i) for the single-agent case, the resource is allocated to the agent whenever the value is larger than the allocation cost, and ii) for the multiple-agent case, the resource is allocated to the agent with the highest realized value in every period. First best can be achieved if 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 if this quantity is larger than the allocation cost. Repeated interactions, however, allow the planner to leverage future allocations when eliciting current period values, which may improve efficiency.

In this dissertation, we design mechanisms without monetary transfers which induce agents to truthfully reveal private information via promises/threats of future favorable/unfavorable allocations. Moreover, we show that our mechanisms asymptotically achieve 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, we take 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 one, 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.

1.2 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; Gershkov and Moldovanu 2010; Bergemann and Välimäki 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 zero. 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 zero. In the third chapter of this dissertation, 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. The study of Guo and Hörner (2015) is closely related to Chapter 2 of this dissertation. However, our model differs from theirs with a continuous value distribution and a random allocation cost. We provide a heuristic mechanism with a good performance while Guo and Hörner (2015) focus on characterizing the optimal dynamic mechanism within a more stylized model. In Chapters 3 and 4 of this dissertation, we study settings with multiple agents and continuous values without allocation cost. If the marginal cost of resource is zero, 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.

As another related study to ours, 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 one. 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 this dissertation, 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) builds upon 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 one. Fudenberg et al. (1994), however, does 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 non-trivial. As a matter of fact, computing equilibrium strategies in repeated games is a challenging problem (see, e.g., Judd et al. 2003; Yeltekin et al. 2017; Abreu 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.

Another stream of literature that is related to our work is the study of "scrip systems," which are non-monetary 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 decreases while others' increase. 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.

1.3 An Overview of Our Approach

Invoking the Revelation Principle, we focus on direct dynamic mechanisms in which allocations in each period depends on reported private values over time. In Chapter 2, a direct dynamic mechanism ensures that the agent reports the value truthfully in each period regardless of the past. In Chapters 3 and 4, 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 next 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.

In Chapter 3, we design an incentive compatible mechanism whose performance approaches the first-best social welfare, which is obtained by implementing the efficient allocation in each period, i.e., allocating the resource whenever the value is higher than the cost, as the discount factor approaches one. Our mechanism is mainly based on the key idea: use the perfect information allocation for a given promised utility whenever this allocation can be complemented by a feasible future promise function. Therefore, we first characterize a subset of the promised utilities where the perfect information allocation can be used. For the other promised utilities, we use simple allocate/no allocate mechanisms so as to ensure feasibility. Our heuristic mechanism asymptotically achieves the first-best social welfare because as the discount factor approaches one i) the allocation of our mechanism gets closer to the efficient allocation, and ii) the number of periods during which this near efficient allocation is implemented increases. Moreover, we show that our mechanism and results can be extended to the setting where the cost distribution is not known but the planner and the agent share the same initial belief.

In Chapters 3 and 4, 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 (In Chapter 2, this set corresponds to the interval between zero and the expected value of the resource.). 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 one, the "perfect information" (PI) achievable set, i.e., 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 one. Notably, using the linear programming approach to approximate dynamic programming we also show that this rate is tight, i.e., no other mechanism can converge faster to first best.

Chapter 2

Single Agent

2.1 Problem Definition and Model

We consider a discrete time infinite horizon setting where a social planner repeatedly allocates a single resource to a single agent in each period without relying on monetary transfers. We denote by v_t the random variable of the agent's (private) value for the resource in period $t \geq 1$. In the context where time period t is not important, we use a generic random variable v to represent the value of the agent v_t. The value of the agent in each period is independent and identically distributed with the cumulative distribution function $F(\cdot)$ and the density function $f(\cdot)$ over the support $[0, \bar{v}]$. The density function is bounded in the domain, i.e., $0 < \underline{f} \leq f(v) \leq \bar{f} < \infty$ for all $v \in [0, \bar{v}]$. The agent's value distribution is common knowledge. The planner incurs a positive cost $c \in \{c_l, c_h\}$ from allocating the resource. For the base case, the allocation cost c is independent and identically distributed in each period; c can take value c_h with probability q and c_l with probability $1 - q$. We assume that the distribution of the cost is common knowledge. Later, in Section 2.4 we extend the base case such that the cost distribution is not known. In period t, social welfare is given by the value of the agent minus the allocation cost, i.e., $v_t - c_t$ if the resource is allocated, and otherwise it is zero. The planner and the agent share the same discount factor $\beta \in (0, 1)$. We assume that the problem parameters satisfy the following assumption.

Assumption 2.1. *For any cost realization $\tilde{c} \in \{c_l, c_h\}$, we assume*

$$\beta > \max\left(\frac{\mathbb{E}[v1\{v \geq \tilde{c}\}]}{\mathbb{E}[v]}, 1 - \frac{\mathbb{E}[v1\{v \geq \tilde{c}\}]}{\bar{v}}, \frac{\bar{v}}{\bar{v} + \mathbb{E}[v]}\right).$$

The overall utility of the agent is given by the discounted sum of the valuations generated by the allocation of the resource across the horizon. The objective of the planner is to maximize the expected discounted sum of social welfare.

2.1.1 Dynamic Mechanisms

Following the Revelation Principle, we assume that the planner commits to an incentive compatible direct dynamic mechanism without loss of generality. That is, in each period, the agent learns her value, and reports 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 the agent's report and allocation in the end of the period.

Formally, a dynamic mechanism $\boldsymbol{\pi}$ is a sequence of allocation rules $\boldsymbol{\pi} = (\pi_t)_{t=1}^{\infty}$, where π_t is the probability that the resource is allocated to the agent in period t. We consider feasible allocation probabilities $0 \leq \pi_t \leq 1$ for all t. For any mechanism $\boldsymbol{\pi}$, the agent submits a report $\hat{v}_t \in [0, \bar{v}]$ in each period t and receives an allocation π_t probability. We define the history available at time $t > 1$, h_t as all reports and previous allocations and the cost realizations up to time t; and at time $t = 1$, $h_1 = \emptyset$. We denote by $\mathcal{H}_t$ the set of histories available in period t. We assume that the history is publicly observed. We say that a dynamic mechanism $\boldsymbol{\pi}$ is non anticipating if π_t depends only on the current period report $\hat{v}_t$, cost realization c_t and the history h_t for each period $t \geq 1$, i.e., $\pi_t : [0, \bar{v}] \times \{c_l, c_h\} \times \mathcal{H}_t \to [0, 1]$. A non-anticipating strategy $\boldsymbol{\sigma} = (\sigma_t)_{t=1}^{\infty}$ for the agent consists of reporting functions for each period that depend only on the value v_t of the agent and the public history h_t up to that period, i.e., $\sigma_t(v_t, h_t) = \hat{v}_t$.

Following the Revelation Principle, without loss of generality, we focus on direct mechanisms in which the agent reports her value truthfully to the planner. Therefore, we enforce the incentive compatibility constraints, which ensure that the agent is better off adopting the truthful reporting strategy than any other reporting strategy. We denote by $V_t(\boldsymbol{\pi}, \boldsymbol{\sigma} | v_t, h_t)$ the agent's utility-to-go in period t when i) the planner implements the mechanism $\boldsymbol{\pi}$, ii) the agent employs the strategy $\boldsymbol{\sigma}$, iii) the value of the agent is v_t and iv) the history is h_t. That is,

$$V_t(\boldsymbol{\pi}, \boldsymbol{\sigma} | v_t, h_t) \triangleq (1 - \beta) \mathbb{E}^{\boldsymbol{\pi}, \boldsymbol{\sigma}} \left[v_t \pi_t(\hat{v}_t, c_t, h_t) + \sum_{\ell = t+1}^{\infty} \beta^{\ell - t} v_\ell \pi_\ell(\hat{v}_\ell, c_\ell, \tilde{h}_\ell) | v_t, h_t \right],$$

where $\mathbb{E}^{\boldsymbol{\pi}, \boldsymbol{\sigma}}[- | v_t, h_t]$ represents the expectation with respect to histories $(\tilde{h}_\ell)_{\ell > t}$ induced by the mechanism $\boldsymbol{\pi}$ and the strategy $\boldsymbol{\sigma}$, given the value of the agent at time t is v_t and the history h_t.

For $\ell \geq t$, we denote by $\hat{v}_\ell = \sigma_\ell(v_\ell, \tilde{h}_\ell)$ the agent's reported value in period ℓ, in which history $\tilde{h}_\ell$ is recursively defined as $\tilde{h}_\ell = (\tilde{h}_{\ell-1}, (\hat{v}_{\ell-1}, c_{\ell-1}, \pi_{\ell-1}(\hat{v}_{\ell-1}, \tilde{h}_{\ell-1})))$ starting from $\tilde{h}_t = h_t$.

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

$$V_t(\boldsymbol{\pi}, \mathbf{I}|v_t, h_t) \geq V_t(\boldsymbol{\pi}, \boldsymbol{\sigma}|v_t, h_t)\,, \forall v_t, h_t, t \tag{2.1}$$

where $\mathbf{I}$ represents the truthful reporting strategy for the agent, i.e., $\mathrm{I}_t(v_t, h_t) = v_t$. The incentive compatibility constraints ensure that the agent is better off reporting truthfully, regardless of past reports and allocations. Given a direct mechanism $\boldsymbol{\pi}$, the discounted social welfare is

$$J_\beta^{\boldsymbol{\pi}} = (1-\beta)\mathbb{E}^{\boldsymbol{\pi}}\left[\sum_{t=1}^{\infty} \beta^{t-1}(v_t - c_t)\pi_t(v_t, c_t, h_t)\right].$$

where the expectation is taken with respect to the history induced by $\boldsymbol{\pi}$. The optimal social welfare J_β^* is found by solving the basic mechanism design problem.

$$J_\beta^* = \max_{\pi_t:[0,\bar{v}]\times\{c_l,c_h\}\times\mathcal{H}_t\to[0,1]} J_\beta^{\boldsymbol{\pi}} \quad \text{subject to (2.1)}\,. \tag{2.2}$$

When we relax constraint (2.1) in (2.2), the optimal solution of the relaxed problem becomes the efficient allocation. That is, the resource is allocated when the value is larger than the cost in all periods. We define first best as the social welfare obtained by the efficient allocation, and it is

$$\mathbb{E}[(v-c)^+]\,. \tag{2.3}$$

Here, we use $(\cdot)^+$ to represent projection at zero, i.e., $(x)^+ = \max(0, x)$.

2.1.2 Promised Utility Framework

We employ the promised utility framework to recursively formulate the mechanism design problem. Because the planner has commitment power and values and costs are independent across periods, we consider the expected discounted utility-to-go of the agent, $u_t = (1-\beta)\mathbb{E}^{\boldsymbol{\pi}}[\sum_{\ell=t}^{\infty} \beta^{\ell-t}\pi_t(v_t, c_t, h_t)v_t]$, as the state variable to guarantee dynamic incentive compatibility.

In the beginning of period t, the planner needs to fulfill the current state u_t as a promise to the agent 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 u_{t+1}. For the agent, the expected value of the current period allocation plus the next period's promised utility has to be equal to 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 u, a stage mechanism is given by an allocation function $p(\cdot, \cdot \,|\, u) : [0, \bar{v}] \times \{c_l, c_h\} \to [0, 1]$ and a future promise function $w(\cdot, \cdot \,|\, u) : [0, \bar{v}] \times \{c_l, c_h\} \to [0, \bar{w}]$, which map the agent's report and the observed allocation cost in the current period to an allocation and a next state promised utility, respectively. Here, $\bar{w} = \mathbb{E}[v]$ is an upper bound on the promised utility which corresponds to the total expected utility under the mechanism that allocates the resource to the agent in all periods. We further define functions $P(v \,|\, u) \triangleq \mathbb{E}_c[p(v, c \,|\, u)]$ and $W(v \,|\, u) \triangleq \mathbb{E}_c[w(v, c \,|\, u)]$ to be the interim allocation and future promise functions, respectively.

A stage mechanism (p, w) needs to satisfy the following constraints. First, the allocation function should be feasible. That is,

$$0 \le p(v, c \,|\, u) \le 1 \quad \forall v \in [0, \bar{v}], c \in \{c_l, c_h\}\,. \tag{FA-1}$$

Additionally, the mechanism satisfies the following promise keeping constraint,

$$u = \mathbb{E}[(1-\beta) v p(v, c \,|\, u) + \beta w(v, c \,|\, u)], \tag{PK(u)-1}$$

which guarantees that the promised utility u is fulfilled by the mechanism. The following incentive compatibility constraint imposes that, for the agent, reporting the value truthfully yields an expected utility at least as large as any other strategy.

$$(1-\beta) v P(v \,|\, u) + \beta W(v \,|\, u) \ge (1-\beta) v P(v' \,|\, u) + \beta W(v' \,|\, u), \quad \forall v, v' \in [0, \bar{v}]\,. \tag{IC-1}$$

Finally, the future promise function w is as follows.

$$0 \le w(v, c \,|\, u) \le \bar{w}, \quad \forall v \in [0, \bar{v}] \text{ and } c \in \{c_l, c_h\}\,. \tag{BC-1}$$

With some abuse of notation, let $J^*_\beta(u)$ be the optimal expected social welfare-to-go of the planner when the promised utility is $u \in [0, \bar{w}]$. We have the following dynamic programming formulation:

$$J^*_\beta(u) = \max_{p,w} \mathbb{E}[(1-\beta) p(v, c \,|\, u)(v - c) + \beta J^*_\beta(w(v, c \,|\, u))] \tag{DMDP}$$

$$\text{s.t. (FA-1), (PK(u)-1), (IC-1), (BC-1)}\,.$$

We denote by u^*_β the utility of the agent under an optimal mechanism, i.e., $u^*_\beta = \text{argmax}_{u \in [0, \bar{u}]} J^*_\beta(u)$. Therefore, the maximum social welfare J^*_β defined in (2.2) is equal to $J^*_\beta(u^*_\beta)$.

2.2 Approximately Optimal Mechanism

In this section we provide an approximate optimal mechanism. We first provide an upper bound for the value function $J^*_\beta(u)$. Next we construct a dynamic incentive compatible mechanism inspired by the upper bound and show that the performance of this mechanism converges to the upper bound as the discount rate converges to one.

2.2.1 Perfect Information Upper Bound

Consider a variant of problem (DMDP) in which we relax incentive compatibility constraints (IC-1) which means that the planner can observe the value of the agent in all periods, i.e., there is no information asymmetry between two players. We define this as the perfect information setting, and the maximum social welfare obtained in the perfect information setting as the first-best social welfare. Define $J^{\text{PI}}(u)$ to be the optimal expected social-welfare-to-go when the promised utility is $u \in [0, \bar{w}]$ in the perfect information setting. That is, $J^{\text{PI}}(u)$ is obtained by relaxing the (IC-1) constraint from the Bellman equation (DMDP). It is therefore clearly an upper bound of $J^*_\beta(u)$ for all $u \in [0, \bar{w}]$. In the following proposition, we formalize this result and provide an expression for $J^{\text{PI}}(u)$.

Proposition 2.1. *The optimal value function $J^*_\beta(u)$ is bounded by the perfect information value function $J^{\text{PI}}(u)$, i.e., $J^*_\beta(u) \le J^{\text{PI}}(u)$, where*

$$J^{\text{PI}}(u) = \min_{x \le 1} \mathbb{E}\left[(v(1-x) - c)^+\right] + xu\,.$$

Note that the perfect information upper bound $J^{\text{PI}}(u)$ is independent of the discount factor β. The next result indicates that the maximum of $J^{\text{PI}}(u)$ coincides with the first-best social welfare.

Corollary 2.1. *Let $u^{\text{PI}} \triangleq \operatorname*{argmax}_{u \in [0,\bar{w}]} J^{\text{PI}}(u)$, and $x^{\text{PI}}(u) \triangleq \operatorname*{argmin}_{x \le 1} \mathbb{E}\left[(v(1-x) - c)^+\right] + xu$. Then, we have*

- $u^{\text{PI}} = \mathbb{E}[v 1\{v \ge c\}]$ *and* $J^{\text{PI}}(u^{\text{PI}}) = \mathbb{E}[(v-c)^+]$, *and*
- $u = \mathbb{E}\left[v 1\{v \ge r^{\text{PI}}(u,c)\}\right]$ *where* $r^{\text{PI}}(u,c) \triangleq c/(1 - x^{\text{PI}}(u))$.

The perfect information upper bound suggests a dynamic mechanism $(p^{\text{PI}}, w^{\text{PI}})$ that allocates according to $\mathbf{1}\left\{v \ge r^{\text{PI}}(u,c)\right\}$ in all periods. That is, when the scaled value $v(1 - x^{\text{PI}}(u))$ is larger than the allocation cost c. The scaling factor $(1 - x^{\text{PI}}(u))$ depends on the promised utility u. If the promised utility u is equal to u^{PI}, then $x^{\text{PI}}(u)$ is zero and the corresponding allocation is the efficient allocation. A promised utility u larger than u^{PI} is delivered to the agent by scaling his value up by $(1 - x^{\text{PI}}(u))$, i.e., $x^{\text{PI}}(u) \leq 0$ for $u \geq u^{\text{PI}}$. Similarly, a promised utility u smaller than u^{PI} is delivered by scaling the value by $(1 - x^{\text{PI}}(u))$, i.e., $x^{\text{PI}}(u) \geq 0$ for $u \leq u^{\text{PI}}$. The corresponding promised utility remains the same, i.e., $w^{\text{PI}}(v, c \,|\, u) = u$. This dynamic mechanism, however, is not guaranteed to be dynamic incentive compatible (because the incentive compatibility constraints are relaxed).

2.2.2 Heuristic Mechanism

In this section, we provide a mechanism $(p^{\text{H}}, w^{\text{H}})$ that satisfies (FA-1), (PK(u)-1), (IC-1), and (BC-1) constraints. The allocation function p^{H} is based on the allocation function p^{PI} whenever it is feasible; and future promise function w^{H} is determined by (IC-1) and (PK(u)-1) constraints.

For any given allocation function $p(v, c \,|\, u)$, incentive compatibility constraints, (IC-1), and the promise keeping constraint, (PK(u)-1), dictate an interim future promise function $W(v \,|\, u)$. Using the envelope theorem, we derive an expression for the interim future promise function $W(v \,|\, u)$ in terms of $P(v \,|\, u)$ and $W(0 \,|\, u)$. Replacing this expression inside the promise keeping constraint, we obtain an interim future promise function $W(v \,|\, u)$ that depends only on the state u and the interim allocation function $P(v \,|\, u)$. Although the interim future promise function $W(v \,|\, u)$ is uniquely determined by p, the ex-post future promise function may not be uniquely determined given 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 $w(v, c \,|\, u)$ to the interim future promise function $W(v \,|\, u)$ for all c. This choice guarantees truthful reporting in the interim sense. By considering an alternative future promise function, we can obtain a stronger notion of incentive compatibility. That is, we ensure that the agent has no incentive to deviate from truthful reporting for any cost realization c. The ex-post incentive compatibility constraints are given by (2.5). We derive this future promise function by replacing $P(v \,|\, u)$ with $p(v, c \,|\, u)$ in $W(v \,|\, u)$ and the resulting expression is provided in the following proposition.

Proposition 2.2. *Let $p(v, c \,|\, u)$ be an arbitrary nondecreasing allocation of v for any $c \in \{c_l, c_h\}$ and $u \in [0, \bar{w}]$. Define the future promise function $w(v, c \,|\, u)$ as*

$$w(v, c \,|\, u) = \frac{u}{\beta} + \frac{(1-\beta)}{\beta} \left[\int_0^v p(x, c \,|\, u)\, dx - p(v, c \,|\, u) v - \int_0^{\bar{v}} \bar{F}(x) p(x, c \,|\, u)\, dx \right]. \tag{2.4}$$

The functions $p(v, c \mid u)$ *and* $w(v, c \mid u)$ *satisfy* (IC-1) *and* (PK(u)-1) *constraints as well as*

$$(1-\beta)vp(v,c \mid u)+\beta w(v,c \mid u) \geq (1-\beta)vp(v',c \mid u)+\beta w(v',c \mid u), \quad \forall v,v' \in [0,\bar{v}], c \in \{c_l, c_h\}\,. \tag{2.5}$$

We construct our heuristic mechanism using the perfect information allocation function p^{PI} and the corresponding future promise function given in (2.4). Although Proposition 2.2 provides a future promise function that satisfies (IC-1) and (PK(u)-1) constraints for any given nondecreasing allocation function, the resulting future promise function is not guaranteed to satisfy (BC-1) constraints, i.e., the future promise function may not be bounded by zero and $\bar{w}$. Therefore, we select the perfect information allocation p^{PI} and the corresponding future promise function given by (2.4) as our heuristic mechanism whenever (BC-1) constraints are satisfied. For the other cases, we complement this mechanism based on p^{PI} with simple allocation and no-allocation mechanisms so as to satisfy (BC-1) constraints. The next proposition formally provides our heuristic mechanism.

Proposition 2.3. *Let* $\bar{u} \triangleq \beta\mathbb{E}[v]$, $\underline{u} = (1-\beta)\bar{v}$ *and define* $(p^{\text{H}}, w^{\text{H}})$ *to be*

$$p^{\text{H}}(v,c \mid u) = \begin{cases} 1 & \bar{u} < u \leq \bar{w}\,, \\ \mathbf{1}\{v \geq r^{\text{PI}}(u,c)\} & \underline{u} \leq u \leq \bar{u}\,, \\ 0 & 0 \leq u < \underline{u}\,, \end{cases} \tag{2.6}$$

and

$$w^{\text{H}}(v,c \mid u) = \begin{cases} \frac{u}{\beta} - \frac{(1-\beta)\mathbb{E}_{\tilde{v}}[\tilde{v}]}{\beta} & \bar{u} < u \leq \bar{w}\,, \\ \frac{u}{\beta} - \frac{(1-\beta)}{\beta}\left[r^{\text{PI}}(u,c)\mathbf{1}\{v \geq r^{\text{PI}}(u,c)\} + \int_{r^{\text{PI}}(u,c)}^{\bar{v}} \bar{F}(y)\,dy\right] & \underline{u} \leq u \leq \bar{u}\,. \\ \frac{u}{\beta} & 0 \leq u < \underline{u}\,. \end{cases} \tag{2.7}$$

The mechanism $(p^{\text{H}}, w^{\text{H}})$ *satisfies* (FA-1), (PK(u)-1), (IC-1) *and* (BC-1) *constraints.*

The heuristic mechanism in Proposition 2.3 has three phases depending on the promised utility u. If the promised utility u is greater than $\bar{u}$, the resource is allocated to the agent regardless of the cost and the value. The future promise function is chosen so as to satisfy (PK(u)-1), (IC-1) constraints. Similarly, if the promised utility is lower than $\underline{u}$, the resource is not allocated and the future promise function is set to u/β. If the promised utility u is between the thresholds, i.e., $\bar{u} \geq u \geq \underline{u}$, then the heuristic mechanism uses the perfect information allocation p^{PI} and the future promise function is determined by (2.4). It is worth noting that as the discount factor β approaches one, $\bar{u}$ converges to $\bar{w}$ and $\underline{u}$ converges to zero. That is, the heuristic mechanism uses p^{PI} more frequently.

2.2.3 Asymptotic Optimality

In this section we show that the social welfare obtained by the dynamic mechanism proposed in Proposition 2.3 converges to the first-best social welfare as the discount factor converges to one. Let $J^{\text{H}}(u^{\text{PI}})$ be the expected performance of the heuristic mechanism when the initial state is u^{PI}.

$$J^{\text{H}}(u^{\text{PI}}) = \mathbb{E}\left[(1-\beta)\sum_{t=1}^{\infty}\beta^{t-1}(v_t - c_t)p^{\text{H}}(v_t, c_t \,|\, u_t^{\text{H}}) \,\Big|\, u_1^{\text{H}} = u^{\text{PI}}\right],$$

$$\text{where } u_{t+1}^{\text{H}} = w^{\text{H}}(v_t, c_t \,|\, u_t^{\text{H}}), \quad \forall t \geq 1.$$

The next result compares the performance of the heuristic mechanism to that of the optimal mechanism.

Theorem 2.1. *We have*

$$J^{\text{H}}(u^{\text{PI}}) \leq J^*_\beta(u^*_\beta) \leq J^{\text{PI}}(u^{\text{PI}}) \leq J^{\text{H}}(u^{\text{PI}}) + O(1-\beta). \tag{2.8}$$

The remaining part of this section is devoted to the proof of Theorem 2.1.

Let $\{u_t^{\text{H}}\}_{t=1}^{\infty}$ denote the stochastic process that governs the evolution of the state under the heuristic policy. This process evolves according to $u_t^{\text{H}} = w^{\text{H}}(v, c \,|\, u_{t-1}^{\text{H}})$ with initial condition $u_1^{\text{H}} = u^{\text{PI}} = \mathbb{E}[v\mathbf{1}\{v \geq c\}]$. Let $\tau = \inf\{t \geq 1 : u_t^{\text{H}} \notin [\underline{u}, \bar{u}]\}$ be the first time that the state falls outside the interval $[\underline{u}, \bar{u}]$.

The next result characterizes the evolution of the stochastic process $\{u_t^{\text{H}}\}_{t=1}^{\infty}$.

Lemma 2.1. *Under the heuristic policy* (2.6)-(2.7)*,*

1. *the process $\{u_t^{\text{H}}\}_{t=1}^{\infty}$ satisfies $\mathbb{E}[u_{t+1}^{\text{H}} \,|\, u_t^{\text{H}}] = u_t^{\text{H}}$ when $u_t^{\text{H}} \in [\underline{u}, \bar{u}]$.*
2. *$\mathbb{E}[|u_t^{\text{H}} - u^{\text{PI}}|^2 \mathbf{1}\{t \leq \tau\}] \leq K_1 t(1-\beta)^2$ for some constant $K_1 > 0$ that is independent of β.*
3. *the c.d.f. of the stopping time τ satisfies $\mathbb{P}(\tau < t) \leq K_2 t(1-\beta)^2$ for some constant $K_2 > 0$ independent of β.*
4. *the probability generating function of the stopping time τ satisfies $\mathbb{E}[\beta^\tau] \leq K_2(1-\beta)$.*

The proofs of Lemma 2.1 and Lemma 2.2 that comes next are presented in Appendix 6.1.

Let $H(x, c) = (1-\beta)\mathbb{E}\left[(v-c)\left(1\{v \geq c\} - 1\{v(1-x) \geq c\}\right)\right]$ be the difference in expected performance between the first-best allocation and the mechanism that weights the report of the agent with x. The next result provides deterministic bounds on the weight price $x^{\text{PI}}(u)$ and the function $H(x,c)$.

Lemma 2.2. *The following holds.*

1. *For all $u \in [\underline{u}, \bar{u}]$, we have*

$$\left| \frac{1}{1 - x^{\text{PI}}(u)} - 1 \right| \le K_3 |u - u^{\text{PI}}| \,,$$

for some constant K_3 independent of β.

2. *The function $H(x,c)$ satisfies for any $c \in \{c_l, c_h\}$:*

$$H(x,c) \le K_4(1-\beta) \left| 1 - \frac{1}{1-x} \right|^2 \,,$$

for some constant K_4 independent of β.

We are now ready to prove the main result. By construction, we have

$$J^{\text{H}}(u^{\text{PI}}) \le J^*_\beta(u^*_\beta) \le J^{\text{PI}}(u^{\text{PI}}) \,,$$

because $(w^{\text{H}}, p^{\text{H}})$ is feasible and $J^{\text{PI}}(u^{\text{PI}})$ is the largest upper bound (see Corollary 2.1). In the remainder of the proof we prove the last inequality of (2.8), that is

$$J^{\text{H}}(u^{\text{PI}}) \ge J^{\text{PI}}(u^{\text{PI}}) - O(1-\beta) \,.$$

Step 1: The expected performance of heuristic policy can be decomposed in as follows:

$$\begin{aligned} J^{\text{H}}(u^{\text{PI}}) &= (1-\beta) \sum_{t=1}^{\infty} \mathbb{E}[\beta^{t-1}(v_t - c_t) p^{\text{H}}(v_t, c_t \,|\, u_t^{\text{H}})] \\ &= \mathbb{E}[(v-c)^+] - (1-\beta) \sum_{t=1}^{\infty} \mathbb{E}\Big[\underbrace{\beta^{t-1}(v_t - c_t)\big(1\{v_t \ge c_t\} - p^{\text{H}}(v_t, c_t \,|\, u_t^{\text{H}})\big)}_{R_t} \Big] \\ &= J^{\text{PI}}(u^{\text{PI}}) - (1-\beta) \sum_{t=1}^{\infty} \mathbb{E}[R_t] \,, \end{aligned}$$

where the second equation follows because v_t's and c_t's are i.i.d. and using that $\sum_{t=0}^{\infty} \beta^{t-1} = 1/(1-\beta)$, the last equation follows because $J^{\text{PI}}(u^{\text{PI}}) = \mathbb{E}[(v-c)^+]$, from Corollary 2.1. The error terms R_t measure the differences in expected performances between the first-best allocation and the mechanism in consideration. We can decompose the error terms R_t as follows

$$\sum_{t=1}^{\infty} \mathbb{E}[R_t] = \underbrace{\sum_{t=1}^{\infty} \mathbb{E}[R_t \mathbf{1}\{t \le \tau\}]}_{R_1} + \underbrace{\sum_{t=1}^{\infty} \mathbb{E}[R_t \mathbf{1}\{t > \tau\}]}_{R_2} .$$

Next, we provide upper bounds for the error terms R_1 and R_2.

Step 2: Here we find an upper bound for the error term R_1. Because when $t \le \tau$ we have that $u_t^{\text{H}} \in [\underline{u}, \bar{u}]$, we use the bounds of Lemma 2.2 to obtain the following upper bound.

$$\begin{aligned} R_1 &\overset{(a)}{=} \sum_{t=1}^{\infty} \mathbb{E}\left[\beta^{t-1}(v_t - c_t)\left(\mathbf{1}\{v_t \ge c_t\} - \mathbf{1}\{v_t \ge r^{\text{PI}}(u_t^{\text{H}}, c_t)\}\right)\mathbf{1}\{t \le \tau\}\right] \\ &\overset{(b)}{=} \frac{1}{(1-\beta)} \sum_{t=1}^{\infty} \beta^{t-1} \mathbb{E}\left[H\left(x^{\text{PI}}(u_t^{\text{H}}), c_t\right) 1\{t \le \tau\}\right] \overset{(c)}{\le} K_4 \sum_{t=1}^{\infty} \beta^{t-1} \mathbb{E}\left[\left|\frac{1}{1 - x^{\text{PI}}(u_t^{\text{H}})} - 1\right|^2 \mathbf{1}\{t \le \tau\}\right] \\ &\overset{(d)}{\le} K_4 K_3^2 \sum_{t=1}^{\infty} \beta^{t-1} \mathbb{E}\left[|u_t^{\text{H}} - u^{\text{PI}}|^2 \mathbf{1}\{t \le \tau\}\right] \\ &\overset{(e)}{\le} K_4 K_3^2 K_1 (1-\beta)^2 \sum_{t=1}^{\infty} \beta^{t-1} t \overset{(f)}{=} K_4 K_3^2 K_1 , \end{aligned}$$

where (a) and (b) follow from the definition of R_1 and $H(\cdot,\cdot)$, (c) and (d) from Lemma 2.2, (e) from the second item of Lemma 2.1, and (f) from the fact that $\sum_{t=1}^{\infty} \beta^{t-1} t = 1/(1-\beta)^2$ because the series converges absolutely when $\beta \in [0, 1)$.

Step 3: We next bound the error term R_2 from above. We have

$$R_2 \overset{(i)}{\le} \bar{v} \sum_{t=1}^{\infty} \mathbb{E}[\beta^{t-1} 1\{t > \tau\}] \overset{(ii)}{=} \bar{v} \mathbb{E}\left[\sum_{t=\tau}^{\infty} \beta^t\right] \overset{(iii)}{=} \frac{\bar{v}}{1-\beta} \mathbb{E}\left[\beta^{\tau}\right] \overset{(iv)}{\le} K_2 \bar{v} ,$$

where (i) follows because $|R_t| \le \beta^{t-1}\bar{v}$ since $c \le \bar{v}$ and the allocation lies in $[0, 1]$, (ii) from Tonelli's Theorem, (iii) because the sum is a geometric series, and (iv) from the last item of Lemma 2.1.

Step 4: Putting everything together we obtain that

$$J^{\text{H}}(u^{\text{PI}}) \ge J^{\text{PI}}(u^{\text{PI}}) - (K_4 K_3^2 K_1 + K_2 \bar{v})(1-\beta) .$$

2.3 Economic Insights

In this section, we shed light on how heuristic mechanism works and the economic insights which can be derived from it. A key feature of the heuristic mechanism is that the agent reports her type truthfully to the social planner in each period. Moreover, as formally established in the previous section, the heuristic mechanism approximately achieves the first-best performance as the discount factor converges to one.

Why is the heuristic mechanism incentive compatible? First note that the allocation and the future promise of the heuristic mechanism do not change depending on the agent's report if $u_t^{\text{H}} \notin [\underline{u}, \bar{u}]$. Therefore, incentive compatibility is directly satisfied for those $u_t^{\text{H}} \notin [\underline{u}, \bar{u}]$. For the promised utility u_t^{H} such that $u_t^{\text{H}} \in [\underline{u}, \bar{u}]$, the heuristic mechanism maintains incentive compatibility by dynamically changing the scaling factor $x^{\text{PI}}(u_t^{\text{H}})$. To see this, note that the resource is allocated if the weighted report of the agent $v_t(1 - x^{\text{PI}}(u_t^{\text{H}}))$ is larger than the cost c_t. Here, a higher (lower) promised utility u_t^{H} corresponds to a lower (higher) scaling factor $x^{\text{PI}}(u_t^{\text{H}})$. In other words, a higher promised utility, is more favorable to the agent. Given this allocation rule, a larger report v_t increases the allocation probability in period t while resulting in a lower promised utility u_{t+1}^{H} and hence a lower allocation probability in period $t + 1$. This in turn creates a tradeoff for the agent which can be explained by intertemporal substitution effect. On one hand, the agent could report a high value to retain the resource now and decrease allocation probability in the next period. On the other hand, the agent could report a low value which potentially causes her to miss the resource but grant a higher future promise for the next period. Hence, there is no incentive for the agent to submit a higher (lower) report when her value is low (high) under the heuristic mechanism.

How does the heuristic mechanism achieve the first-best? The first-best performance is achieved by implementing the efficient allocation rule in each period that allocates the resource whenever the value of the agent is higher than the cost. However, this mechanism is not incentive compatible because the agent is always better off reporting a value higher than the maximum realization of the cost. On the other hand, the heuristic mechanism is incentive compatible. If the promised utility is $u_t^{\text{H}} \in [\underline{u}, \bar{u}]$, the heuristic mechanism dynamically determines its scaling factor $(1 - x^{\text{PI}}(u_t^{\text{H}}))$ so as to achieve incentive compatibility. In the other cases, the resource is allocated if u_t^{H} is higher than $\bar{u}$ and not allocated if u_t^{H} is lower than $\underline{u}$ regardless of the report of the agent in order to achieve incentive compatibility. The heuristic mechanism is clearly not efficient when $u_t^{\text{H}} \notin [\underline{u}, \bar{u}]$. Moreover, if $u_t^{\text{H}} \in [\underline{u}, \bar{u}]$, we use a scaling factor $(1 - x^{\text{PI}}(u_t^{\text{H}}))$ to weight the value of the agent, and the allocation is not efficient whenever $(1 - x^{\text{PI}}(u_t^{\text{H}})) \neq 1$. Despite these facts, the heuristic mechanism asymptotically achieves the first-best performance because the promised utility u_t^{H} remains close to the initial state u^{PI}, at which $x^{\text{PI}}(u^{\text{PI}}) = 0$, as the discount factor β increases (see Lemma 2.1). Being closer to u^{PI} implies that i) $(1 - x^{\text{PI}}(u_t^{\text{H}}))$ deviates around one, i.e, the allocation is close to efficient allocation, and ii) it takes longer for the promised utility u_t^{H} to jump outside $[\underline{u}, \bar{u}]$. In other words, as the discount factor β approaches one, the number of periods during which the resource allocation tends to be close to efficient allocation increases, and the heuristic mechanism approximately achieves the total expected discounted social welfare of first best.

2.4 Unknown Cost Distribution

In this section, we extend the base model of Section 2.1 to a setting where the probability distribution of the cost c is unknown to the planner and the agent. Specifically, the realization of the cost c can be either c_h or c_l, i.e., $c \in \{c_h, c_l\}$, and the probability of c_h is either q_1 or q_2, both satisfying Assumption 2.1. At the beginning of the first period, the planner and the agent share the same belief about the probability distribution. We use b_t to denote the probability that c_h appears with probability q_1 in period t. In each period, the planner and the agent update the belief using the Bayes' rule after observing the realized cost.

In order to recursively formulate the mechanism design problem, we introduce the belief as the second dimension of the state space after the promised utility. Therefore, for a generic promised utility u and belief b, a stage mechanism consists of an allocation function $p(\cdot, \cdot \,|\, u, b) : [0, \bar{v}] \times \{c_l, c_h\} \to [0, 1]$ and a future promise function $w(\cdot, \cdot \,|\, u, b) : [0, \bar{v}] \times \{c_l, c_h\} \to [0, \bar{w}]$. We denote by $P(v \,|\, u, b) \triangleq \mathbb{E}_c[p(v, c \,|\, u, b) \,|\, b]$ and $W(v \,|\, u, b) \triangleq \mathbb{E}_c[w(v, c \,|\, u, b) \,|\, b]$ the interim allocation and future promise functions which are obtained by taking expectation of p and w functions over cost c, respectively. Here, we use the notation $\mathbb{E}_c[\cdot \,|\, b]$ to represent the expectation over the cost c for a given belief b with some abuse of notation.

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

$$0 \leq p(v, c \,|\, u, b) \leq 1 \quad \forall v \in [0, \bar{v}], c \in \{c_l, c_h\}\,. \tag{FA(b)-1}$$

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

$$u = \mathbb{E}[(1-\beta)vp(v,c\,|\,u,b) + \beta w(v,c\,|\,u,b)\,|\,b], \tag{PK(u,b)-1}$$

which guarantees that the promised utility u is fulfilled by the mechanism with respect to the current belief. An agent should not have an incentive to misreport the true type. The following incentive compatibility constraint imposes that, for the agent, reporting the value truthfully yields an expected utility at least as large as any other strategy. The incentive compatibility constraints are

$$(1-\beta)vP(v\,|\,u,b) + \beta W(v\,|\,u,b) \geq (1-\beta)vP(v'\,|\,u,b) + \beta W(v'\,|\,u,b), \quad \forall v,v' \in [0,\bar{v}]\,. \tag{IC(b)-1}$$

Finally, the future promise function w should be bounded as follows.

$$0 \leq w(v,c\,|\,u,b) \leq \bar{w}, \quad \forall v \in [0,\bar{v}] \text{ and } c \in \{c_l, c_h\}\,. \tag{BC(b)-1}$$

Let $J^*_\beta(u,b)$ be the optimal expected social welfare-to-go of the planner when the promised utility is $u \in [0,\bar{w}]$ and the belief is b. The corresponding dynamic programming formulation is

$$\begin{aligned} J^*_\beta(u,b) = \max_{p,w} \mathbb{E}[(1-\beta)p(v,c\,|\,u,b)(v-c) + \beta J^*_\beta(w(v,c\,|\,u,b), \mathsf{B}(b,c))\,|\,b] \qquad &\text{(DMDP(b))} \\ \text{s.t. (FA(b)-1), (PK(u,b)-1), (IC(b)-1), (BC(b)-1)} \qquad & \end{aligned}$$

where $\mathsf{B}(b,c)$ represents the posterior belief after observing the realized cost c from prior b. Specifically, we have

$$\mathsf{B}(b,c) = \frac{q_1 b}{q_1 b + q_2(1-b)}\mathbf{1}\{c = c_h\} + \frac{(1-q_1)b}{(1-q_1)b + (1-q_2)(1-b)}\mathbf{1}\{c = c_l\}$$

Assume that the initial prior belief is b_1. We define u^*_{β,b_1} to be the utility of the agent under an optimal mechanism with the initial belief b_1, i.e., $u^*_{\beta,b_1} = \text{argmax}_{u\in[0,\bar{u}]} J^*_\beta(u,b_1)$. Therefore, the maximum social welfare which can be obtained following a direct mechanism is $J^*_\beta(u^*_{\beta,b_1}, b_1)$.

2.4.1 Perfect Information Upper Bound

Similar to Section 2.2.1, we consider the perfect information upper bound by relaxing the incentive compatibility constraints, which corresponds to assuming that the planner observes the value of the agent in all periods. That is, there is no information asymmetry between the agent and the planner.

Let $J^{\text{PI}}(u,b)$ be the optimal expected social-welfare-to-go when the promised utility is $u \in [0,\bar{w}]$

and the belief is $b \in [0,1]$ in the perfect information setting. Because $J^{\text{PI}}(u,b)$ is obtained by removing constraints (IC(b)-1) from the Bellman equation of $J^*_\beta(u,b)$, it is clearly an upper bound, i.e., $J^{\text{PI}}(u,b) \geq J^*_\beta(u,b)$ for all $u \in [0,\bar{w}]$ and $b \in [0,1]$. In the following proposition, we formalize this result and provide an expression for $J^{\text{PI}}(u,b)$.

Proposition 2.4. *The optimal value function $J^*_\beta(u,b)$ is bounded from above by the perfect information value function $J^{\text{PI}}(u,b)$, i.e., $J^*_\beta(u,b) \leq J^{\text{PI}}(u,b)$, where*

$$J^{\text{PI}}(u,b) = \min_{x \leq 1} \mathbb{E}\left[(v(1-x)-c)^+ \mid b\right] + xu\,. \tag{2.9}$$

Note that the perfect information upper bound $J^{\text{PI}}(u,b)$ is independent of the discount factor β. Optimizing over the initial state, we obtain that the maximum of $J^{\text{PI}}(u,b)$ coincides with the first-best social welfare.

Corollary 2.2. *Let $u^{\text{PI}}(b) \triangleq \operatorname{argmax}_{u\in[0,\bar{w}]} J^{\text{PI}}(u,b)$, and $x^{\text{PI}}(u,b) \triangleq \operatorname{argmin}_{x\leq 1} \mathbb{E}\left[(v(1-x)-c)^+ \mid b\right] + xu$. Then, we have*

- $u^{\text{PI}}(b) = \mathbb{E}[v\mathbf{1}\{v \geq c\} \mid b]$ *and* $J^{\text{PI}}(u^{\text{PI}}(b),b) = \mathbb{E}[(v-c)^+ \mid b]$, *and*
- $u = \mathbb{E}\left[v\mathbf{1}\{v \geq r^{\text{PI}}(u,c,b)\} \mid b\right]$ *where* $r^{\text{PI}}(u,c,b) \triangleq c/(1-x^{\text{PI}}(u,b))$.

The perfect information upper bound suggests a dynamic mechanism that uses the allocation rule $\mathbf{1}\left\{v \geq r^{\text{PI}}(u,c,b)\right\}$ in all periods.

2.4.2 Heuristic Mechanism

In this section, we provide a heuristic mechanism, i.e., a feasible solution for the mechanism design problem (DMDP(b)). Let $u_1^{\text{PI}} \triangleq u^{\text{PI}}(1)$ and $u_2^{\text{PI}} \triangleq u^{\text{PI}}(0)$. Therefore, it trivially follows that $u^{\text{PI}}(b) = bu_1^{\text{PI}} + (1-b)u_2^{\text{PI}}$.

Proposition 2.5. *Define $\bar{\rho} = \max_b u^{\text{PI}}(b)$ and $\underline{\rho} = \min_b u^{\text{PI}}(b)$. Let $\bar{u}(b) = \beta\mathbb{E}[v] - \bar{\rho} + u^{\text{PI}}(b)$ and $\underline{u}(b) = (1-\beta)\bar{v} - \underline{\rho} + u^{\text{PI}}(b)$. Define $(p^{\text{H}}, w^{\text{H}})$ to be*

$$p^{\text{H}}(v,c \mid u,b) = \begin{cases} 1 & \textit{if } \bar{u}(b) < u \leq \bar{w} \\ \mathbf{1}\{v \geq r^{\text{PI}}(u,c,b)\} & \textit{if } \underline{u}(b) \leq u \leq \bar{u}(b) \\ 0 & \textit{if } 0 \leq u < \underline{u}(b) \end{cases} \tag{2.10}$$

and

$$w^{\mathrm{H}}(v,c\,|\,u,b)=\begin{cases}\frac{u}{\beta}-\frac{(1-\beta)\mathbb{E}_{\tilde v}[\tilde v]}{\beta} & \text{if } \bar u(b)<u\le \bar w\\ \frac{u+u^{\mathrm{PI}}(\mathsf{B}(b,c))-u^{\mathrm{PI}}(b)}{\beta}-\frac{(1-\beta)\mathcal{T}\left(r^{\mathrm{PI}}(u,c,b),v\right)}{\beta} & \text{if } \underline{u}(b)\le u\le \bar u(b)\\ \frac{u}{\beta} & \text{if } 0\le u<\underline{u}(b)\end{cases} \tag{2.11}$$

where $\mathcal{T}(x,v)=x\mathbf{1}\{v\ge x\}+\int_x^{\bar v}\bar F(y)dy$. *Then the mechanism* $(p^{\mathrm{H}},w^{\mathrm{H}})$ *satisfies* (FA(b)-1), (PK(u,b)-1), (IC(b)-1), *and* (BC(b)-1).

Note that Proposition 2.3 provides a heuristic mechanism for the known cost distribution setting. The known cost distribution setting in fact is a special case of the unknown cost distribution setting where the initial belief b_1 is either zero or one. Similarly, the mechanism provided in Proposition 2.3 is a special case of the mechanism of Proposition 2.5. Specifically, the term $u^{\mathrm{PI}}(\mathsf{B}(b,c))-u^{\mathrm{PI}}(b)$ inside the future promise function in (2.11) would disappear when b is equal to either zero or one because $\mathsf{B}(1,c)=1$ and $\mathsf{B}(0,c)=0$ for any $c\in\{c_l,c_h\}$.

Define $J^{\mathrm{H}}(u^{\mathrm{PI}}(b_1),b_1)$ to be the expected performance of the heuristic mechanism with initial promised utility $u^{\mathrm{PI}}(b_1)$ and initial belief b_1. That is,

$$J^{\mathrm{H}}(u^{\mathrm{PI}}(b_1),b_1)=\mathbb{E}\left[(1-\beta)\sum_{t=1}^{\infty}\beta^{t-1}(v_t-c_t)p^{\mathrm{H}}(v_t,c_t\,|\,u_t^{\mathrm{H}},b_t)\,\middle|\,u_1^{\mathrm{H}}=u^{\mathrm{PI}}(b_1),b_1\right]$$

$$\text{where } u_{t+1}^{\mathrm{H}}=w^{\mathrm{H}}(v_t,c_t\,|\,u_t^{\mathrm{H}},b_t),\text{and } b_{t+1}=\mathsf{B}(b_t,c_t)\qquad\forall t\ge 1\,.$$

The following result is parallel to Theorem 2.1.

Theorem 2.2. *We have*

$$J^{\mathrm{H}}(u^{\mathrm{PI}}(b_1),b_1)\le J^{\mathrm{PI}}(u^{\mathrm{PI}}(b_1),b_1)\le J^{\mathrm{H}}(u^{\mathrm{PI}}(b_1),b_1)+O(1-\beta)\,.$$

Simlar to Section 2.1, the performance of our heuristic mechanism converges to the perfect information upper bound as the discount factor β approaches one. Note that the mechanism and the analysis is reduced to the analysis of the known cost distribution case if the belief b_1 takes value zero or one. In that sense, the unknown cost distribution is a generalization of the known cost distribution. Moreover, Theorem 2.1 and Theorem 2.2 characterize the convergence rate of the performance of the heuristic mechanisms to the first best. We can show that this convergence rate is optimal, i.e., there is no other mechanism that can converge to the first best at a faster rate, following Theorem 3.2. In the proof of Theorem 3.2, we provide an upper bound for the maximum social welfare J^*_β that converges to $J^{\mathrm{H}}(u^{\mathrm{PI}})$ at a rate of $(1-\beta)$. This result implies that the convergence rate of the heuristic mechanism in the known cost distribution case is optimal. Observing that the known cost distribution setting has the optimal rate of convergence $(1-\beta)$, we conclude that the convergence rate of the heuristic mechanism in the unknown cost distribution case is optimal, too.

Chapter 3

Two Agents

The research in this chapter was conducted under the supervision of Santiago R. Balseiro and Peng Sun and the analysis is largely included in Balseiro et al. (2017).

3.1 Problem Definition and Model

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, i.e., $0 < \underline{f} \leq f_i(v_i) \leq \bar{f} < \infty$ 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.

[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 in order 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.

Notation. For a sequence of vectors $\mathbf{a} = \big(\big(a_{i,t}\big)_{i=1}^{n}\big)_{t=1}^{\infty} \in \mathbb{R}^{n\times\infty}$, we denote by $a_{i,1:t} = \big(a_{i,\ell}\big)_{\ell=1}^{t} \in \mathbb{R}^t$ the i^{th} components of the 1-st to the t-th vectors, and by $\mathbf{a}_{1:t} = \big(\big(a_{i,\ell}\big)_{i=1}^{n}\big)_{\ell=1}^{t} \in \mathbb{R}^{n\times t}$ the entire 1-st to the t-th vectors. For a given vector $\boldsymbol{x}$, we denote by $\boldsymbol{x}_{-i}$ the vector obtained by removing x_i from $\boldsymbol{x}$, and $\boldsymbol{x}^\intercal$ 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.

3.1.1 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.

Non-anticipating Mechanisms and Strategies. Formally, a dynamic mechanism $\boldsymbol{\pi}$ is a sequence of allocation rules $\boldsymbol{\pi} = \big(\big(\pi_{i,t}\big)_{i=1}^{n}\big)_{t=1}^{\infty}$, 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, i.e., $\mathcal{P} \triangleq \Big\{\boldsymbol{\pi} \in [0,1]^n : \boldsymbol{\pi} \geq \mathbf{0}, \sum_{i=1}^{n} \pi_i \leq 1\Big\}$, 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, $\mathbf{h}_t = \big(h_{i,t}\big)_{i=1}^{n}$, as all reports and previous allocations up to time t, where $h_{i,t} = \big(\hat{v}_{i,1:t-1}, \pi_{i,1:t-1}\big)$.[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 non-anticipating strategy profile $\boldsymbol{\sigma} = \big((\sigma_{i,t})_{t=1}^{\infty}\big)_{i=1}^{n}$ 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., $\sigma_{i,t}\big(v_{i,t}, \mathbf{h}_t\big) = \hat{v}_{i,t}$. We say a dynamic mechanism $\boldsymbol{\pi}$ is non-anticipating if $\pi_{i,t}$ depends only on the current period reports $\hat{\mathbf{v}}_t$ and the history $\mathbf{h}_t$, that is, $\boldsymbol{\pi}_t : [0,\bar{v}]^n \times \mathcal{H}_t \to \mathcal{P}$. Since agents' values in each period are

---

[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 chapter, it is clear that asymptotically the planner can do no better through such manipulations.

independent, we restrict attention to non-anticipating 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 one 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,

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

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 $\tilde{\mathbf{h}}_\ell = \Big(\tilde{\mathbf{h}}_{\ell-1},\big(\hat{\mathbf{v}}_{\ell-1},\boldsymbol{\pi}_{\ell-1}(\hat{\mathbf{v}}_{\ell-1},\tilde{\mathbf{h}}_{\ell-1})\big)\Big)$ starting from $\tilde{\mathbf{h}}_t = \mathbf{h}_t$. In order 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:

$$V_{i,t}\Big(\boldsymbol{\pi},\mathbf{I}|v_{i,t},\mathbf{h}_t\Big) \geq V_{i,t}\Big(\boldsymbol{\pi},(\sigma_i,\mathbf{I}_{-i})|v_{i,t},\mathbf{h}_t\Big)\,, \quad \forall v_{i,t},\sigma_i,\mathbf{h}_t,i,t\ , \tag{3.1}$$

where I_i represents the truthful reporting strategy for agent i, i.e., $\mathrm{I}_{i,t}(v_{i,t},\mathbf{h}_t) = v_{i,t}$. Hereafter, we say a non-anticipating mechanism $\boldsymbol{\pi}$ is perfect Bayesian incentive compatible (PBIC) if $\boldsymbol{\pi}$ satisfies (3.1). 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.

$$\mathcal{U}_\beta \triangleq \{\mathbf{u} \in \mathbb{R}^n \mid u_i = V_i(\boldsymbol{\pi}, \mathbf{I}), \text{ for a PBIC mechanism } \boldsymbol{\pi}\} , \tag{3.2}$$

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 non-anticipating direct (dynamic) PBIC mechanism $\boldsymbol{\pi}$ which 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 one 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 $\mathbf{u}_\beta^* \in \operatorname{argmax}_{\mathbf{u} \in \mathcal{U}_\beta} \sum_{i=1}^n u_i$ to represent the utilities of the agents under an optimal mechanism. (Although in this chapter 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 non-anticipating 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.

3.1.2 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 (3.2) in the end of this subsection.

Note that in order 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,

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

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

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

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:

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

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:

$$B_\beta(\mathcal{A}) = \{\mathbf{u} \in \mathbb{R}^n_+ \mid \exists(\mathbf{p}, \mathbf{w}) \text{ satisfying (IC), (FA), (PK}(\mathbf{u})), \text{ and } \mathbf{w}(\mathbf{v}) \in \mathcal{A} \text{ for all } \mathbf{v}\}. \tag{3.3}$$

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 3.1. *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 the promised utilities in such a set $\mathcal{A}$ can be fulfilled with future promises from within the same set. Proposition 3.1 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 3.1. *It is worth noting that the "lower triangle" set $\mathtt{L} \triangleq \left\{\mathbf{u} \in \mathbb{R}^n_+ \;\middle|\; \sum_{i=1}^n u_i/\mathbb{E}[v_i] \leq 1\right\}$ is self-generating. In fact, for any state $\mathbf{u} \in \mathtt{L}$, consider the random allocation rule $p_i^{\mathtt{L}}(\mathbf{v}|\mathbf{u}) \triangleq u_i/\mathbb{E}[v_i]$ regardless of the value $\mathbf{v}$. Such an allocation rule $\mathbf{p}^{\mathtt{L}}$, together with the future promise functions $\mathbf{w}^{\mathtt{L}}(\mathbf{v}|\mathbf{u}) \triangleq \mathbf{u}$, achieves utilities $\mathbf{u}$, i.e., satisfy (PK($\mathbf{u}$)). Therefore, the lower triangle set $\mathtt{L}$ is self-generating with respect to mechanism $(\mathbf{p}^{\mathtt{L}}, \mathbf{w}^{\mathtt{L}})$, and, therefore, is a subset of $\mathcal{U}_\beta$ for any discount factor β.*

Proposition 3.1 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($\mathbf{u}$)) and (IC) in a recursive manner, and with the future promise functions $\mathbf{w}$ lying in the same set $\mathcal{U}_\beta$.

Proposition 3.2. *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 Appendix 6.8 we provide a procedure to numerically compute the set $\mathcal{U}_\beta$ using value iteration in the space of support functions.

3.1.3 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:

$$\mathcal{U} \triangleq \{\mathbf{u} \in \mathbb{R}^n_+ \mid u_i = \mathbb{E}[v_i p_i(\mathbf{v})] \text{ for all } i, \text{ for some } \mathbf{p} \text{ satisfying (FA)}\} . \tag{3.4}$$

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

Following Luenberger (1969, p. 44), any convex set $\mathcal{U}$ 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

$$\phi(\boldsymbol{\alpha}) \triangleq \sup_{\mathbf{u}\in\mathcal{U}} \boldsymbol{\alpha}^\intercal \mathbf{u} = \sup_{\mathbf{p} \text{ s.t. (FA)}} \sum_{i=1}^n \mathbb{E}_{\mathbf{v}}\left[\alpha_i v_i p_i(\mathbf{v})\right] = \mathbb{E}_{\mathbf{v}}\left[\max_{i=1,\dots,n} \alpha_i v_i\right] , \tag{3.5}$$

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 3.3. *The support function $\phi(\boldsymbol{\alpha})$ given in (3.5) 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*

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

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:

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

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 on p. 47).

Furthermore, Proposition 3.3 implies that all points on the efficient frontier of the PI set are achievable by allocations of the form $p_i(\mathbf{v}) = \mathbf{1}\Big\{\alpha_i v_i \geq \max_{j\neq i} \alpha_j v_j\Big\}$. 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:

$$\boldsymbol{\alpha}^*(\mathbf{u}) \in \underset{\boldsymbol{x}:\|\boldsymbol{x}\|_1=1,\boldsymbol{x}\geq\mathbf{0}}{\operatorname{argmax}} \left\{\boldsymbol{x}^\intercal\mathbf{u} - \phi(\boldsymbol{x})\right\}. \tag{3.6}$$

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 $\nabla\phi\big(\boldsymbol{\alpha}^*(\mathbf{u})\big) = \mathbf{u}$.

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

$$J^{\text{FB}} = \max_{\mathbf{u}\in\mathcal{U}} \sum_{i=1}^{n} u_i \geq \max_{\mathbf{u}\in\mathcal{U}_\beta} \sum_{i=1}^{n} u_i = J^*_\beta\,. \tag{3.7}$$

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

3.1.4 An Overview of Our Approach

In general, it is not possible to fully characterize the optimal mechanism in closed-form. Thus, in order 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 chapter 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. Be-

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

cause 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 3.1 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 one as β converges to one, which implies that the mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$ achieves the efficient allocation asymptotically as the discount factor approaches one.

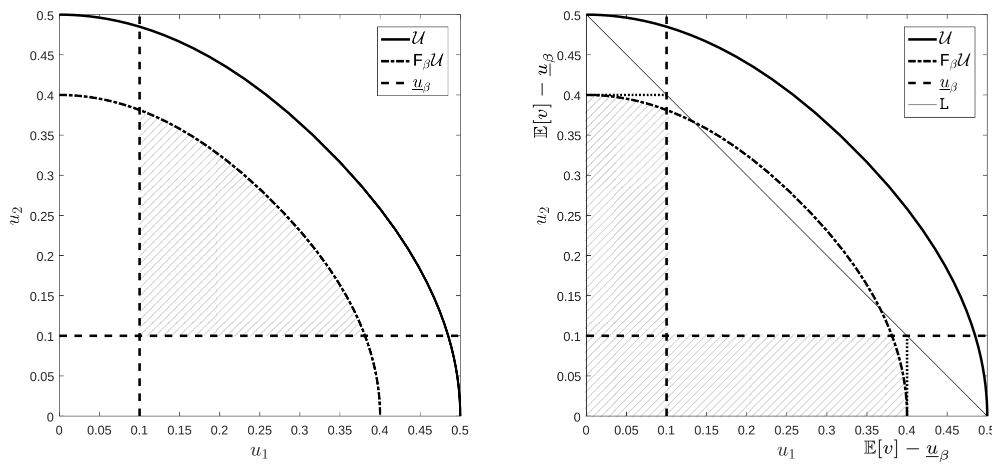

(a) Central Region

(b) Boundary Region

Figure 3.1: The central region and boundary region in a two-agent case.

3.2 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. In order to ensure that the future promises are within the self-generating set, they need to be non-negative and beneath the efficient frontier of the self-generating set. This motivates different designs when the promised utility is close to zero 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 3.1 to deliver the promised utility. This guarantees that, when the promised utility is close to zero, future promises do not to jump too far from the current promised utility and remain non-negative. The crucial design issue, which is the focus on this section, is the main phase mechanism in the central region. (See Figure 3.1 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 zero and F_β approaches one as β approaches one. Define the central region $\hat{\mathcal{U}}_\beta$ as

$$\hat{\mathcal{U}}_\beta = \mathsf{F}_\beta \mathcal{U} \cap \{\mathbf{u} \in \mathbb{R}^n \mid u_i \geq \underline{u}_\beta, \forall i\}, \tag{3.8}$$

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 3.1a illustrates such a central region, where we implement the main phase mechanism, to be described next.

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($\mathbf{u}$)), 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.

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 (3.6). Motivated from this, for any state $\mathbf{u}$ in the central region $\hat{\mathcal{U}}_\beta$ defined in (3.8), we define the allocation function to be

$$\hat{p}_i(\mathbf{v}|\mathbf{u}) = \mathbf{1}\Big\{\alpha_i^*(\mathbf{u}/\mathsf{F}_\beta)v_i \geq \max_{j\neq i} \alpha_j^*(\mathbf{u}/\mathsf{F}_\beta)v_j\Big\}, \tag{3.9}$$

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 3.3, the expected utility of this allocation satisfies $\mathbb{E}_{\mathbf{v}}\left[v_i \hat{p}_i(\mathbf{v}|\mathbf{u})\right] = \dfrac{\partial \phi}{\partial \alpha_i}\big(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)\big)$, i.e., the allocation delivers the utility of the point on the efficient frontier of the set $\mathcal{U}$ with normal $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$.

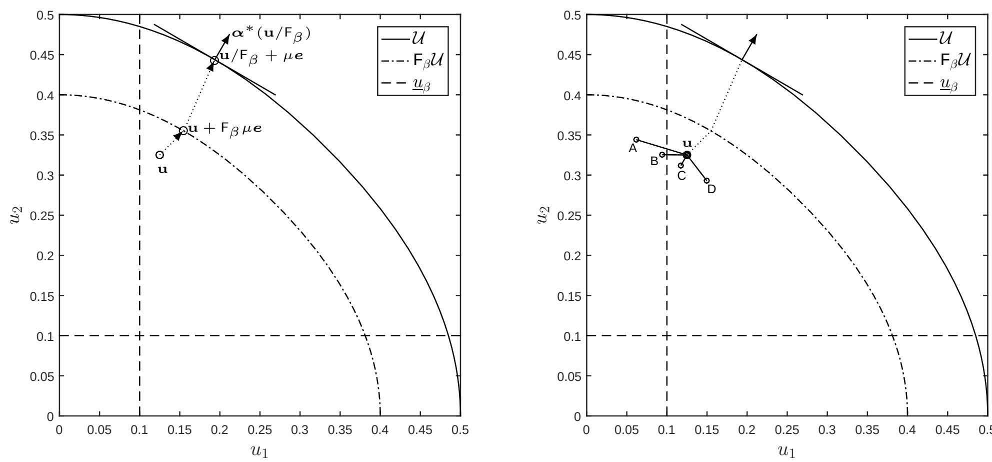

(a) Illustration of the weight $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$ **(b)** Illustration of the future promises $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$

Figure 3.2: The weight $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$ and the future promises $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$

Remark 3.2. *It is worth illustrating geometrically how the weights $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$ are determined. According to the optimality conditions for (3.6), 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 1's in $\mathbb{R}^n$. (Here μ corresponds to the dual variable of the $\|\boldsymbol{x}\|_1 = 1$ 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 3.2a, 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 (3.9) 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 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 in order to induce truthful revelation. Following our design of the future promise functions, as the time discount factor gets closer to one, however, promised utility trajectories tend to concentrate around the 45° line for longer periods of time, resulting in more efficient allocations.

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($\mathbf{u}$)) constraints yields the following interim future promise function:

$$\hat{W}_i(v_i|\mathbf{u}) = \frac{1}{\beta}\left(u_i + (1-\beta)\left(\int_0^{v_i} \hat{P}_i(y|\mathbf{u})\mathrm{d}y - \hat{P}_i(v_i|\mathbf{u})v_i - \int_0^{\bar{v}} \bar{F}_i(y)\hat{P}_i(y|\mathbf{u})\mathrm{d}y\right)\right), \forall i. \tag{3.10}$$

Because the interim allocation is increasing in an agent's report, equation (3.10) is necessary and sufficient to guarantee incentive compatibility.

Although the interim future promise functions $\hat{\mathbf{W}}$ in (3.10) 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, i.e., $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}) \in \mathsf{F}_\beta\mathcal{U}$, 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 3.3. In fact, the threshold $\underline{u}_\beta$ is designed to resolve issue (i).

In order 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 (d'Aspremont and Gérard-Varet, 1979; Arrow, 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:

$$\hat{w}_i(\mathbf{v}|\mathbf{u}) = \hat{W}_i(v_i|\mathbf{u}) - \frac{1}{n-1}\sum_{j\neq i}\frac{\alpha_j^*(\mathbf{u}/\mathsf{F}_\beta)}{\alpha_i^*(\mathbf{u}/\mathsf{F}_\beta)}\left[\hat{W}_j(v_j|\mathbf{u}) - \mathbb{E}_{\tilde{v}_j}[\hat{W}_j(\tilde{v}_j|\mathbf{u})]\right] . \tag{3.11}$$

It is clear from the expression (3.11) that the summation term has an expectation of zero, 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 (3.11).

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 3.4. *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,*

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

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 3.4 are useful towards 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 $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)^\intercal(\mathbf{u}-\mathbf{x})$ as the "signed directional distance" of a point $\mathbf{x} \in \mathbb{R}^n$ to the hyperplane supporting $\mathbf{u}$, we obtain that $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)^\intercal(\mathbf{u}-\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}))$ measures how close the future promises are to the hyperplane supporting $\mathbf{u}$. Among all ex-post future promise functions which satisfy the interim future promise functions $\hat{W}_i(v_i|\mathbf{u})$, the one according to expression (3.11) is "pushed" the farthest from the efficient frontier $\mathcal{E}(\mathsf{F}_\beta\mathcal{U})$. This is formalized in the next result.

Proposition 3.5. *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* (3.10)*. Then, the ex-post future promise* $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ *defined in* (3.11) *is an optimal solution of the following optimization problem:*

$$\begin{aligned} \max_{\tilde{\mathbf{w}}(\cdot)} \quad & \min_{\mathbf{v}} \boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)^\intercal(\mathbf{u} - \tilde{\mathbf{w}}(\mathbf{v})) \\ \text{s.t.} \quad & \mathbb{E}_{\mathbf{v}_{-i}}[\tilde{w}_i(v_i, \mathbf{v}_{-i})] = \hat{W}_i(v_i|\mathbf{u}) \quad \forall v_i, i \end{aligned} \tag{3.12}$$

where $\boldsymbol{\alpha}^*(\cdot)$ *is defined in* (3.6).

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 3.4 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 3.2b 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 $\mathsf{F}_\beta\mathcal{U}$.

3.3 Boundary Region and Mechanism for the Two Agent Case

In the previous section we have described the mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$ for the central region. In order 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 Chapter 4.

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

$$\mathsf{F}_\beta = 1 - \frac{\underline{u}_\beta}{\underline{m}} \quad \text{and} \quad \underline{u}_\beta = \xi(1-\beta)\,, \tag{3.13}$$

where ξ is a constant scalar independent of β, and is given in equation (6.11) in the 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 $\mathtt{L}$. See Figure 3.1b 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,

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

Following Remark 3.1, the mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$ satisfies (IC), (FA), (PK($\mathbf{u}$)) 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 $\mathbf{u} \in \mathsf{F}_\beta\mathcal{U}$, we 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 (3.13), 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 3.6. *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 study relies crucially on the design of the constant ξ in (3.13), 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 (3.11), 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

$$I_{\mathcal{U}}(\mathbf{u}) = \max_{\boldsymbol{x}:\|\boldsymbol{x}\|_1=1,\boldsymbol{x}\geq\mathbf{0}} \left\{\boldsymbol{x}^\intercal\mathbf{u} - \phi(\boldsymbol{x})\right\}. \tag{3.14}$$

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, zero if on the efficient frontier of the set, and negative otherwise. Following (3.6), we obtain that $I_{\mathcal{U}}(\mathbf{u}) = \boldsymbol{\alpha}^*(\mathbf{u})^\intercal\mathbf{u} - \phi(\boldsymbol{\alpha}^*(\mathbf{u})) = \boldsymbol{\alpha}^*(\mathbf{u})^\intercal(\mathbf{u} - \nabla\phi(\boldsymbol{\alpha}^*(\mathbf{u})))$, because $\phi(\boldsymbol{\alpha}) = \nabla\phi(\boldsymbol{\alpha})^\intercal\boldsymbol{\alpha}$. 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 $I_{\mathcal{U}}(\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})/\mathsf{F}_\beta) \leq 0$ (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 $\mathbf{u}/\mathsf{F}_\beta$,

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

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 zeroth-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 3.4. 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 $\mathcal{E}(\mathsf{F}_\beta\mathcal{U})$ starting from any point $\mathbf{u} \in \mathsf{F}_\beta \mathcal{U}$.

Proposition 3.6 implies that starting from any state $\mathbf{u} \in \mathsf{F}_\beta \mathcal{U}$, we can construct a sequence of allocation and future promises which 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

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

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 3.1. 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 3.1. *Let $\underline{\beta} \in (0,1)$ be such that $\mathsf{F}_{\underline{\beta}} \ge 0.5$. For any $\beta \ge \underline{\beta}$, we have*

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

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

It is clear that F_β as defined in (3.13) approaches one as β approaches one. 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 one, 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 3.2. *Suppose agents' values are i.i.d., then*

$$\left(J^{\text{FB}} - J^*_\beta\right) / J^{\text{FB}} \ge \Omega(1-\beta) .$$

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

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 $V(1) = J^{\text{FB}}$ we obtain that $V(\beta) = J^{\text{FB}} - (1-\beta)V'(1_{-}) + o(1-\beta)$ 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).

3.4 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 zero. 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 $\Omega\left((1-\beta)^{1/2}\right)$.

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 value of times that this type is realized over τ_β periods. Using this simple mechanism, an agent following an "approximately truthful strategy" (JS07, pp. 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.

In order 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 $\lim_{\beta\uparrow 1}\sup_{\sigma_i}\inf_{\boldsymbol{\sigma}_{-i}} V_i(\boldsymbol{\pi},(\sigma_i,\boldsymbol{\sigma}_{-i})) = u_i^*$. 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. In order 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 one.

Proposition 3.7. *The budget-based mechanism with cycles of length τ_β secures efficient utility levels if and only if*

$$\lim_{\beta\uparrow 1}\tau_\beta = \infty \quad \textit{and} \quad \lim_{\beta\uparrow 1}\beta^{\tau_\beta} = 1\,. \tag{3.15}$$

Proposition 3.7 implies that the number of time periods in each cycle should be $\tau_\beta = o\big(1/(1-\beta)\big)$ in order 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 (3.15). By design, the mechanism guarantees that the empirical distribution of the competitor's reports over a cycle coincides with the distribution one would expect if the reports are truthful. When β^{τ_β} is close to one, 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 (3.15) 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 one, 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 (3.15), 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 $J^{\text{JS}}_{\beta,\tau_\beta}(\boldsymbol{\sigma})$ be the expected social welfare of the budget-based mechanism with cycle length τ_β when agents employ strategy profile $\boldsymbol{\sigma}$.

Theorem 3.3. *If the length of each cycle τ_β satisfies condition (3.15) and agents have two types, then* $\left(J^{\text{FB}} - \sup_{\boldsymbol{\sigma}} J^{\text{JS}}_{\beta,\tau_\beta}(\boldsymbol{\sigma})\right)/J^{\text{FB}} \geq \Omega\left((1-\beta)^{1/2}\right)$.

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 $\Omega\big(\tau_\beta^{1/2}\big)$ 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 $\Omega\left((1-\beta)^{1/2}\right)$ under condition (3.15).

Proposition 3.7 and Theorem 3.3 together imply that the best possible rate of convergence to first best of a budget-based mechanism that secures efficient utility levels is $\Omega\left((1-\beta)^{1/2}\right)$. This separates our mechanism with the budget-based mechanism extended from JS07. That is, in light of Theorem 3.1, our mechanism converges to first best at a faster rate than the budget-based mechanism as the discount factor β approaches one 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 (3.15) is $O\left((1-\beta)^{1/2}\right)$.

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 Chapter 4.

3.5 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 one).

3.5.1 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 3.8. *Agent i's allocation $\hat{p}_i(\mathbf{v}|\mathbf{u})$ is non-decreasing in the agent's report v_i (for fixed $\mathbf{v}_{-i}$ and $\mathbf{u}$), and non-decreasing in the agent's promised utility u_i (for fixed $\mathbf{v}$ and $\mathbf{u}_{-i}$).*

According to (3.9), 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 3.2, 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 3.8 formalizes these intuitive features of our allocation rule.

Incentive compatibility implies that the interim future promise function of each agent is non-increasing 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 3.9. *The future promise function $\hat{w}_i(\mathbf{v}|\mathbf{u})$ is non-increasing in v_i and non-decreasing 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 3.2b 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.

3.5.2 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 3.3, 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 3.3a, 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 Lemma 6.5 in the appendix). In Figure 3.3a and 3.3b, the first 250 steps of each trajectory are marked black while later steps are colored grey. As we can see from Figure 3.3b, after 250 time periods the promised utilities are still concentrated around the initial state, and it takes longer to reach the boundary region.

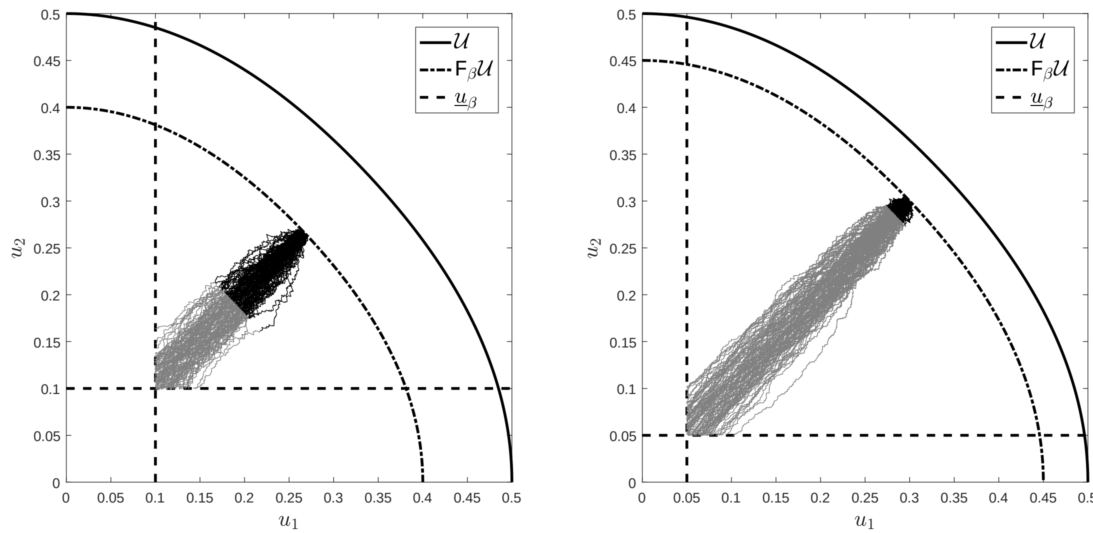

(a) $\mathsf{F}_\beta = 0.8$, $\underline{u}_\beta = 0.1$ and $\beta = 0.9969$

(b) $\mathsf{F}_\beta = 0.9$, $\underline{u}_\beta = 0.05$ $\beta = 0.9985$

Figure 3.3: Each figure demonstrates 100 sample trajectories.

Figure 3.4 illustrates the expected social welfare per round generated by our mechanism, which is given by $\mathbb{E}_{\mathbf{u}_t,\mathbf{v}_t}\left[\sum_{i=1}^{n} v_{i,t}\hat{p}_i(\mathbf{v}_t|\mathbf{u}_t)\right]$, 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 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 3.4a, it takes between 400 and 500 periods for all trajectories to reach the boundary region. In comparison, for a higher discount factor, Figure 3.4b shows that it takes between 1000 and 1500 time periods to reach the boundary region. As the time discount factor approaches one, the boundary mechanism is pushed further into the future and the mechanism allocates approximately efficiently for longer periods of time.

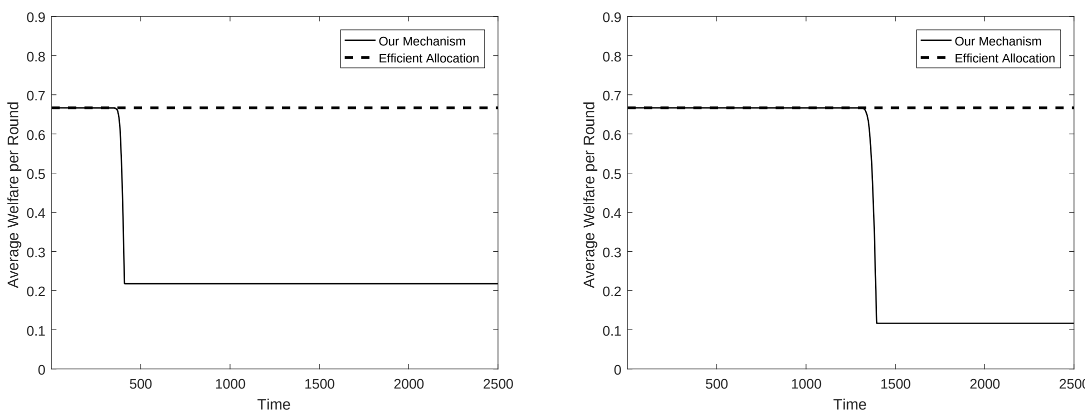

(a) $\mathsf{F}_\beta = 0.8$, $\underline{u}_\beta = 0.1$ and $\beta = 0.9969$

(b) $\mathsf{F}_\beta = 0.9$, $\underline{u}_\beta = 0.05$ and $\beta = 0.9985$

Figure 3.4: The average time to reach the boundary region

Chapter 4

Generalization to n Agents

The research in this chapter was conducted under the supervision of Santiago R. Balseiro and Peng Sun and the analysis is largely included in Balseiro et al. (2017).

4.1 Comparison with the Two-agent Mechanism

In this chapter, we generalize the analysis of the third chapter 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 non-identical 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 $(\hat{\mathbf{p}}^{\mathsf{L}}, \hat{\mathbf{w}}^{\mathsf{L}})$. If $n > 2$, however, the boundary region is not always contained within the lower triangle set $\mathtt{L}$ anymore, as long as the lower bound for the n agent case, $\underline{u}_{\beta}^{(n)}$, approaches zero with β approaching one. 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 4.1. Here we focus on a situation where one agent's promised utility is below the threshold $\underline{u}_{\beta}^{(3)}$. In Figure 4.1a, we plot the efficient frontier of the PI set $\mathcal{E}(\mathcal{U})$, its scaled version $\mathcal{E}(\mathsf{F}_{\beta}\mathcal{U})$, and the plane corresponding to $u_1 = 0.01 < \underline{u}_{\beta}^{(3)} = 0.0167$. Figure 4.1b demonstrates the intersections between the plane with efficient frontiers of sets $\mathcal{U}$, $\mathsf{F}_\beta\mathcal{U}$, and $\mathtt{L}$, 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, $u_1 + u_2 + u_3 > \mathbb{E}[v]$, which implies that the randomized allocation $\hat{\mathbf{p}}^{\mathtt{L}}$ 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.

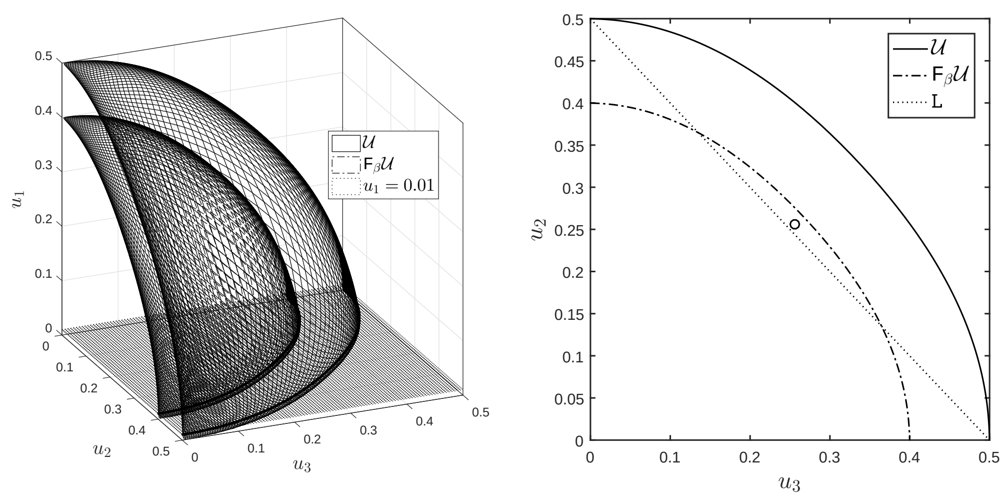

(a) Efficient frontiers

(b) Level set

Figure 4.1: Efficient frontiers and the level set

More generally, suppose the main phase mechanism defined in Section 3.2 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.

4.2 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 6.3 in the appendix.[1]

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 3.2.

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 (3.4), (3.5), and (3.6), respectively. (In these definitions the state vector $\mathbf{u}$ corresponds to a k dimensional vector.) Similar to Section 3.2, 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 3.2 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

[1] Note that the values of the threshold $\underline{u}_\beta$ and the constant ξ for the two-agent mechanism provided in Section 3.3 are different from $\underline{u}_\beta^{(n)}$ and $\xi^{(n)}$ (i.e., $\underline{u}_\beta^{(n)} \neq \underline{u}_\beta$ and $\xi^{(n)} \neq \xi$) because we can provide stronger guarantees in the two-agent case.

active agent i is defined similarly to the main phase mechanism (3.9) for the k agent case. That is,

$$p_i^{(k)}(\mathbf{v}|\mathbf{u}) = s(\mathbf{u})\mathbf{1}\left\{\alpha_i^{(k)}\left(\frac{\mathbf{u}^{(k)}}{s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}}\right)v_i \geq \max_{j\neq i}\alpha_j^{(k)}\left(\frac{\mathbf{u}^{(k)}}{s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}}\right)v_j\right\}. \tag{4.1}$$

The corresponding future utility is defined similar to (3.11), as

$$w_i^{(k)}(\mathbf{v}|\mathbf{u}) = W_i^{(k)}(v_i|\mathbf{u}) - \frac{1}{k-1}\sum_{j\neq i}\frac{\alpha_j^{(k)}\left(\frac{\mathbf{u}^{(k)}}{s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}}\right)}{\alpha_i^{(k)}\left(\frac{\mathbf{u}^{(k)}}{s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}}\right)}\left\{W_j^{(k)}(v_j|\mathbf{u}) - \mathbb{E}_{\tilde{v}_j}\left[W_j^{(k)}(\tilde{v}_j|\mathbf{u})\right]\right\}. \tag{4.2}$$

Here, and the interim future promise function $W_i^{(k)}$ is as defined in (3.10) of Section 3.2, 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

$$\hat{p}_i(\mathbf{v}|\mathbf{u}) = \begin{cases} p_i^{(k)}(\mathbf{v}|\mathbf{u}), & \text{if } u_i \geq \underline{u}_\beta^{(n)}, \\ u_i/\mathbb{E}[v], & \text{otherwise.} \end{cases} \tag{4.3}$$

The future promise function is given by

$$\hat{w}_i(\mathbf{v}|\mathbf{u}) = \begin{cases} w_i^{(k)}(\mathbf{v}|\mathbf{u}), & \text{if } u_i \geq \underline{u}_\beta^{(n)}, \\ u_i, & \text{otherwise.} \end{cases} \tag{4.4}$$

4.3 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 4.1b, for example, the solid curve marks the boundary of the intersection between the 3-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 2-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

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

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 4.1. *The set Ω_β is self-generating with respect to the mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$ defined in (4.3)-(4.4).*

In order 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 (4.3)-(4.4) satisfies the conditions given in (3.3). The (IC), (FA), and (PK($\mathbf{u}$)) 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 4.1, we first show that the subvector of future promises satisfies

$$\mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u}) \in s(\mathbf{u}) \mathsf{F}_\beta^{(k,n)} \mathcal{U}^{(k)} . \tag{4.5}$$

This step extends the geometric approach of Proposition 3.6 to higher dimensions by using the fact that the mechanism for active agents also satisfies the properties in Section 3.2. (In particular, in Appendix 6.9 we prove extensions of Propositions 3.4 and 3.5 that account for the scaling factor.)

Condition (4.5) 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, (4.5) 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 4.1. *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 one as β approaches one, Theorem 4.1 implies that the maximum expected discounted social welfare achieved by the mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$ approaches to first best. In particular the convergence rate is $1-\mathsf{F}_\beta^{(n,n)} = O\left((1-\beta)^{\frac{1}{n+4}}\right)$. Similar to the two agent case, because $\underline{u}_\beta^{(n)}$ converges to zero as β converges to one, set Ω_β converges to the PI achievable set $\mathcal{U}$ as β approaches one. 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 4.1 for $n = 2$ is slower than the one in Theorem 3.1. The proof for the two-agent case given in Section 3.3 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, in order 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 a future research direction.

Chapter 5

Conclusion

In this dissertation, we study resource allocation with asymmetric information and no monetary transfer in a dynamic setting. In the second chapter, there is an allocation cost in each period while the marginal cost for the resource is zero in each period in the third and fourth chapters. In both cases, the mechanism designer focuses on allocation efficiency. We propose mechanisms that achieve asymptotic efficiency as the time discount factor approaches one. From an algorithmic perspective, our mechanism is readily implementable. We provide the explicit sub-optimality of our mechanism, that is, the rate of convergence to first best is $O(1-\beta)$, $O(1-\beta)$ and $O\left((1-\beta)^{\frac{1}{n+4}}\right)$ for the single-agent case and the two-agent case, and the n-agent case, respectively.

In the settings with two agents and n agents, 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 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 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 see if one can still establish asymptotic efficiency in this case.

Chapter 6

Appendix

6.1 Appendix for Chapter 1

6.1.1 Proof of Proposition 2.1

Proof. First, we equivalently represent the optimal value function $J^*_\beta(u)$ in a form that is similar to the basic mechanism design problem provided in (2.2).

$$\begin{aligned} J^*_\beta(u) = & \max_{\pi_t:[0,\bar{v}]\times\mathcal{H}_t\to[0,1]} J^{\boldsymbol{\pi}}_\beta \\ & \text{st. } u = V_t(\boldsymbol{\pi}, \mathbf{I}) \text{ and (2.1)} \end{aligned} \tag{6.1}$$

By its definition, the perfect information value function $J^{\text{PI}}(u)$ is obtained by relaxing constraints (2.1) in the optimization problem above.

$$\begin{aligned} J^{\text{PI}}(u) = & \max_{\pi_t:[0,\bar{v}]\times\mathcal{H}_t\to[0,1]} J^{\boldsymbol{\pi}}_\beta \\ & \text{st. } u = V_t(\boldsymbol{\pi}, \mathbf{I}) \end{aligned} \tag{6.2}$$

Therefore, it readily follows that $J^*_\beta(u) \le J^{\text{PI}}(u)$ for all $u \in [0, \bar{w}]$. Next, we solve this optimization problem. In particular, we show that the perfect information function $J^{\text{PI}}(u)$ can be evaluated by solving the optimization problem $\min_{x\le 1} \mathbb{E}[(v(1-x)-c)^+] + xu$.

We consider a Lagrangian relaxation in which we dualize the constraint $u = V_t(\boldsymbol{\pi}, \mathbf{I})$ in (6.2). Let x be the corresponding Lagrange multiplier. The Lagrangian is given by:

$$\mathcal{L}(\boldsymbol{\pi}, x) = J^{\boldsymbol{\pi}}_\beta - xV_t(\boldsymbol{\pi}, \mathbf{I}) + xu = (1-\beta)\mathbb{E}^{\boldsymbol{\pi}}\left[\sum_{t=1}^{\infty} \beta^{t-1}(v_t(1-x) - c_t)\pi_t(v_t, h_t)\right] + ux\,.$$

Optimizing the Lagrangian function over $\boldsymbol{\pi}$, we obtain the dual objective function:

$$\begin{aligned}\max_{\pi_t:[0,\bar{v}]\times\mathcal{H}_t\to[0,1]} \mathcal{L}(\boldsymbol{\pi}, x) =& ux + \max_{\pi_t:[0,\bar{v}]\times\mathcal{H}_t\to[0,1]} (1-\beta)\mathbb{E}^{\boldsymbol{\pi}}\left[\sum_{t=1}^{\infty}\beta^{t-1}(v_t(1-x)-c_t)\pi_t(v_t,h_t)\right] \\ =& ux + \mathbb{E}[(v(1-x)-c)^+]\end{aligned}$$

Here, the second equality follows from the fact that the distributions of v_t and c_t are i.i.d. across time. Because $u = V_t(\boldsymbol{\pi}, \mathbf{I})$ is a linear constraint and the objective function $J_\beta^{\boldsymbol{\pi}}$ is linear, strong duality implies that $J^{\text{PI}}(u) = \min_x ux + \mathbb{E}[(v(1-x)-c)^+]$. Finally, we can restrict attention to $x \le 1$ because the support of v is nonnegative and $u \ge 0$, and so $J^{\text{PI}}(u) = \min_{x\le 1} ux + \mathbb{E}[(v(1-x)-c)^+]$. □

6.1.2 Proof of Corollary 2.1

Proof. We prove this result in two steps. Before proceeding with the steps, let us define

$$B_u(x) \triangleq \mathbb{E}\left[(v(1-x)-c)^+\right] + xu\,.$$

By Proposition 2.1, we have $J^{\text{PI}}(u) = \min_{x\le 1} B_u(x)$.

Step 1: Characterizing the optimal x. In this step, we characterize the optimal solution of the minimization problem that is used to determine $J^{\text{PI}}(u) = \min_{x\le 1} B_u(x)$. Recognizing the second-order condition, we show that the function $B_u(x)$ is convex in $x \in (-\infty, 1]$. By this observation, we use the first-order condition to characterize $x^{\text{PI}}(u) = \arg\min_{x\le 1} B_u(x)$ for a given value of u.

We start by defining the function $G(r) \triangleq \mathbb{E}[v\mathbf{1}\{v \ge r\}]$.

Lemma 6.1. *$G(r) = \mathbb{E}[v\mathbf{1}\{v \ge r\}] = \int_r^{\bar{v}} \bar{F}(x)dx + r\bar{F}(r)$.*

Lemma 6.1 follows directly from the following derivation, in which the third inequality follows from the Tonelli's Theorem.

$$\mathbb{E}[v\mathbf{1}\{v \ge r\}] = \mathbb{E}[(v-r)^+] + r\bar{F}(r) = \int_r^{\bar{v}}\int_r^{x} \mathrm{d}y f(x)\mathrm{d}x + r\bar{F}(r) = \int_r^{\bar{v}} \bar{F}(x)\mathrm{d}x + r\bar{F}(r)\,.$$

By using $G(\cdot)$, the derivative of $B_u(x)$ with respect to x is given as follows.

$$\frac{\mathrm{d}B_u}{\mathrm{d}x}(x) = -qG\left(\frac{c_h}{1-x}\right) - (1-q)G\left(\frac{c_l}{1-x}\right) + u\,.$$

The first-order condition, i.e., $B_u'(x) = 0$ implies that

$$u = qG\left(\frac{c_h}{1-x}\right) + (1-q)G\left(\frac{c_l}{1-x}\right) = \mathbb{E}\left[v\mathbf{1}\{v \geq c/(1-x)\}\right] .$$

Now, we show that the first-order condition is sufficient for the optimality by showing that $B_u(x)$ is convex in $x \in (-\infty, 1]$. The second derivative of $B_u(x)$ with respect to x is given as follows.

$$\frac{\mathrm{d}^2 B_u}{\mathrm{d}x^2}(x) = qf\left(\frac{c_h}{1-x}\right)\frac{c_h^2}{(1-x)^3} + (1-q)f\left(\frac{c_l}{1-x}\right)\frac{c_l^2}{(1-x)^3} .$$

Note that the p.d.f. is always nonnegative and $x \in (-\infty, 1]$. Therefore, the second derivative of $B_u(x)$ is nonnegative and the convexity follows.

Hence for a given u, we get following results.

- The optimal solution $x^{\text{PI}}(u)$ satisfies $u = \mathbb{E}\left[v\mathbf{1}\{v \geq r^{\text{PI}}(u,c)\}\right]$ where $r^{\text{PI}}(u,\tilde{c}) = \tilde{c}/(1 - x^{\text{PI}}(u))$ for $\tilde{c} \in \{c_l, c_h\}$.
- The optimal solution $x^{\text{PI}}(u)$ is decreasing in u.

Step 2: Characterizing the optimal u. In this step, we prove the main claim by using the results derived in the former step. By the envelope theorem, we get $\mathrm{d}J^{\text{PI}}(u)/\mathrm{d}u = x^{\text{PI}}(u)$. Moreover, we know that $x^{\text{PI}}(u)$ is decreasing in u, therefore, it is sufficient to check the first-order condition in order to find $u^{\text{PI}} = \operatorname{argmax}_{u\in[0,\bar{w}]} J^{\text{PI}}(u)$.

The first-order condition is given by $\mathrm{d}J^{\text{PI}}(u)/\mathrm{d}u|_{u=u^{\text{PI}}} = x^{\text{PI}}(u^{\text{PI}}) = 0$. By the former step, we know that $u = \mathbb{E}\left[v\mathbf{1}\{v \geq r^{\text{PI}}(u,c)\right]$ for all u. Therefore, it follows that

$$u^{\text{PI}} = \mathbb{E}\left[v\mathbf{1}\{v \geq c/(1 - x^{\text{PI}}(u^{\text{PI}}))\}\right] = \mathbb{E}\left[v\mathbf{1}\{v \geq c)\}\right] .$$

We complete the proof by showing that the value of $J^{\text{PI}}(u^{\text{PI}})$ is given by

$$J^{\text{PI}}(u^{\text{PI}}) = B_{u^{\text{PI}}}(x^{\text{PI}}(u^{\text{PI}})) = B_{u^{\text{PI}}}(0) = \mathbb{E}[(v-c)^+] .$$ □

6.1.3 Proof of Proposition 2.2

Proof. We first prove that the allocation and the future promise functions in the statement of the proposition satisfy (2.5) and (PK(u)-1). The classical envelope theorem argument by Myerson

(1981) implies that the ex-post incentive compatibility constraints (2.5) hold if and only if $p(v, c \,|\, u)$ is nondecreasing in v for all c and

$$\beta w(v, c \,|\, u) = (1-\beta)\left[\int_0^v p(y, c \,|\, u)\mathrm{d}y - vp(v, c \,|\, u)\right] + \beta w(0, c \,|\, u)\,.$$

Replacing $\beta w(v, c \,|\, u)$ in (PK(u)-1), we obtain that an allocation function p and the future promise function w satisfy (2.5) and (PK(u)-1) if and only if the allocation function $p(v, c \,|\, u)$ is nondecreasing in v for all c and

$$\beta w(v, c \,|\, u) = u + (1-\beta)\left[\int_0^v p(y, c \,|\, u)\mathrm{d}y - vp(v, c \,|\, u) - \mathbb{E}_{\tilde{v}}\left[\int_0^{\tilde{v}} p(y, c \,|\, u)\mathrm{d}y\right]\right]\,.$$

Note that (IC-1) is obtained by taking expectation of both sides of (2.5) over c. Because p and w satisfy (2.5), they also satisfy (IC-1). □

6.1.4 Proof of Proposition 2.3

Proof. Note that the heuristic mechanism $(p^{\text{H}}, w^{\text{H}})$ satisfies (FA-1), (PK(u)-1), (IC-1) because p^{H} is an allocation function that is nondecreasing in v, and w^{H} is obtained using Proposition 2.2. Therefore, we find the interval $[\underline{u}, \bar{u}]$ which satisfies that the promised utility falls inside of the feasible set $[0, \bar{w}]$, i.e., (BC-1). By the monotonicity of w^{H} in v for $u \in [\underline{u}, \bar{u}]$, it is sufficient for us to check the boundary conditions $w^{\text{H}}(0, \tilde{c} \,|\, u) \leq \bar{w}$ and $w^{\text{H}}(\bar{v}, \tilde{c} \,|\, u) \geq 0$ for all $u \in [\underline{u}, \bar{u}]$ and $\tilde{c} \in \{c_l, c_h\}$.

By definition, for a given value of u and $\tilde{c}$, we get

$$\beta w^{\text{H}}(\bar{v}, \tilde{c} \,|\, u) = u-(1-\beta)\left[r^{\text{PI}}(u, \tilde{c}) + \int_{r^{\text{PI}}(u,\tilde{c})}^{\bar{v}} \bar{F}(y)\mathrm{d}y\right] \text{ and } \beta w^{\text{H}}(0, \tilde{c} \,|\, u) = u-(1-\beta)\int_{r^{\text{PI}}(u,\tilde{c})}^{\bar{v}} \bar{F}(y)\mathrm{d}y\,.$$

First, the conditions $w^{\text{H}}(0, \tilde{c} \,|\, u) \leq \bar{w}$ and $w^{\text{H}}(\bar{v}, \tilde{c} \,|\, u) \geq 0$ can equivalently be expressed as follows.

$$u \geq (1-\beta)\left[r^{\text{PI}}(u, \tilde{c}) + \int_{r^{\text{PI}}(u,\tilde{c})}^{\bar{v}} \bar{F}(y)\mathrm{d}y\right] \text{ and } u \leq \beta\bar{w} + (1-\beta)\int_{r^{\text{PI}}(u,\tilde{c})}^{\bar{v}} \bar{F}(y)\mathrm{d}y\,.$$

Note that if the right-hand side of the first inequality is maximized over $r^{\text{PI}}(u, \tilde{c})$, we get the lower bound $\underline{u}$. The optimal value of this maximization problem is obtained when $r = \bar{v}$ because the function $r+\int_r^{\bar{v}}\bar{F}(y)\mathrm{d}y$ is increasing in r. Therefore the lower bound $\underline{u}$ is given as follows.

$$\underline{u} \geq (1-\beta)\left[\sup_{r\in[0,\bar{v}]} r + \int_r^{\bar{v}} \bar{F}(y)\mathrm{d}y\right] = (1-\beta)\bar{v} \Rightarrow \underline{u} = (1-\beta)\bar{v}\,,$$

Similarly, when the right-hand side of the second inequality is minimized over $r^{\mathrm{PI}}(u,\tilde{c})$, we get

$$\bar{u} \leq (1-\beta)\left[\inf_{r\in[0,\bar{v}]} \frac{\beta\mathbb{E}[v]}{1-\beta} + \int_r^{\bar{v}} \bar{F}(y)\mathrm{d}y\right] \Rightarrow \bar{u} = \beta\mathbb{E}[v]\,.$$

Eventually, we check that $\bar{u} \geq \underline{u}$ for feasibility. This is satisfied when

$$\beta\mathbb{E}[v] \geq (1-\beta)\bar{v} \Rightarrow \beta \geq \frac{\bar{v}}{\bar{v}+\mathbb{E}[v]}\,.$$

This condition holds under Assumption 2.1. □

6.1.5 Proof of Lemma 2.1

Proof. We provide the proof of each item separately.

1. The process $\{u_t^{\mathrm{H}}\}_{t=1}^{\infty}$ satisfies $\mathbb{E}[u_{t+1}^{\mathrm{H}} \,|\, u_t^{\mathrm{H}}] = u_t^{\mathrm{H}}$ when $u_t^{\mathrm{H}} \in [\underline{u}, \bar{u}]$.

   First, note that by promise keeping constraint for any feasible mechanism (w,p) and state u the following condition holds.

   $$\mathbb{E}[w(v,c\,|\,u)] = \frac{u - \mathbb{E}[(1-\beta)p(v,c\,|\,u)v]}{\beta}\,.$$

   Second, by Corollary 2.1 we know that for a given state $u_t^{\mathrm{H}} \in [\underline{u},\bar{u}]$, the allocation of the heuristic mechanism satisfies,

   $$u_t^{\mathrm{H}} = \mathbb{E}[vp^{\mathrm{H}}(v,c\,|\,u_t^{\mathrm{H}})]\,.$$

   Considering these observations together, the result follows.

   $$\mathbb{E}[u_{t+1}^{\mathrm{H}} \,|\, u_t^{\mathrm{H}}] = \mathbb{E}[w^{\mathrm{H}}(v,c\,|\,u_t^{\mathrm{H}}) \,|\, u_t^{\mathrm{H}}] = u_t^{\mathrm{H}}\,.$$

2. Let $K_1 \triangleq (\bar{v}+\bar{w})^2$. Define martingale difference sequence $Y_i^{\mathrm{H}} = u_i^{\mathrm{H}} - u_{i-1}^{\mathrm{H}}$, then we have

$u^{\mathrm{H}}_t - u^{\mathrm{PI}} = \sum_{i=1}^{t} Y^{\mathrm{H}}_i$. Therefore

$$\mathbb{E}[|u^{\mathrm{H}}_t - u^{\mathrm{PI}}|^2 \mathbf{1}\{t \leq \tau\}] \leq \mathbb{E}[|u^{\mathrm{H}}_{t\wedge\tau} - u^{\mathrm{PI}}|^2] = \mathbb{E}\left[\left|\sum_{i=1}^{t} Y^{\mathrm{H}}_i \mathbf{1}\{i \leq \tau\}\right|^2\right] = \sum_{i=1}^{t} \mathbb{E}\left[|Y^{\mathrm{H}}_i|^2\, \mathbf{1}\{i \leq \tau\}\right],$$

where the last equality follows because the stopped martingale $u^{\mathrm{H}}_{t\wedge\tau}$ is a martingale and martingale differences are orthogonal. We prove the result by showing that when $i \leq \tau$ we have that $|Y^{\mathrm{H}}_i| \leq \sqrt{K_1}$ almost surely. In particular, Proposition 2.3 implies that when $i \leq \tau$, $|Y^{\mathrm{H}}_i|$ is equivalently expressed as follows.

$$|Y^{\mathrm{H}}_i| = |u^{\mathrm{H}}_i - u^{\mathrm{H}}_{i-1}| = \left|\left(\frac{1-\beta}{\beta}\right)\left\{u^{\mathrm{H}}_{i-1} - r(u^{\mathrm{H}}_{i-1}, c)\mathbf{1}\{v \geq r^{\mathrm{PI}}(u^{\mathrm{H}}_{i-1}, c)\} - \int_{r^{\mathrm{PI}}(u^{\mathrm{H}}_{i-1}, c)}^{\bar{v}} \bar{F}(y)\mathrm{d}y\right\}\right|.$$

Note any state u^{H}_{i-1} is in $[0, \bar{w}]$ and we know that $\sup_r r + \int_r^{\bar{v}} \bar{F}(y)\mathrm{d}y \leq \bar{v}$ and $\bar{w} \leq \inf_r r + \int_r^{\bar{v}} \bar{F}(y)\mathrm{d}y$. Therefore, it follows that

$$|Y^{\mathrm{H}}_i| \leq \left(\frac{1-\beta}{\beta}\right)\max(0, \bar{v}) = \left(\frac{1-\beta}{\beta}\right)\bar{v}$$

Using the Assumption 2.1, we can replace β in the denominator with $\bar{v}/(\bar{v} + \mathbb{E}[v])$ and obtain the following bound.

$$|Y^{\mathrm{H}}_i| \leq (1-\beta)(\bar{v} + \mathbb{E}[v]) = (1-\beta)(\bar{v} + \bar{w}).$$

Therefore we obtain that,

$$\mathbb{E}[|u^{\mathrm{H}}_t - u^{\mathrm{PI}}|^2 \mathbf{1}\{t \leq \tau\}] \leq (\bar{v} + \bar{w})^2 (1-\beta)^2 \sum_{i=1}^{t} \mathbb{E}\left[|\mathbf{1}\{i \leq \tau\}|^2\right] \leq K_1 t (1-\beta)^2.$$

3. The stopping time being less than t implies that either the maximum or the minimum state under heuristic policy exceeds the bounds $\bar{u}$ or $\underline{u}$, respectively. The state under the heuristic policy exceeds the bounds if and only if one of the following two cases occurs:

   - the distance between the initial state u^{PI} and $\max_{i=1,\ldots,t} u^{\mathrm{H}}_{i\wedge\tau}$ is greater than $\bar{u} - u^{\mathrm{PI}}$.
   - the distance between the initial state u^{PI} and $\min_{i=1,\ldots,t} u^{\mathrm{H}}_{i\wedge\tau}$ is greater than $u^{\mathrm{PI}} - \underline{u}$.

Therefore the bound on the c.d.f. of the stopping time is expressed as follows.

$$\mathbb{P}(\tau < t) \leq \mathbb{P}\left(\max_{i=1,\ldots,t} |u^{\text{PI}} - u^{\text{H}}_{i\wedge\tau}| \geq \min(\bar{u} - u^{\text{PI}}, u^{\text{PI}} - \underline{u})\right) .$$

Recall that by the first step we have $u^{\text{H}}_t \in [\underline{u}, \bar{u}]$ is a martingale; this in turn implies that the stopped process given by $\{u^{\text{H}}_{t\wedge\tau}\}_{t=1}^{\infty}$, and the process obtained by subtracting constant u^{PI} from it (i.e., $\{u^{\text{H}}_{t\wedge\tau} - u^{\text{PI}}\}_{t=1}^{\infty}$) are also martingales.

Recognizing the corollary given by Ross (1996, pp. 315) to bound the c.d.f. of the absolute value of a martingale's highest-order statistic, we get the following expression:

$$\mathbb{P}\left(\max_{i=1,\ldots,t} |u^{\text{PI}} - u^{\text{H}}_{i\wedge\tau}| \geq \min(\bar{u} - u^{\text{PI}}, u^{\text{PI}} - \underline{u})\right) \leq \frac{\mathbb{E}[|u^{\text{PI}} - u^{\text{H}}_{t\wedge\tau}|^2]}{[\min(\bar{u} - u^{\text{PI}}, u^{\text{PI}} - \underline{u})]^2} ,$$

where $\min(\bar{u} - u^{\text{PI}}, u^{\text{PI}} - \underline{u}) = \min\left(\beta\mathbb{E}[v] - \mathbb{E}[v1\{v \geq c\}], \mathbb{E}[v1\{v \geq c\}] - (1-\beta)\bar{v}\right)$.

Assumption 2.1 implies the following two inequalities.

- $\beta\mathbb{E}[v] - \mathbb{E}[v\mathbf{1}\{v \geq c\}] > 0$.
- $\mathbb{E}[v\mathbf{1}\{v \geq c\}] - (1-\beta)\bar{v} > 0$.

It follows that for β close enough to one, there exists $\varepsilon > 0$, independent of β, such that $\min(\bar{u} - u^{\text{PI}}, u^{\text{PI}} - \underline{u}) \geq \varepsilon$. Note that in the former step we show that $\mathbb{E}[|u^{\text{PI}} - u^{\text{H}}_{t\wedge\tau}|^2] \leq K_1 t(1-\beta)^2$. These observations together imply that

$$\mathbb{P}(\tau < t) \leq \frac{\mathbb{E}[|u^{\text{PI}} - u^{\text{H}}_{t\wedge\tau}|^2]}{\varepsilon^2} \leq K_2 t(1-\beta)^2 ,$$

in which $K_2 \triangleq K_1/\varepsilon^2 > 0$ and is independent of β.

4. The expectation in the statement of the lemma is given as follows.

$$\mathbb{E}[\beta^\tau] = \sum_{t=1}^{\infty} \beta^t \mathbb{P}(\tau = t) = \sum_{t=1}^{\infty} \beta^t \mathbb{P}(\tau < t+1) - \sum_{t=1}^{\infty} \beta^t \mathbb{P}(\tau < t) .$$

Note that the first term can equivalently be expressed in a similar form as the second term by change of variable.

$$\sum_{t=1}^{\infty} \beta^t \mathbb{P}(\tau < t+1) = \sum_{t=1}^{\infty} \beta^{t-1} \mathbb{P}(\tau < t) - \mathbb{P}(\tau < 1)$$

Observing that Assumption 2.1 implies $u^{\text{PI}} \in [\underline{u}, \bar{u}]$, therefore the stopping time is a positive integer. Using the result proved in the former item, the expectation is bounded as follows.

$$\mathbb{E}[\beta^\tau] = \frac{1-\beta}{\beta} \sum_{t=1}^{\infty} \beta^t \mathbb{P}(\tau < t) \le K_2(1-\beta)^3 \sum_{t=1}^{\infty} \beta^{t-1} t = K_2(1-\beta)\,.$$

where the last equality is obtained by taking derivatives with respect to β in the geometric series $\sum_{t=0}^{\infty} \beta^t = 1/(1-\beta)$. □

6.1.6 Proof of Lemma 2.2

Proof. We provide the proof of each item separately

1. Let $K_3 = \underline{f}(qc_h^2 + (1-q)c_l^2)/2$ Fix $u \in [\underline{u}, \bar{u}]$. We know that the reserve price corresponding to state u according to heuristic method is given by $r^{\text{PI}}(u, c) = c/(1 - x^{\text{PI}}(u))$. By Corollary 2.1, we also know that $x^{\text{PI}}(u)$ satisfies $u = \mathbb{E}[v\mathbf{1}\{v(1 - x^{\text{PI}}(u)) \ge c\}]$. These two observations together imply that $u = \mathbb{E}[v\mathbf{1}\{v \ge r^{\text{PI}}(u, c)\}]$.

   As a reminder, we denote $G(r) = \mathbb{E}[v\mathbf{1}\{v \ge r\}]$. First assume that $u \le u^{\text{PI}}$.

$$\begin{aligned}
|u - u^{\text{PI}}| = u^{\text{PI}} - u &= q \int_{c_h}^{r^{\text{PI}}(u,c_h)} y f(y) \mathrm{d}y + (1-q) \int_{c_l}^{r^{\text{PI}}(u,c_l)} y f(y) \mathrm{d}y \\
&\ge q \underline{f} \int_{c_h}^{r^{\text{PI}}(u,c_h)} y \mathrm{d}y + (1-q) \underline{f} \int_{c_l}^{r^{\text{PI}}(u,c_l)} y \mathrm{d}y \\
&\ge q c_h \underline{f} \int_{c_h}^{r^{\text{PI}}(u,c_h)} \mathrm{d}y + (1-q) c_l \underline{f} \int_{c_l}^{r^{\text{PI}}(u,c_l)} \mathrm{d}y \\
&= \left(\frac{1}{1 - x^{\text{PI}}(u)} - 1 \right) \underline{f}(qc_h^2 + (1-q)c_l^2)\,.
\end{aligned}$$

   Here, the first inequality is obtained by replacing the p.d.f. with its lower bound $\underline{f}$, and the second inequality is obtained by replacing the integral variable y's with the lower bounds of the integrals.

Next, assume that $u \geq u^{\text{PI}}$.

$$
\begin{aligned}
|u - u^{\text{PI}}| = u - u^{\text{PI}} &= q \int_{r^{\text{PI}}(u,c_h)}^{c_h} y f(y) \mathrm{d}y + (1-q) \int_{r^{\text{PI}}(u,c_l)}^{c_l} y f(y) \mathrm{d}y \\
&\geq q \underline{f} \int_{r^{\text{PI}}(u,c_h)}^{c_h} y \mathrm{d}y + (1-q) \underline{f} \int_{r^{\text{PI}}(u,c_l)}^{c_l} y \mathrm{d}y \\
&= \frac{\underline{f}}{2} (q c_h^2 + (1-q) c_l^2) \left(1 - \frac{1}{1 - x^{\text{PI}}(u)}\right) \left(1 + \frac{1}{1 - x^{\text{PI}}(u)}\right) \\
&\geq \frac{\underline{f}}{2} (q c_h^2 + (1-q) c_l^2) \left(1 - \frac{1}{1 - x^{\text{PI}}(u)}\right) .
\end{aligned}
$$

Here, we obtain the first inequality similarly to the previous step. The last inequality follows from the fact that $(1 + 1/(1 - x^{\text{PI}}(u))) \geq 1$.

2. Let $K_4 \triangleq \bar{f} c_h^2$. Fix $\tilde{c} \in \{c_l, c_h\}$. We first focus on the case where $1 > x \geq 0$. The term $H(x, \tilde{c})$ is bounded as follows.

$$
H(x, \tilde{c}) = (1-\beta) \int_{\tilde{c}}^{\tilde{c}/(1-x)} (y - \tilde{c}) f(y) \mathrm{d}y \leq (1-\beta) \int_{\tilde{c}}^{\tilde{c}/(1-x)} (y - \tilde{c}) \bar{f} \mathrm{d}y \leq (1-\beta) \bar{f} \tilde{c}^2 \left(\frac{1}{1-x} - 1\right)^2 .
$$

Now, consider $x \leq 0$. We have

$$
H(x, \tilde{c}) = (1-\beta) \int_{\tilde{c}/(1-x)}^{\tilde{c}} (\tilde{c} - y) f(y) \mathrm{d}y \leq (1-\beta) \bar{f} \tilde{c}^2 \left(1 - \frac{1}{1-x}\right)^2
$$

which completes the proof. □

6.1.7 Proof of Proposition 2.4

Proof. By its definition, the perfect information value function $J^{\text{PI}}(u, b)$ is given by the following optimization problem.

$$
\begin{aligned}
J^{\text{PI}}(u, b) = \max_{\pi_t : [0, \bar{v}] \times \mathcal{H}_t \to [0,1]} (1-\beta) \mathbb{E}^{\boldsymbol{\pi}} \left[\sum_{t=1}^{\infty} \beta^{t-1} (v_t (1-x) - c_t) \pi_t (v_t, h_t) \,|\, b \right] \\
\text{st. } u = (1-\beta) \mathbb{E}^{\boldsymbol{\pi}} \left[\sum_{t=1}^{\infty} \beta^{t-1} v_t (1-x) \pi_t (v_t, h_t) \,|\, b \right]
\end{aligned}
\tag{6.3}
$$

Therefore, it readily follows that $J^*_\beta(u, b) \leq J^{\text{PI}}(u, b)$ for all $u \in [0, \bar{w}]$ and $b \in [0, 1]$. Next, we solve this optimization problem. In particular, we show that the perfect information function $J^{\text{PI}}(u, b)$

can be evaluated by solving the optimization problem $\min_{x\leq 1} \mathbb{E}[(v(1-x)-c)^+ \,|\, b] + xu$.

We consider a Lagrangian relaxation in which we dualize the promise keeping constraint. Let x be the corresponding Lagrange multiplier. The Lagrangian is given by:

$$\mathcal{L}(\boldsymbol{\pi}, x) = (1-\beta)\mathbb{E}^{\boldsymbol{\pi}}\left[\sum_{t=1}^{\infty}\beta^{t-1}(v_t(1-x)-c_t)\pi_t(v_t,h_t) \,|\, b\right] + ux$$

Optimizing the Lagrangian function over $\boldsymbol{\pi}$, we obtain the dual objective function:

$$\begin{aligned}\max_{\pi_t:[0,\bar{v}]\times\mathcal{H}_t\to[0,1]} \mathcal{L}(\boldsymbol{\pi}, x) =& ux + \max_{\pi_t:[0,\bar{v}]\times\mathcal{H}_t\to[0,1]}(1-\beta)\mathbb{E}^{\boldsymbol{\pi}}\left[\sum_{t=1}^{\infty}\beta^{t-1}(v_t(1-x)-c_t)\pi_t(v_t,h_t) \,|\, b\right]\\ =& ux + \mathbb{E}[(v(1-x)-c)^+ \,|\, b]\end{aligned}$$

Here, the second equality follows from the fact that the distributions of v_t and c_t are i.i.d. across time from the perspective of the first period and the strong duality as discussed in the proof of Proposition 2.1. □

6.1.8 Proof of Corollary 2.2

Proof. This corollary is basically obtained by using Proposition 2.4 and following the same steps in Corollary 2.1 when the probability of c's being c_h is $bq_1 + (1-b)q_2$ and c_l is $1 - bq_1 - (1-b)q_2$. □

6.1.9 Proof of Proposition 2.5

Proof. By construction, $(p^{\text{H}}, w^{\text{H}})$ satisfy the incentive compatibility, the promise keeping and the feasible allocation constraints. Therefore, we find the interval $[\underline{u}(b), \bar{u}(b)]$ which satisfies that the promised utility falls inside of the feasible set (i.e., $[0, \bar{w}]$). By the monotonicity of w^{H} in v, it is sufficient for us to check the boundary conditions $w^{\text{H}}(0, c \,|\, u, b) \leq \bar{w}$ and $w^{\text{H}}(\bar{v}, c \,|\, u, b) \geq 0$ for all $u \in [\underline{u}(b), \bar{u}(b)]$ and $b \in [0, 1]$.

For a given value of u, b and c, we get the following by using the definition of w^{H}.

$$w^{\text{H}}(0, c \,|\, u, b) = \frac{u + u^{\text{PI}}(\mathsf{B}(b,c)) - u^{\text{PI}}(b)}{\beta} - \left(\frac{1-\beta}{\beta}\right)\mathbb{E}\left[(v - r^{\text{PI}}(u,c,b))^+\right]$$

and

$$w^{\mathrm{H}}(\bar{v}, c|u, b) =$$

$$\frac{u + u^{\mathrm{PI}}(\mathsf{B}(b,c)) - u^{\mathrm{PI}}(b)}{\beta} - \frac{1-\beta}{\beta}\left[r^{\mathrm{PI}}(u,c,b)\mathbf{1}\{\bar{v} \geq r^{\mathrm{PI}}(u,c,b)\} + \mathbb{E}\left[(v - r^{\mathrm{PI}}(u,c,b))^{+}\right]\right].$$

We first equivalently represent the conditions $w^{\mathrm{H}}(0, c|u, b) \leq \bar{w}$ and $w^{\mathrm{H}}(\bar{v}, c|u, b) \geq 0$ by using the expressions.

$$u \leq \beta\mathbb{E}[v] + (1-\beta)\,\mathbb{E}\left[(v-r)^{+}\right] - \left(u^{\mathrm{PI}}(\mathsf{B}(b,c)) - u^{\mathrm{PI}}(b)\right)$$

and

$$u \geq (1-\beta)\left[r^{\mathrm{PI}}(u,c,b)\mathbf{1}\{\bar{v} \geq r^{\mathrm{PI}}(u,c,b)\} + \mathbb{E}\left[(v - r^{\mathrm{PI}}(u,c,b))^{+}\right]\right] - \left(u^{\mathrm{PI}}(\mathsf{B}(b,c)) - u^{\mathrm{PI}}(b)\right)$$

Next, we optimize over the bounds to find values of $\underline{u}(b)$ and $\bar{u}(b)$ to that would satisfy the conditions. We minimize the right-hand side of the first inequality over $r^{\mathrm{PI}}(u, c, b)$ and b and c in order to find $\bar{u}(b)$. Therefore, it follows that

$$\bar{u}(b) \leq \beta E[v] - \bar{\rho} + u^{\mathrm{PI}}(b)\,.$$

For the second inequality, we maximize the right-hand side similarly, and complete the proof.

$$\underline{u}(b) \geq (1-\beta)\bar{v} - \underline{\rho} + u^{\mathrm{PI}}(b)\,. \qquad \square$$

6.1.10 Proof of Theorem 2.2

Proof. We prove this theorem by following similar steps as in the proof of Theorem 2.1. We defer the proofs of lemmas to the end of this section.

$$J^{\mathrm{H}}(p^{\mathrm{H}}, w^{\mathrm{H}} \,|\, u^{\mathrm{PI}}(b_1), b_1) = (1-\beta)\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}(v_t - c_t)p^{\mathrm{H}}(v_t, c_t \,|\, u^{\mathrm{H}}_t, b_t) \,|\, u^{\mathrm{H}}_1 = u^{\mathrm{PI}}(b_1), b_1\right]$$

$$\text{where } b_{t+1} = \mathsf{B}(b_t, c_t) \text{ and } u^{\mathrm{H}}_{t+1} = w^{\mathrm{H}}(v_t, c_t \,|\, u^{\mathrm{H}}_t, b_t)\,, \forall t \geq 1\,.$$

Let $\tau = \inf\{t \geq 1 : u^{\mathrm{H}}_t \notin [\underline{u}(b_t), \bar{u}(b_t)]\}$.

Lemma 6.2. *Under the heuristic policy* (2.10)-(2.11)*,*

1. *the process* $\{u^{\mathrm{H}}_t, b_t\}_{t=1}^{\infty}$ *satisfies* $\mathbb{E}[u^{\mathrm{H}}_{t+1} \,|\, u^{\mathrm{H}}_t, b_t] = u^{\mathrm{H}}_t$ *when* $u^{\mathrm{H}}_t \in [\underline{u}(b_t), \bar{u}(b_t)]$.
2. $\mathbb{E}[|u^{\mathrm{H}}_t - u^{\mathrm{PI}}(b_t)|^2\mathbf{1}\{t \leq \tau\}] \leq \tilde{K}_1(1-\beta)^2 t$ *for some constant* $\tilde{K}_1$ *that is independent of* β.

3. *the c.d.f. of the stopping time* τ *satisfies* $\mathbb{P}(\tau < t) \leq \tilde{K}_2 t(1-\beta)^2$ *for some constant* $\tilde{K}_2$ *that is independent of* β.

4. *the probability generating function of the stopping time* τ *satisfies* $\mathbb{E}[\beta^\tau] \leq \tilde{K}_2(1-\beta)$.

Recall that $H(x,c) = \mathbb{E}[(v-c)[\mathbf{1}\{v \geq c\} - \mathbf{1}\{v(1-x) \geq c\}]]$ is the difference in expected performance when the cost realization c is observed between the efficient allocation and the allocation at which the value is weighted by $(1-x)$. The next result provides a deterministic bound.

Lemma 6.3. *For* $u \in [0, \bar{w}]$, *it follows that* $\left|1 - \frac{1}{1-x^{\text{PI}}(u,b)}\right| \leq \tilde{K}_3 |u - u^{\text{PI}}(b)|$ *for* $\tilde{K}_3$ *that is independent of* β.

We are now ready to prove the theorem.

Step 1: The expected performance of the heuristic policy can be decomposed in as follows:

$$J^{\text{H}}(u^{\text{PI}}(b_1), b_1) = (1-\beta)\mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1}(v_t - c_t)p^{\text{H}}(v_t, c_t | u_t^{\text{H}}, b_t) \,|\, u_1^{\text{H}} = u^{\text{PI}}(b_1), b_1\right]$$

$$= \mathbb{E}[(v-c)^+ | b_1] - (1-\beta)\mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1}(v_t - c_t)[\mathbf{1}\{v_t \geq c_t\} - p^{\text{H}}(v_t, c_t | u_t^{\text{H}}, b_t)] \,|\, u_1^{\text{H}} = u^{\text{PI}}(b_1), b_1\right]$$

$$= \mathbb{E}[(v-c)^+ | b_1] - (1-\beta)\sum_{t=1}^{\infty} \mathbb{E}[R_t]$$

We can decompose the error terms R_t as follows:

$$\sum_{t=1}^{\infty} \mathbb{E}[R_t] = \underbrace{\sum_{t=1}^{\infty} \mathbb{E}[R_t \mathbf{1}\{t \leq \tau\}]}_{R_1} + \underbrace{\sum_{t=1}^{\infty} \mathbb{E}[R_t \mathbf{1}\{t > \tau\}]}_{R_2}$$

Next, we find upper bounds for each error terms that are independent of the discount factor β.

Step 2: We first find an upper bound for the error term R_1. Because when $t \leq \tau$ we have that $u_t^{\text{H}} \in [\underline{u}(b_t), \bar{u}(b_t)]$ we can use the bounds of Lemma 6.2 and Lemma 6.3 to obtain the following upper bound.

$$R_1 = \sum_{t=1}^{\infty} \mathbb{E}\left[\beta^{t-1}(v_t - c_t)[\mathbf{1}\{v_t \geq c_t\} - p^{\text{H}}(v_t, c_t \,|\, u_t^{\text{H}}, b_t)]\mathbf{1}\{t \leq \tau\} \,|\, u_1^{\text{H}} = u^{\text{PI}}(b_1), b_1\right]$$

$$= \sum_{t=1}^{\infty} \mathbb{E}\left[\beta^{t-1} H(x^{\text{PI}}(u_t^{\text{H}}, b_t), c_t)\mathbf{1}\{t \leq \tau\} \,|\, u_1^{\text{H}} = u^{\text{PI}}(b_1), b_1\right]$$

$$\leq K_4 \sum_{t=1}^{\infty} \beta^{t-1} \mathbb{E}\left[\left|1 - \frac{1}{1 - x^{\text{PI}}(u_t^{\text{H}}, b_t)}\right|^2 \mathbf{1}\{t \leq \tau\} \,|\, u_1^{\text{H}} = u^{\text{PI}}(b_1), b_1\right]$$

$$\leq K_4 \tilde{K}_3^2 \sum_{t=1}^{\infty} \beta^{t-1} \mathbb{E}\left[|u_t^{\text{H}} - u^{\text{PI}}(b_t)|^2 \mathbf{1}\{t \leq \tau\} \,|\, u_1^{\text{H}} = u^{\text{PI}}(b_1), b_1\right]$$

$$\leq K_4 \tilde{K}_3^2 \tilde{K}_1 (1-\beta)^2 \sum_{t=1}^{\infty} \beta^{t-1} t = K_4 \tilde{K}_3^2 \tilde{K}_1$$

where the first two equalities follow from the definitions of R_1 and $H(\cdot, \cdot)$. The first inequality follows from the second item of Lemma 2.2. The second and the third inequalities are results of Lemma 6.2. The last equality follows from $\sum_{t=1}^{\infty} \beta^{t-1} t = 1/(1-\beta)^2$ because the series converges absolutely when $\beta \in [0, 1)$. Therefore we obtain the following upper bound for the error term:

$$\sum_{t=1}^{\infty} \mathbb{E}[R_t] \leq K_4 \tilde{K}_3^2 \tilde{K}_1 + R_2$$

Step 3: We next bound R_2. We have

$$R_2 \leq \bar{v} \sum_{t=1}^{\infty} \mathbb{E}[\beta^{t-1}\mathbf{1}\{t > \tau\}] = \bar{v}\mathbb{E}\left[\sum_{t=\tau}^{\infty} \beta^t\right] = \frac{\bar{v}}{1-\beta}\mathbb{E}[\beta^\tau] \leq \tilde{K}_2 \bar{v}$$

where the first inequality follows because $|R_t| \leq \beta^{t-1}\bar{v}$ since $c \leq \bar{v}$ and the allocation is bounded by one. The first equality follows from Tonelli's Theorem. The second equality is because the sum is a geometric series, and the last inequality follows from the last item of Lemma 6.2.

Step 4: Putting everything together we obtain that

$$J^{\text{H}}(u^{\text{PI}}(b_1), b_1) \geq \mathbb{E}[(v-c)^+ \,|\, b_1] - (1-\beta)(K_4 \tilde{K}_3^2 \tilde{K}_1 + \tilde{K}_2 \bar{v}) \,. \qquad \square$$

Proof of Lemma 6.2. We provide the proof of each item separately.

1. The process $\{u_t^{\text{H}}, b_t\}_{t=1}^{\infty}$ satisfies $\mathbb{E}[u_{t+1}^{\text{H}} \,|\, u_t^{\text{H}}, b_t] = u_t^{\text{H}}$ for $u_t^{\text{H}} \in [\underline{u}(b_t), \bar{u}(b_t)]$. First, note that by

promise keeping constraint when the belief is b for any feasible mechanism (p, w) and state u the following condition holds.

$$\mathbb{E}[w(v, c \,|\, u, b) \,|\, b] = \frac{u - (1-\beta)\mathbb{E}[vp(v, c \,|\, u, b) \,|\, b]}{\beta} \,.$$

Second, we know that for a given state $u_t^{\mathrm{H}} \in [\underline{u}(b-t), \bar{u}(b_t)]$, the allocation of the heuristic mechanism satisfies,

$$u_t^{\mathrm{H}} = \mathbb{E}[vp^{\mathrm{H}}(v, c \,|\, u_t^{\mathrm{H}}, b_t) \,|\, b_t] \,.$$

Considering these observations together, the result follows:

$$\mathbb{E}[u_{t+1}^{\mathrm{H}} \,|\, u_t^{\mathrm{H}}, b_t] = u_t^{\mathrm{H}} \,.$$

2. Let $\tilde{K}_1 \triangleq (\bar{v} + 2\bar{w})^2$. Define the martingale sequence $Y_i^{\mathrm{H}} = u_i^{\mathrm{H}} - u_{i-1}^{\mathrm{H}} - (u^{\mathrm{PI}}(b_i) - u^{\mathrm{PI}}(b_{i-1}))$, then we have $u_t^{\mathrm{H}} - u^{\mathrm{PI}}(b_t) = \sum_{i=1}^{t} Y_i^{\mathrm{H}}$ because $u_1^{\mathrm{H}} = u^{\mathrm{PI}}(b_1)$. Note also that because $u^{\mathrm{PI}}(\cdot)$ is a linear function and b_t is a martingale. Therefore

$$\begin{aligned}\mathbb{E}[|u_t^{\mathrm{H}} - u^{\mathrm{PI}}(b_t)|^2 \mathbf{1}\{t \le \tau\}] &\le \mathbb{E}[|u_{t\wedge\tau}^{\mathrm{H}} - u^{\mathrm{PI}}(b_{t\wedge\tau})|^2] \\ &= \mathbb{E}\left[\left|\sum_{i=1}^{t} Y_i^{\mathrm{H}} \mathbf{1}\{i \le \tau\}\right|^2\right] \\ &= \sum_{i=1}^{t} \mathbb{E}\left[|Y_i^{\mathrm{H}}|^2 \mathbf{1}\{i \le \tau\}\right] \,.\end{aligned}$$

We prove the result by showing that when $i \le \tau$, $|Y_i^{\mathrm{H}}|$ is equivalently expressed as follows:

$$\begin{aligned}|Y_i^{\mathrm{H}}| &= \left(\frac{1-\beta}{\beta}\right)\left|u_{i-1}^{\mathrm{H}} + u^{\mathrm{PI}}(b_i) - u^{\mathrm{PI}}(b_{i-1}) - \mathcal{T}(r^{\mathrm{PI}}(u_{i-1}^{\mathrm{H}}, c_{i-1}, b_{i-1}), v_{i-1})\right| \\ &\le (1-\beta)\tilde{K}_1^{1/2}\end{aligned}$$

where $\mathcal{T}(x, v) = x\mathbf{1}\{v \ge x\} + \int_x^{\bar{v}} \bar{F}(y)\mathrm{d}y$. By definition $\mathcal{T}(x, v)$ is bounded by $\bar{v} + \bar{w}$ for any x, $v \in [0, \bar{v}]$. Note also that u_{i-1}^{H}, $u^{\mathrm{PI}}(b_i)$ and $u^{\mathrm{PI}}(b_{i-1})$ are smaller than $\bar{w}$. Therefore, the last bound follows, and we obtain

$$\mathbb{E}[|u_t^{\mathrm{H}} - u^{\mathrm{PI}}(b_t)|^2 \mathbf{1}\{t \le \tau\}] \le \tilde{K}_1(1-\beta)^2 \sum_{i=1}^{t} \mathbb{E}\left[|\mathbf{1}\{i \le \tau\}|^2\right] \le \tilde{K}_1(1-\beta)^2 t$$

because $\mathbf{1}\{i \leq \tau\} \leq 1$.

3. The stopping time being less than t implies that the state under heuristic policy exceeds the bounds $\bar{u}(b_t)$ or $\underline{u}(b_t)$. The states under heuristic policy exceeds the bounds if one of the following two cases occurs:

   - $\max_{i=1,\ldots,t} |u^{\text{H}}_{i\wedge\tau} - u^{\text{PI}}(b_{i\wedge\tau})|$ is greater than $\beta\mathbb{E}[v] - \bar{\rho}$.
   - $\max_{i=1,\ldots,t} |u^{\text{H}}_{i\wedge\tau} - u^{\text{PI}}(b_{i\wedge\tau})|$ is greater than $\underline{\rho} - (1-\beta)\bar{v}$.

   Therefore, the bound on the c.d.f. of the stopping time is expressed as follows:

   $$\mathbb{P}(\tau < t) \leq \mathbb{P}\left(\max_{i=1,\ldots,t} |u^{\text{H}}_{i\wedge\tau} - u^{\text{PI}}(b_{i\wedge\tau})| \geq \min(\beta\mathbb{E}[v] - \bar{\rho}, \underline{\rho} - (1-\beta)\bar{v})\right)$$

   Recall that by the first step we have $u^{\text{H}}_t \in [\underline{u}(b_t), \bar{u}(b_t)]$ is a martingale; this in turn implies that the stopped process given by $\{u^{\text{H}}_{t\wedge\tau}\}_{t=1}^{\infty}$, and the process obtained by subtracting another martingale $\{u^{\text{PI}}(b_{t\wedge\tau})\}_{t=1}^{\infty}$ is also martingale. Recognizing the corollary given by Ross (1996, pp. 315) to bound the c.d.f. of the absolute value of a martingale's highest-order statistic, we get the following expression:

   $$\mathbb{P}\left(\max_{i=1,\ldots,t} |u^{\text{H}}_{i\wedge\tau} - u^{\text{PI}}(b_{i\wedge\tau})| \geq \min(\beta\mathbb{E}[v] - \bar{\rho}, \underline{\rho} - (1-\beta)\bar{v})\right) \leq \frac{\mathbb{E}[|u^{\text{H}}_{t\wedge\tau} - u^{\text{PI}}(b_{t\wedge\tau})|^2]}{[\min(\beta\mathbb{E}[v] - \bar{\rho}, \underline{\rho} - (1-\beta)\bar{v})]^2}.$$

   Note that in the former step we show that $\mathbb{E}[|u^{\text{H}}_{t\wedge\tau} - u^{\text{PI}}(b_{t\wedge\tau})|^2] \leq \tilde{K}_1(1-\beta)^2$. Because $[\min(\beta\mathbb{E}[v] - \bar{\rho}, \underline{\rho} - (1-\beta)\bar{v})] > \varepsilon > 0$. These observations together imply that

   $$\mathbb{P}(\tau < t) \leq \tilde{K}_2 t(1-\beta)^2$$

   where $\tilde{K}_2 \triangleq \tilde{K}_1/\varepsilon$ and the claim follows.

4. The expectation in the statement of the lemma is given as follows:

   $$\mathbb{E}[\beta^\tau] = \sum_{t=1}^{\infty} \beta^t \mathbb{P}(\tau = t) = \sum_{t=1}^{\infty} \beta^t \mathbb{P}(\tau < t+1) - \sum_{t=1}^{\infty} \beta^t \mathbb{P}(\tau < t)$$

   Note that the first term can equivalently be expressed in a similar form as the second term by change of variables.

$$\sum_{t=1}^{\infty} \beta^t \mathbb{P}(\tau < t+1) = \sum_{t=1}^{\infty} \beta^{t-1} \mathbb{P}(\tau < t) - \mathbb{P}(\tau < 1)$$

Observing the stopping time is a positive integer and the result proved in the former item, the expectation is bounded as follows:

$$\mathbb{E}[\beta^\tau] = \left(\frac{1-\beta}{\beta}\right) \sum_{t=1}^{\infty} \beta^t \mathbb{P}(\tau < t) \leq \tilde{K}_2 (1-\beta)^3 \sum_{t=1}^{\infty} \beta^{t-1} t\,.$$

By taking derivatives with respect to β in the geometric series $\sum_{t=0}^{\infty} \beta^t = 1/(1-\beta)$ because the series converges absolutely when $\beta \in [0,1)$, the result follows.

$$\mathbb{E}[\beta^\tau] \leq \tilde{K}_2 (1-\beta)\,. \qquad \square$$

Proof of Lemma 6.3. Fix u and b such that $u \in [\underline{u}(b), \bar{u}(b)]$. We know that the weight corresponding to state u according to the heuristic mechanism satisfies $u = \mathbb{E}[v\mathbf{1}\{v(1 - x^{\text{PI}}(u,b)) \geq c\} \,|\, b]$ and $u^{\text{PI}}(b) = \mathbb{E}[v\mathbf{1}\{v \geq c\} \,|\, b]$. Therefore, the rest of the proof follows from the first item of Lemma 2.2 when the probability of c's being c_h is $bq_1 + (1-b)q_2$ and c_l is $1 - bq_1 - (1-b)q_2$. Therefore, the constant $\tilde{K}_3 \triangleq \underline{f}(bq_1 + (1-b)q_2 c_h^2 + (1 - bq_1 + (1-b)q_2)c_l^2)/2$ suffices the result. $\square$

6.2 Appendix for Chapter 3 Section 3.1

6.2.1 Proof of Proposition 3.1

Proof. We need to show that if a set $\mathcal{A}$ satisfies $\mathcal{A} \subseteq B_\beta(\mathcal{A})$, then $B_\beta(\mathcal{A}) \subseteq \mathcal{U}_\beta$. We prove this by constructing, for each point $\mathbf{u} \in B_\beta(\mathcal{A})$, a mechanism $\boldsymbol{\pi}$ satisfying (3.1) such that $u_i = V_i(\boldsymbol{\pi}, \mathbf{I})$.

Step 1 (constructing mechanism $\boldsymbol{\pi}$). Let $(\hat{\mathbf{v}}_t)_{t\geq 1}$ be a sequence of reports. We construct a mechanism recursively that depends exclusively on past reports (and not on the allocations) as follows. First, let $\mathbf{u}_1 = \mathbf{u}$. Because $\mathbf{u}_1 \in B_\beta(\mathcal{A})$, there exists a pair of functions $(\mathbf{p}_1^{\mathbf{u}_1}, \mathbf{w}_1^{\mathbf{u}_1})$ satisfying (IC), (PK($\mathbf{u}_1$)), and (FA) and $\mathbf{w}_1^{\mathbf{u}_1}(\mathbf{v}) \in \mathcal{A}$ for all $\mathbf{v}$. We proceed for $t > 1$ by setting $\mathbf{u}_t = \mathbf{w}_{t-1}^{\mathbf{u}_{t-1}}(\hat{\mathbf{v}}_{t-1})$. Because $\mathbf{u}_t \in \mathcal{A} \subseteq B_\beta(\mathcal{A})$, there exists a pair of functions $(\mathbf{p}_t^{\mathbf{u}_t}, \mathbf{w}_t^{\mathbf{u}_t})$ satisfying (IC), (PK($\mathbf{u}_t$)), and (FA) and $\mathbf{w}_t^{\mathbf{u}_t}(\mathbf{v}) \in \mathcal{A}$ for all $\mathbf{v}$. The mechanism is thus given by $\boldsymbol{\pi}_t(\mathbf{v}, \mathbf{h}_t) = \mathbf{p}_t^{\mathbf{u}_t}(\mathbf{v})$ with the history $\mathbf{h}_t = \hat{\mathbf{v}}_{1:t-1}$ given by the past reports (which recursively determines the promise utility $\mathbf{u}_t$).

Step 2 (mechanism $\boldsymbol{\pi}$ satisfies $u_i = V_i(\boldsymbol{\pi}, \mathbf{I})$). This result follows from the fact that $(\mathbf{p}_t^{\mathbf{u}_t}, \mathbf{w}_t^{\mathbf{u}_t})$ satisfy (PK($\mathbf{u}_t$)) for all $t \geq 1$. In particular, we replace the continuation values with each $w_{i,t}^{\mathbf{u}_t}(\mathbf{v})$ by using the promise keeping constraints. More formally, when agents bid truthfully and $\hat{\mathbf{v}}_t = \mathbf{v}_t$, we obtain by iterating up to a time $\ell < \infty$

$$
\begin{aligned}
u_i &= (1-\beta)\mathbb{E}[v_{i,1} p_{i,1}^{\mathbf{u}_1}(\mathbf{v}_1)] + \beta \mathbb{E}[w_{i,1}^{\mathbf{u}_1}(\mathbf{v}_1)] = (1-\beta)\mathbb{E}[v_{i,1} p_{i,1}^{\mathbf{u}_1}(\mathbf{v}_1)] + \beta \mathbb{E}[u_{2,1}] \\
&= (1-\beta) \sum_{t=1}^{\ell} \beta^{t-1} \mathbb{E}[v_{i,t} p_{i,t}^{\mathbf{u}_t}(\mathbf{v}_t)] + \beta^\ell \mathbb{E}[u_{i,\ell+1}] \\
&= (1-\beta) \sum_{t=1}^{\ell} \beta^{t-1} \mathbb{E}^{\boldsymbol{\pi},\mathbf{I}}[v_{i,t} \pi_{i,t}(\mathbf{v}_t, \mathbf{h}_t)] + \beta^\ell \mathbb{E}[u_{i,\ell+1}] \,,
\end{aligned}
$$

where the first, second and third equalities follow from the fact that $(\mathbf{p}_t^{\mathbf{u}_t}, \mathbf{w}_t^{\mathbf{u}_t})$ satisfy (PK($\mathbf{u}_t$)) for all $t \geq 1$ and using that $\mathbf{u}_t = \mathbf{w}_{t-1}^{\mathbf{u}_{t-1}}(\mathbf{v}_{t-1})$ together with the fact that values are identically distributed; and the fourth equality follows from our definition of the mechanism $\boldsymbol{\pi}$. Because values have finite means and $\beta \in (0,1)$, we obtain that the series in absolutely convergent. Taking ℓ to infinity, we conclude that $u_i = V_i(\boldsymbol{\pi}, \mathbf{I})$.

Step 3 (mechanism $\boldsymbol{\pi}$ is PBIC). Fix a time period $t \geq 1$. We need to show that

$$V_{i,t}(\boldsymbol{\pi}, \mathbf{I}|v_{i,t}, \mathbf{h}_t) \geq V_{i,t}(\boldsymbol{\pi}, (\sigma_i, \mathbf{I}_{-i})|v_{i,t}, \mathbf{h}_t)\,, \tag{6.4}$$

for every history $\mathbf{h}_t$ and strategy σ_i. By the one-shot deviation principle, it suffices to show that agent i has no incentive to deviate at time t. Therefore, we restrict attention to strategies σ_i that reports $\sigma_{i,t}(v_{i,t}, \mathbf{h}_t) = \hat{v}_{i,t}$ at time t and truthfully at time $\ell > t$. Additionally, because the mechanism $\boldsymbol{\pi}$ only depends on the past reports via the promise utility $\mathbf{u}_t$, it suffices to consider $\mathbf{u}_t$ as the state.

Because under the strategies $(\sigma_i, \mathbf{I}_{-i})$ the other agents report truthfully at time t and all agents report truthfully at times $\ell > t$, we obtain using a similar argument to that of step 2 that

$$V_{i,t}(\boldsymbol{\pi}, (\sigma_i, \mathbf{I}_{-i})|v_{i,t}, \mathbf{h}_t) = (1-\beta)v_{i,t}\mathbb{E}_{\mathbf{v}_{-i,t}}[p^{\mathbf{u}_t}_{i,t}(\hat{v}_{i,t}, \mathbf{v}_{-i,t})] + \beta\mathbb{E}_{\mathbf{v}_{-i,t}}[w^{\mathbf{u}_t}_{i,t}(\hat{v}_{i,t}, \mathbf{v}_{-i,t})]\,,$$

where $\hat{v}_{i,\ell} = \sigma_{i,\ell}(v_{i,\ell}, \mathbf{h}_\ell)$ is the report of agent i at time t. Therefore, we can write (6.4) as

$$\begin{aligned} &(1-\beta)v_{i,t}\mathbb{E}_{\mathbf{v}_{-i,t}}[p^{\mathbf{u}_t}_{i,t}(\mathbf{v}_t)] + \beta\mathbb{E}_{\mathbf{v}_{-i,t}}[w^{\mathbf{u}_t}_{i,t}(\mathbf{v}_t)] \\ &\geq (1-\beta)v_{i,t}\mathbb{E}_{\mathbf{v}_{-i,t}}[p^{\mathbf{u}_t}_{i,t}(\hat{v}_{i,t}, \mathbf{v}_{-i,t})] + \beta\mathbb{E}_{\mathbf{v}_{-i,t}}[w^{\mathbf{u}_t}_{i,t}(\hat{v}_{i,t}, \mathbf{v}_{-i,t})]\,, \end{aligned}$$

which follows because the mechanism $(\mathbf{p}^{\mathbf{u}_t}_t, \mathbf{w}^{\mathbf{u}_t}_t)$ satisfies (IC) and agents' values are identically distributed. □

6.2.2 Proof of Proposition 3.2

Proof. Recall that the set of achievable utilities is given by

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

while the operator $B_\beta(\cdot)$ is defined in (3.3). We prove that $\mathcal{U}_\beta = B_\beta(\mathcal{U}_\beta)$ by showing that $\mathcal{U}_\beta \subseteq B_\beta(\mathcal{U}_\beta)$, since by Proposition 3.1 the latter implies that $B_\beta(\mathcal{U}_\beta) \subseteq \mathcal{U}_\beta$. For a given point $\mathbf{u} \in \mathcal{U}_\beta$, we show that $\mathbf{u} \in B_\beta(\mathcal{U}_\beta)$ by constructing an allocation function $\mathbf{p}(\cdot)$ and a future promise function $\mathbf{w}(\cdot)$, such that $(\mathbf{p}, \mathbf{w})$ satisfies (IC), (FA),(PK($\mathbf{u}$)) and $\mathbf{w}(\mathbf{v}) \in \mathcal{U}_\beta$ for all $\mathbf{v}$.

Step 1 (constructing stage mechanism $(\mathbf{p}, \mathbf{w})$). Because $\mathbf{u} \in \mathcal{U}_\beta$, there exists a PBIC mechanism $\boldsymbol{\pi}$ satisfying (3.1) and $\boldsymbol{\pi}_t \in \mathcal{P}$ for all t, and $u_i = V_i(\boldsymbol{\pi}, \mathbf{I})$. We use the first period allocation and the continuation values induced by mechanism $\boldsymbol{\pi}$ to construct the stage mechanism $(\mathbf{p}, \mathbf{w})$. For all $\mathbf{v}$, we set $\mathbf{p}(\mathbf{v}) = \boldsymbol{\pi}_1(\mathbf{v}, \emptyset)$ and

$$w_i(\mathbf{v}) = (1-\beta)\sum_{\ell=2}^{\infty} \beta^{\ell-2}\mathbb{E}^{\boldsymbol{\pi},\mathbf{I}}[v_{i,\ell}\pi_{i,\ell}(\mathbf{v}_\ell, \mathbf{h}_\ell)|\mathbf{h}_2 = (\mathbf{v}, \boldsymbol{\pi}_1(\mathbf{v}, \emptyset))]\,. \tag{6.5}$$

Note that the function $\mathbf{p}$ trivially satisfies (FA) since $\mathbf{p}(\mathbf{v}) = \boldsymbol{\pi}_1(\mathbf{v}, \emptyset)$ and $\boldsymbol{\pi}_1(\mathbf{v}, \emptyset) \in \mathcal{P}$ for all $\mathbf{v}$. The promise keeping constraint holds because

$$u_i = V_i(\boldsymbol{\pi}, \mathbf{I}) = (1-\beta)\mathbb{E}_{\mathbf{v}}\left[v_i p_i(\mathbf{v})\right] + \beta\mathbb{E}_{\mathbf{v}}\left[w_i(\mathbf{v})\right]\,,$$

where the last equation follows from our definition of the allocation function $\mathbf{p}$ and the future promise function $\mathbf{w}$. Thus the stage mechanism $(\mathbf{p}, \mathbf{w})$ satisfies (PK($\mathbf{u}$)).

Step 2 (stage mechanism $(\mathbf{p}, \mathbf{w})$ satisfies (IC)). Since the mechanism $\boldsymbol{\pi}$ is PBIC, we obtain that at time $t = 1$ agent i has no incentive to misreport his value:

$$V_{i,1}(\boldsymbol{\pi}, (\sigma_i, \mathbf{I}_{-i})|v_{i,1}, \emptyset) \leq V_{i,1}(\boldsymbol{\pi}, \mathbf{I}|v_{i,1}, \emptyset), \quad \forall v_{i,1}\,.$$

Consider the strategy σ_i that reports $\sigma_{i,1}(v_{i,1}, \emptyset) = \hat{v}_{i,1}$ at time t and truthfully at time $t > 1$. We obtain using our definition of the stage mechanism $(\mathbf{p}, \mathbf{w})$ that

$$V_{i,1}(\boldsymbol{\pi}, (\sigma_i, \mathbf{I}_{-i})|v_{i,1}, \emptyset) = (1-\beta)v_{i,1}\mathbb{E}_{\mathbf{v}_{-i,1}}[p_i(\hat{v}_{i,1}, \mathbf{v}_{-i,1})] + \beta\mathbb{E}_{\mathbf{v}_{-i,1}}[w_i(\hat{v}_{i,1}, \mathbf{v}_{-i,1})]\,,$$

and

$$V_{i,1}(\boldsymbol{\pi}, \mathbf{I}|v_{i,1}, \emptyset) = (1-\beta)v_{i,1}\mathbb{E}_{\mathbf{v}_{-i,1}}[p_i(v_{i,1}, \mathbf{v}_{-i,1})] + \beta\mathbb{E}[w_i(v_{i,1}, \mathbf{v}_{-i,1})]\,.$$

This implies that the stage mechanism $(\mathbf{p}, \mathbf{w})$ satisfies (IC).

Step 3 (for all $\mathbf{v}$, $\mathbf{w}(\mathbf{v}) \in \mathcal{U}_\beta$). Fix a report $\hat{\mathbf{v}}_1 \in [0, \bar{v}]^n$ for period 1. We prove that $\mathbf{w}(\hat{\mathbf{v}}_1) \in \mathcal{U}_\beta$, by constructing a PBIC mechanism $\bar{\boldsymbol{\pi}}$ such that $w_i(\hat{\mathbf{v}}_1) = V_i(\bar{\boldsymbol{\pi}}, \mathbf{I})$. The proposed mechanism $\bar{\boldsymbol{\pi}}$ involves implementing at time t the allocation of mechanism $\boldsymbol{\pi}$ at $t+1$ given that the reports of the first period $\hat{\mathbf{v}}_1$, that is, we shift all allocations one time period. For any history $\bar{\mathbf{h}}_t \in \mathcal{H}_t$ for the shifted mechanism $\bar{\boldsymbol{\pi}}$ consider the history $\mathbf{h}_{t+1} \in \mathcal{H}_{t+1}$ for the original mechanism $\boldsymbol{\pi}$ given by $\mathbf{h}_{t+1} = \Big((\hat{\mathbf{v}}_1, \boldsymbol{\pi}_1(\hat{\mathbf{v}}_1, \emptyset)), \bar{\mathbf{h}}_t\Big)$ and $\mathbf{h}_2 = (\hat{\mathbf{v}}_1, \boldsymbol{\pi}_1(\hat{\mathbf{v}}_1, \emptyset))$. Using this notation, we can denote the shifted mechanism as $\bar{\boldsymbol{\pi}}_t(\mathbf{v}, \bar{\mathbf{h}}_t) = \boldsymbol{\pi}_{t+1}(\mathbf{v}, \mathbf{h}_{t+1})$ for all $t \geq 1$.

First, we show that $w_i(\hat{\mathbf{v}}_1) = V_i(\bar{\boldsymbol{\pi}}, \mathbf{I})$. From (6.5), we have that

$$
\begin{aligned}
w_i(\hat{\mathbf{v}}_1) =& (1-\beta)\sum_{\ell=2}^{\infty} \beta^{\ell-2}\mathbb{E}^{\boldsymbol{\pi},\mathbf{I}}[v_{i,\ell}\pi_{i,\ell}(\mathbf{v}_\ell, \mathbf{h}_\ell)|\mathbf{h}_2 = (\hat{\mathbf{v}}_1, \boldsymbol{\pi}_1(\hat{\mathbf{v}}_1, \emptyset))] \\
=& (1-\beta)\sum_{\ell=1}^{\infty} \beta^{\ell-1}\mathbb{E}^{\boldsymbol{\pi},\mathbf{I}}[v_{i,\ell}\pi_{i,\ell+1}(\mathbf{v}_\ell, \mathbf{h}_{\ell+1})|\mathbf{h}_2 = (\hat{\mathbf{v}}_1, \boldsymbol{\pi}_1(\hat{\mathbf{v}}_1, \emptyset))] \\
=& (1-\beta)\sum_{\ell=1}^{\infty} \beta^{\ell-1}\mathbb{E}^{\bar{\boldsymbol{\pi}},\mathbf{I}}[v_{i,\ell}\bar{\pi}_{i,\ell}(\mathbf{v}_\ell, \bar{\mathbf{h}}_\ell)|\emptyset] = V_i(\bar{\boldsymbol{\pi}}, \mathbf{I})\,,
\end{aligned}
$$

where the second equation follows from shifting the index in the summation and using that values are identically distributed, and the third equation follows from our definition of the shifted mechanism and its history.

Second, we show that the shifted mechanism $\bar{\boldsymbol{\pi}}$ is PBIC. For any strategy $\bar{\sigma}_i$ for the shifted mechanism $\bar{\boldsymbol{\pi}}$ consider the strategy σ_i for the original mechanism $\boldsymbol{\pi}$ given by $\sigma_{i,t+1}(\cdot, \mathbf{h}_{t+1}) = \bar{\sigma}_{i,t}(\cdot, \bar{\mathbf{h}}_t)$ that shifts reports by one time period (since we consider the original mechanism at times $t > 1$, the reports of the strategy $\sigma_{i,t}$ at time $t = 1$ are irrelevant). Because agents only condition their actions on previous reports and allocations, and this information is publicly observed, we claim that for all v_i, $\bar{\sigma}_i$, $\bar{\mathbf{h}}_t$, i and $t \geq 1$

$$
V_{i,t}\left(\bar{\boldsymbol{\pi}}, (\bar{\sigma}_i, \mathbf{I}_{-i})|v_i, \bar{\mathbf{h}}_t\right) = V_{i,t+1}\left(\boldsymbol{\pi}, (\sigma_i, \mathbf{I}_{-i})|v_i, \mathbf{h}_{t+1}\right)\,.
$$

This readily implies that the shifted mechanism $\bar{\boldsymbol{\pi}}$ is PBIC, since the original mechanism $\boldsymbol{\pi}$ is PBIC.

We conclude by proving the claim. We have that $V_{i,t+1}\left(\boldsymbol{\pi}, (\sigma_i, \mathbf{I}_{-i})|v_i, \mathbf{h}_{t+1}\right) = (I) + (II)$ where the first term is given by

$$
\begin{aligned}
(I) &= (1-\beta)v_i\mathbb{E}^{\boldsymbol{\pi},(\sigma_i,\mathbf{I}_{-i})}\left[\pi_{i,t+1}\big((\sigma_{i,t+1}(v_i, \mathbf{h}_{t+1}), \mathbf{v}_{-i,t+1}), \mathbf{h}_{t+1}\big)|v_i, \mathbf{h}_{t+1}\right] \\
&= (1-\beta)v_i\mathbb{E}^{\bar{\boldsymbol{\pi}},(\bar{\sigma}_i,\mathbf{I}_{-i})}\left[\bar{\pi}_{i,t}\big((\bar{\sigma}_{i,t}(v_i, \bar{\mathbf{h}}_t), \mathbf{v}_{-i,t}), \bar{\mathbf{h}}_t\big)|v_i, \bar{\mathbf{h}}_t\right]\,,
\end{aligned}
$$

where the second equality follows from our definition of the shifted strategy, mechanism and histories and using that values are identically distributed. The second term is given by

$$
\begin{aligned}
(II) &= (1-\beta)\sum_{\ell=t+2}^{\infty}\beta^{\ell-t-1}\mathbb{E}^{\boldsymbol{\pi},(\sigma_i,\mathbf{I}_{-i})}\left[v_{i,\ell}\pi_{i,\ell}\big((\sigma_{i,\ell}(v_{i,\ell},\mathbf{h}_\ell),\mathbf{v}_{-i,\ell}),\mathbf{h}_\ell\big)|v_i,\mathbf{h}_{t+1}\right] \\
&= (1-\beta)\sum_{\ell=t+1}^{\infty}\beta^{\ell-t}\mathbb{E}^{\boldsymbol{\pi},(\sigma_i,\mathbf{I}_{-i})}\left[v_{i,\ell}\pi_{i,\ell+1}\big((\sigma_{i,\ell+1}(v_{i,\ell},\mathbf{h}_{\ell+1}),\mathbf{v}_{-i,\ell}),\mathbf{h}_{\ell+1}\big)|v_i,\mathbf{h}_{t+1}\right] \\
&= (1-\beta)\sum_{\ell=t+1}^{\infty}\beta^{\ell-t}\mathbb{E}^{\bar{\boldsymbol{\pi}},(\bar{\sigma}_i,\mathbf{I}_{-i})}\left[v_{i,\ell}\bar{\pi}_{i,\ell}\big((\bar{\sigma}_{i,\ell}(v_{i,\ell},\bar{\mathbf{h}}_\ell),\mathbf{v}_{-i,\ell}),\bar{\mathbf{h}}_\ell\big)|v_i,\bar{\mathbf{h}}_t\right]
\end{aligned}
$$

where the second equation follows from shifting the index in the summation and using that values are identically distributed, and the third equation follows from our definition of the shifted mechanism, the shifted strategies and the shifted history. We thus conclude that $(I)+(II) = V_{i,t}\left(\bar{\boldsymbol{\pi}},(\bar{\sigma}_i,\mathbf{I}_{-i})|v_i,\bar{\mathbf{h}}_t\right)$ and the claim follows. □

6.2.3 Proof of Proposition 3.3

Proof. As noted by Boyd and Vandenberghe (2009), the convexity of support function follows because it is the point-wise supremum of a family of linear functions. We next show that ϕ is twice differentiable in the following steps.

Step 1. In this step, we show that $\phi(\boldsymbol{\alpha})$ is differentiable and provide its derivative with respect to $\boldsymbol{\alpha}$. Shapiro et al. (2014, Theorem 7.46, p. 371) show that a finite valued, well defined function, $\tilde{f}(\boldsymbol{\alpha}) = \mathbb{E}_{\mathbf{v}}[\tilde{F}(\boldsymbol{\alpha},\mathbf{v})]$, is differentiable at a point $\boldsymbol{\alpha}_0$ if and only if the random function, $\tilde{F}(\boldsymbol{\alpha},\mathbf{v})$, is convex and differentiable at $\boldsymbol{\alpha}_0$ with probability 1. Here, the convexity of the random function means that $\tilde{F}(\cdot,\mathbf{v})$ is convex for almost every $\mathbf{v}$. In our case, the random function is $\max_i(\alpha_i v_i)$, and the corresponding expected value function is $\phi(\boldsymbol{\alpha})$. First, note that $\phi(\boldsymbol{\alpha})$ is well defined and finite valued because $0 \le \phi(\boldsymbol{\alpha}) \le \max_j \alpha_j \mathbb{E}[\max_i v_i]$ for all $\boldsymbol{\alpha} \in \mathbb{R}^n_+$. Second, the random function $\max_i \alpha_i v_i$ is a convex function of $\boldsymbol{\alpha}$ for all $\mathbf{v}$ because it is maximum of linear functions of $\boldsymbol{\alpha}$ (see Boyd and Vandenberghe, 2009). Moreover, it is differentiable at a fixed $\boldsymbol{\alpha}_0$ with probability 1, because for a given $\mathbf{v}$, the partial derivative of $\max_i \alpha_i v_i$ with respect to α_i at $\boldsymbol{\alpha}_0$ does not exist only when $\alpha_{0,i}v_i = \alpha_{0,j}v_j$ for some $j \ne i$; and the probability of this event is 0 for a continuous distribution of $\mathbf{v}$. Therefore, the derivative of $\phi(\boldsymbol{\alpha})$ exists and is given by

$$
h_i(\boldsymbol{\alpha}) \triangleq \frac{\partial\phi}{\partial\alpha_i}(\boldsymbol{\alpha}) = \mathbb{E}_{\mathbf{v}}\Big[v_i\mathbf{1}\{\alpha_i v_i \ge \max_{j\ne i}\alpha_j v_j\}\Big]\,.
$$

Step 2. We next consider the second derivative. Taking expectation of $v_i \mathbf{1}\{\alpha_i v_i \geq \max_{j \neq i} \alpha_j v_j\}$ with respect to $\mathbf{v}_{-i}$ in $h_i(\boldsymbol{\alpha})$, we obtain $h_i(\boldsymbol{\alpha}) = \mathbb{E}_{v_i}\left[v_i \prod_{j \neq i} F_j\left(\frac{\alpha_i v_i}{\alpha_j}\right)\right]$. The derivatives of the integrand with respect to α_i and α_j exist for almost all v_i, and are given by

$$\frac{\partial}{\partial \alpha_i}\left(v_i \prod_{j \neq i} F_j\left(\frac{\alpha_i v_i}{\alpha_j}\right) f_i(v_i)\right) = \sum_{j \neq i} \frac{v_i^2}{\alpha_j} f_j\left(\frac{\alpha_i v_i}{\alpha_j}\right) \mathbf{1}\{\alpha_i v_i \leq \alpha_j \bar{v}\} \prod_{k \neq i,j} F_k\left(\frac{\alpha_i v_i}{\alpha_k}\right) f_i(v_i)\,,$$

$$\frac{\partial}{\partial \alpha_j}\left(v_i \prod_{j \neq i} F_j\left(\frac{\alpha_i v_i}{\alpha_j}\right) f_i(v_i)\right) = -\frac{\alpha_i v_i^2}{\alpha_j^2} f_j\left(\frac{\alpha_i v_i}{\alpha_j}\right) \mathbf{1}\{\alpha_i v_i \leq \alpha_j \bar{v}\} \prod_{k \neq i,j} F_k\left(\frac{\alpha_i v_i}{\alpha_k}\right) f_i(v_i)\,.$$

Moreover, these derivatives are locally integrable, that is, for all compact intervals $[\underline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\alpha}}]$ contained in $\mathbb{R}^n_{>0}$ the following integrals are finite.

$$\int_{\underline{\boldsymbol{\alpha}}}^{\overline{\boldsymbol{\alpha}}} \int_0^{\bar{v}} \sum_{j \neq i} \frac{v_i^2}{\alpha_j} f_j\left(\frac{\alpha_i v_i}{\alpha_j}\right) \mathbf{1}\{\alpha_i v_i \leq \alpha_j \bar{v}\} \prod_{k \neq i,j} F_k\left(\frac{\alpha_i v_i}{\alpha_k}\right) f_i(v_i) \mathrm{d}v_i \mathrm{d}\boldsymbol{\alpha} \leq \frac{\bar{v}^3 \bar{f}^2 (n-1)}{\min_{j=1\ldots n} \underline{\alpha}_j} \prod_{j=1}^n \overline{\alpha}_j < \infty$$

$$\int_{\underline{\boldsymbol{\alpha}}}^{\overline{\boldsymbol{\alpha}}} \int_0^{\bar{v}} \frac{\alpha_i v_i^2}{\alpha_j^2} f_j\left(\frac{\alpha_i v_i}{\alpha_j}\right) \mathbf{1}\{\alpha_i v_i \leq \alpha_j \bar{v}\} \prod_{k \neq i,j} F_k\left(\frac{\alpha_i v_i}{\alpha_k}\right) f_i(v_i) \mathrm{d}v_i \mathrm{d}\boldsymbol{\alpha} \leq \frac{\max_{j=1\ldots n} \overline{\alpha}_j}{\min_{j=1\ldots n} \overline{\alpha}_j} \bar{v}^3 \bar{f}^2 \prod_{j=1}^n \overline{\alpha}_j < \infty$$

In these inequalities, we use the fact that the density functions $f_i(\cdot)$'s are bounded by $\bar{f}$ and the support of the values are bounded by $\bar{v}$.

Therefore Leibniz rule implies that $h_i(\boldsymbol{\alpha})$ is differentiable and hence $\phi(\boldsymbol{\alpha})$ is twice differentiable. Finally, we provide the components of the Hessian matrix of $\phi(\boldsymbol{\alpha})$.

$$\frac{\partial h_i}{\partial \alpha_i}(\boldsymbol{\alpha}) = \frac{1}{(\alpha_i)^3} \sum_{j \neq i} \frac{1}{\alpha_j} \int_0^{\bar{v} \min(\alpha_i, \alpha_j)} x^2 f_i\left(\frac{x}{\alpha_i}\right) f_j\left(\frac{x}{\alpha_j}\right) \prod_{t \neq i,j} F_t\left(\frac{x}{\alpha_t}\right) \mathrm{d}x\,,$$

$$\frac{\partial h_i}{\partial \alpha_j}(\boldsymbol{\alpha}) = -\frac{\alpha_i}{(\alpha_j)^2 (\alpha_i)^3} \int_0^{\bar{v} \min(\alpha_i, \alpha_j)} x^2 f_i\left(\frac{x}{\alpha_i}\right) f_j\left(\frac{x}{\alpha_j}\right) \prod_{t \neq i,j} F_t\left(\frac{x}{\alpha_t}\right) \mathrm{d}x\,. \qquad \square$$

6.3 Appendix for Chapter 3 Section 3.2

6.3.1 Proof of Proposition 3.4

Proof. We show that there exist $\boldsymbol{a}$ and b such that $\boldsymbol{a}^\intercal \hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}) + b = 0$. Suppose that $\boldsymbol{a} = \boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$ and $b = \sum_{i=1}^{n} a_i \big[(1-\beta)(\nabla\phi(\boldsymbol{a}))_i - u_i\big]/\beta$.

Using the definition of $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$, we obtain $\boldsymbol{a}^\intercal \hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}) = \sum_{i=1}^{n} a_i \mathbb{E}_{\hat{v}_i}[\hat{W}_i(\hat{v}_i|\mathbf{u})]$. Moreover, (PK($\mathbf{u}$)) implies that $\beta\mathbb{E}_{\hat{v}_i}[\hat{W}_i(\hat{v}_i|\mathbf{u})] = \Big[u_i - (1-\beta)\mathbb{E}[v_i \mathbf{1}\{a_i v_i \geq \max_{j\neq i} a_j v_j\}]\Big]$. Finally, recall that the support function $\phi(\cdot)$ satisfies $\left(\nabla\phi(\boldsymbol{a})\right)_i = \mathbb{E}\Big[v_i \mathbf{1}\{a_i v_i \geq \max_{j\neq i} a_j v_j\}\Big]$. Combining these three observations, it follows that $\boldsymbol{a}^\intercal \hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}) + b = 0$.

The second part of the proposition follows from the following properties of the support function $\phi(\cdot)$. Note that $\boldsymbol{x}^\intercal \nabla\phi(\boldsymbol{x}) = \phi(\boldsymbol{x})$ for all $\boldsymbol{x}$ by the definition of $\phi(\cdot)$. Furthermore, for a vector $\boldsymbol{x}$, the support function satisfies that $\phi(\boldsymbol{x}) \geq \boldsymbol{x}^\intercal \boldsymbol{z}$ for all $\boldsymbol{z} \in \hat{\mathcal{U}}_\beta$ (see Schneider, 2013). Because $\mathbf{u} \in \hat{\mathcal{U}}_\beta$, we obtain the following inequality, which concludes the proof.

$$\big(\nabla\phi(\boldsymbol{a}) - \mathbf{u}\big)^\intercal \boldsymbol{a} = \nabla\phi(\boldsymbol{a})^\intercal \boldsymbol{a} - \mathbf{u}^\intercal \boldsymbol{a} = \phi(\boldsymbol{a}) - \mathbf{u}^\intercal \boldsymbol{a} \geq 0\,. \qquad \square$$

6.3.2 Proof of Proposition 3.5

Proof. For simplicity, we use $\boldsymbol{a}$ to represent $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$ in this proof. First, we reformulate problem (3.12) as the following linear programming problem:

$$\max_{\tilde{\mathbf{w}}(\cdot)} \quad z$$

$$\text{st.} \quad \mathbb{E}[\tilde{w}_i(v_i, \mathbf{v}_{-i})] = \hat{W}_i(v_i|\mathbf{u}) \quad \forall v_i, i \tag{6.6}$$

$$z \leq \boldsymbol{a}^\intercal(\mathbf{u} - \tilde{\mathbf{w}}(\mathbf{v})) \qquad \forall \mathbf{v} \tag{6.7}$$

Next, we find an upper bound for this problem by relaxing the constraints over z. Because $\mathbf{u} \in \mathcal{E}(\hat{\mathcal{U}}_\beta)$, we have $\mathbf{u}/\mathsf{F}_\beta \in \mathcal{E}(\mathcal{U})$, which implies $\mathbf{u}/\mathsf{F}_\beta = \nabla\phi(\boldsymbol{a})$, i.e., $u_i = \mathsf{F}_\beta \mathbb{E}[v_i \mathbf{1}\{a_i v_i \geq \max_{j\neq i} a_j v_j\}]$. Moreover, $u_i = (1-\beta)\mathbb{E}[v_i \mathbf{1}\{a_i v_i \geq \max_{j\neq i} a_j v_j\}] + \beta\mathbb{E}[\hat{W}_i(v_i|\mathbf{u})]$, following (PK($\mathbf{u}$)). These two observations imply that $u_i = \Upsilon \mathbb{E}[\hat{W}_i(v_i|\mathbf{u})]$, where $\Upsilon = \dfrac{\beta\mathsf{F}_\beta}{\mathsf{F}_\beta - (1-\beta)}$. Therefore, constraint (6.7) becomes $z \leq \sum_{i=1}^{n} a_i \big(\Upsilon \mathbb{E}[\hat{W}_i(v_i|\mathbf{u})] - \tilde{w}_i(\mathbf{v})\big)$ for all $\mathbf{v}$. Because z has to satisfy this constraint for all $\mathbf{v}$, we can relax it by replacing these constraints with a single constraint where the right hand side is the expectation over $\mathbf{v}$. This operation corresponds to taking expectation of $\tilde{w}_i(\mathbf{v})$. Following constraint (6.6), we obtain $z \leq \sum_{i=1}^{n} a_i(\Upsilon - 1)\mathbb{E}[\hat{W}_i(v_i|\mathbf{u})]$. Therefore, it follows that $\sum_{i=1}^{n} a_i(\Upsilon - 1)\mathbb{E}[\hat{W}_i(v_i|\mathbf{u})]$ is an upper bound for the optimal value.

To show $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ is an optimal solution, we first show that $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ is feasible, and the objective function evaluated at $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ is equal to this upper bound. By its definition, the expectation of $\hat{w}_i(\mathbf{v}|\mathbf{u})$ with respect to $\mathbf{v}_{-i}$ is $\hat{W}_i(v_i|\mathbf{u})$, which implies feasibility. By Proposition 3.4, we also know that $\boldsymbol{a}^\intercal(\mathbf{u} - \hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})) = \sum_{i=1}^{n} a_i(\Upsilon - 1)\mathbb{E}[\hat{W}_i(v_i|\mathbf{u})]$ which, in turn, implies that $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ is an optimal solution for (3.12). $\square$

6.4 Appendix for Chapter 3 Section 3.3

In Section 3.3, we introduce the signed distance function for the perfect information achievable set $\mathcal{U}$, in (3.14). In fact, this function can be defined for any convex set $\mathcal{C}$. In Appendix 6.4.1 we provide the definition and a key property of the function, which is used to show that a point belongs to a set. We then focus on the signed distance function for (scaled) perfect information achievable sets, and prove the main results of Section 3.3.

6.4.1 Signed Distance Function for Convex Sets

For a convex set $\mathcal{C} \in \mathbb{R}^n$, Schneider (2013) defines its support function as

$$\psi(\boldsymbol{\lambda}) \triangleq \sup\{\boldsymbol{\lambda}^\intercal \mathbf{x} : \mathbf{x} \in \mathcal{C}\}\,.$$

We further define the following function

$$I_{\mathcal{C}}(\mathbf{x}) \triangleq \sup_{\boldsymbol{\lambda} \geq 0\,, \|\boldsymbol{\lambda}\|=1} \{\mathbf{x}^\intercal \boldsymbol{\lambda} - \psi(\boldsymbol{\lambda})\}\,.$$

The following lemma demonstrates that function $I_{\mathcal{C}}$ is indeed a signed distance function.

Lemma 6.4. *For any convex set $\mathcal{C} \subset \mathbb{R}^n$, we have*

$$I_{\mathcal{C}}(\mathbf{x}) = \begin{cases} \leq 0 & \text{if } \mathbf{x} \in hyp(\mathcal{C})\,, \\ > 0 & \text{if } \mathbf{x} \notin hyp(\mathcal{C})\,, \end{cases}$$

where $hyp(\mathcal{C}) = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x} \leq \bar{\mathbf{x}}\,, \exists \bar{\mathbf{x}} \in \mathcal{C}\}$.

Proof. We prove this result by considering two separate cases which determines the sign of the indicator function.

Step 1. In this step, we assume that $\mathbf{x} \in \text{hyp}(\mathcal{C})$. The definition of the $\text{hyp}(\mathcal{C})$ implies the existence of $\bar{\mathbf{x}} \in \mathcal{C}$ satisfying $\mathbf{x} \leq \bar{\mathbf{x}}$. Fixing an arbitrary nonnegative vector $\boldsymbol{\lambda}$ satisfying $\|\boldsymbol{\lambda}\| = 1$, we have

$$\boldsymbol{\lambda}^\intercal \mathbf{x} \leq \boldsymbol{\lambda}^\intercal \bar{\mathbf{x}} \leq \psi(\boldsymbol{\lambda})\,.$$

Therefore, we have $I_{\mathcal{C}}(\mathbf{x}) \leq 0$.

Step 2. Now assume that $\mathbf{x} \notin \text{hyp}(\mathcal{C})$. The separating hyperplane theorem implies that there exists a nonzero vector $\boldsymbol{\lambda}$ satisfying $\boldsymbol{\lambda}^\intercal \mathbf{x} > \boldsymbol{\lambda}^\intercal \bar{\mathbf{x}}$ for all $\bar{\mathbf{x}} \in \text{hyp}(\mathcal{C})$. We argue that $\boldsymbol{\lambda} \geq 0$, following the fact that $\text{hyp}(\mathcal{C})$ is unbounded from below. Otherwise, for the component $\lambda_i < 0$, we can always find a low enough negative $\bar{x}_i$ to violate $\boldsymbol{\lambda}^\intercal \mathbf{x} > \boldsymbol{\lambda}^\intercal \bar{\mathbf{x}}$.

Let $\bar{\boldsymbol{\lambda}} \triangleq \boldsymbol{\lambda}/\|\boldsymbol{\lambda}\|$, such that $\|\bar{\boldsymbol{\lambda}}\| = 1$. Moreover, for all $\bar{\mathbf{x}} \in \text{hyp}(\mathcal{C})$, we have

$$\bar{\boldsymbol{\lambda}}^\intercal \mathbf{x} > \bar{\boldsymbol{\lambda}}^\intercal \bar{\mathbf{x}} \,.$$

These observations together imply that $\bar{\boldsymbol{\lambda}}^\intercal \mathbf{x} > \psi(\bar{\boldsymbol{\lambda}})$, thus $I_{\mathcal{C}}(\mathbf{x}) > 0$. □

6.4.2 The Signed Distance Function for the Scaled Perfect Information Set

In this section, we consider the perfect information achievable set $\mathcal{U}$ defined in (3.4). We provide the signed distance function for this set and the sets obtained by scaling it down. First, we start by showing that the perfect information achievable set and the sets obtained by scaling it down are convex.

Proposition 6.1. *The perfect information achievable set $\mathcal{U}$ defined in* (3.4) *is a convex set. Moreover, for any $s \in [0,1]$, the set $s\mathcal{U}$ is also convex, and $s\mathcal{U} = hyp(s\mathcal{U}) \cap \mathbb{R}^n_+$.*

Proof. For any two points $\mathbf{u}', \mathbf{u}'' \in \mathcal{U}$ and a scalar $x \in (0,1)$, let $\mathbf{p}'$ and $\mathbf{p}''$ be the allocation functions corresponding to the points $\mathbf{u}'$ and $\mathbf{u}''$, respectively. Then, the point $x\mathbf{u}' + (1-x)\mathbf{u}''$ is in $\mathcal{U}$ because $x\mathbf{p}' + (1-x)\mathbf{p}''$ is an allocation satisfying (FA), and $xu'_i + (1-x)u''_i = \mathbb{E}[v_i(xp'_i(\mathbf{v}) + (1-x)p''_i(\mathbf{v}))]$. Because $\mathcal{U}$ is convex, $s\mathcal{U}$ is also convex.

Any set is a subset of its hypograph. And $s\mathcal{U} \subseteq \mathbb{R}^n_+$. Thus, it is sufficient to show that $\text{hyp}(s\mathcal{U}) \cap \mathbb{R}^n_+ \subseteq s\mathcal{U}$. Let $\mathbf{u} \in \text{hyp}(s\mathcal{U}) \cap \mathbb{R}^n_+$. This implies that $\mathbf{u} \geq \mathbf{0}$ and there exists $\bar{\mathbf{u}} \in s\mathcal{U}$ such that $\bar{\mathbf{u}} \geq \mathbf{u}$. Let $\bar{\mathbf{p}}$ be the allocation corresponding to $\bar{\mathbf{u}}$.

Consider the 2^n points that can be obtained by replacing each subset of components of $\bar{\mathbf{u}}$ with zero. For example, when $n = 2$, we get $\bar{\mathbf{u}}^0 = (0,0)$, $\bar{\mathbf{u}}^1 = (0, \bar{u}_2)$, $\bar{\mathbf{u}}^2 = (\bar{u}_1, 0)$, $\bar{\mathbf{u}}^4 = \bar{\mathbf{u}}$. All these 2^n states are in $s\mathcal{U}$. In particular, for a state $\bar{\mathbf{u}}^m$, we can construct a feasible allocation $\bar{\mathbf{p}}^m$ satisfying $\bar{\mathbf{u}}^m_i = \mathbb{E}[v_i p^m_i(\mathbf{v})]$ for all i by setting $\bar{p}^m_i = 0$ if $\bar{u}^m_i = 0$ and $\bar{p}^m_i = \bar{p}_i$ otherwise. Therefore, $\bar{\mathbf{u}}^m \in s\mathcal{U}$ for all m. The polytope whose extreme points are the states $\bar{\mathbf{u}}^m$ given by $\{\tilde{\mathbf{u}} \in \mathbb{R}^n_+ : 0 \leq \tilde{u}_i \leq \bar{u}_i \ \ \forall i\}$

is a subset of $s\mathcal{U}$ because all extreme points are in $s\mathcal{U}$ and $s\mathcal{U}$ is convex. The result follows because $\mathbf{u}$ lies in this polytope. $\square$

We now provide the signed distance function corresponding to $\mathsf{F}\mathcal{U}$ for a given scalar F as follows:

$$\mathfrak{J}(\mathbf{u},\mathsf{F}) \triangleq I_{\mathcal{U}}\left(\frac{\mathbf{u}}{\mathsf{F}}\right) = \sup_{\|\boldsymbol{x}\|_1=1,\boldsymbol{x}\geq\mathbf{0}} \left\{\frac{\mathbf{u}^{\intercal}\boldsymbol{x}}{\mathsf{F}} - \phi(\boldsymbol{x})\right\}\,, \tag{6.8}$$

where $\phi(\cdot)$ is the support function of $\mathcal{U}$ given in (3.5).

By Lemma 6.4 in Appendix 6.4.1, for any nonnegative $\mathbf{u}$, it follows that

$$\mathfrak{J}(\mathbf{u},\mathsf{F}) = \begin{cases} \leq 0 & \mathbf{u} \in \operatorname{hyp}(\mathsf{F}\mathcal{U})\,, \\ > 0 & \mathbf{u} \notin \operatorname{hyp}(\mathsf{F}\mathcal{U})\,. \end{cases}$$

Corollary 6.1. *For a point* $\mathbf{u} \in \mathbb{R}^n_+$*,* $\mathbf{u} \in \mathsf{F}\mathcal{U}$ *if and only if* $\mathfrak{J}(\mathbf{u},\mathsf{F}) \leq 0$*.*

Proof. Proposition 6.1 implies that $\mathsf{F}\mathcal{U} = \operatorname{hyp}(\mathsf{F}\mathcal{U}) \cap \mathbb{R}^n_+$, thus for any $\mathbf{u} \geq \mathbf{0}$, $\mathbf{u} \in \mathsf{F}\mathcal{U}$ if and only if $\mathbf{u} \in \operatorname{hyp}(\mathsf{F}\mathcal{U})$. Moreover, Lemma 6.4 implies that $\mathfrak{J}(\mathbf{u},\mathsf{F}) \leq 0$ if and only if $\mathbf{u} \in \operatorname{hyp}(\mathsf{F}\mathcal{U})$. $\square$

Finally, we provide two properties of $\mathfrak{J}(\mathbf{u},\mathsf{F})$ by the following propositions. Before stating these results, we introduce the following matrix that would be used in later results:

$$\Pi(\boldsymbol{a}) \triangleq \operatorname{Hess}(\phi(\boldsymbol{a})) + \boldsymbol{e}\boldsymbol{e}^{\intercal}\,, \tag{6.9}$$

where $\boldsymbol{a} > \mathbf{0}$ and $\boldsymbol{e}$ is the n dimensional vector of ones.

For any scalar $\mathsf{F} \in (0,1]$, define set $\mathcal{S}(\mathsf{F}) \in \mathbb{R}^n$ as

$$\mathcal{S}(\mathsf{F}) \triangleq \{\boldsymbol{z} + a\boldsymbol{e} : \forall a \geq 0, \boldsymbol{z} \in \mathsf{F}\mathcal{U},\ \text{and}\ \boldsymbol{z} > 0\}\,.$$

Set $\mathcal{S}(\mathsf{F})$ is useful in the following result as well as later in the Appendices. Specifically, there are several technicalities in our analysis for which it is necessary to have a positive optimal solution in $\mathfrak{J}(\mathbf{u},\mathsf{F})$ (i.e., $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}) > \mathbf{0}$) such as uniqueness of $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F})$ and twice differentiability of $\mathfrak{J}(\mathbf{u},\mathsf{F})$ (see Proposition 6.3). Although this condition trivially holds for positive points $\mathbf{u} \in \mathsf{F}\mathcal{U}$, it is not necessarily true for all positive $\mathbf{u} \notin \mathsf{F}\mathcal{U}$. For example, our construction of the future promises $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ are not automatically guaranteed to be in the set $\mathsf{F}\mathcal{U}$. According to Remark 3.2, $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}) > \mathbf{0}$ if we can reach the scaled efficient frontier by moving from $\mathbf{u}$ along a 45° line. The set $\mathcal{S}(\mathsf{F})$ exactly gives the set of all points that satisfy this property. Proposition 6.2 formally establishes that points in $\mathcal{S}(\mathsf{F})$ satisfy $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}) > \mathbf{0}$.

Proposition 6.2. *Let* $\mathbf{u} \in \mathcal{S}(\mathsf{F})$*. We have* $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}) > \mathbf{0}$*, in which* $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F})$ *is defined in (3.6).*

Proof. Consider a point $\mathbf{u} = \boldsymbol{z} + a\boldsymbol{e}$ for some fixed $\boldsymbol{z} \in \mathsf{F}\mathcal{U}$, $\boldsymbol{z} > 0$ and $a \geq 0$. The optimization problem in (3.6) becomes

$$\sup_{\|\boldsymbol{x}\|_1=1, \boldsymbol{x} \geq \mathbf{0}} \left\{ \frac{(\boldsymbol{z} + a\boldsymbol{e})^\intercal \boldsymbol{x}}{\mathsf{F}} - \phi(\boldsymbol{x}) \right\} = \frac{a}{\mathsf{F}} + \sup_{\|\boldsymbol{x}\|_1=1, \boldsymbol{x} \geq \mathbf{0}} \left\{ \frac{\boldsymbol{z}^\intercal \boldsymbol{x}}{\mathsf{F}} - \phi(\boldsymbol{x}) \right\}.$$

Therefore, $\boldsymbol{\alpha}^*((\boldsymbol{z} + a\boldsymbol{e})/\mathsf{F}) = \boldsymbol{\alpha}^*(\boldsymbol{z}/\mathsf{F})$. Now focus on $\boldsymbol{z} > 0$ such that $\boldsymbol{z} \in \mathsf{F}\mathcal{U}$. The Lagrangian is given by $\mathcal{L}(\boldsymbol{x}, \mu, \boldsymbol{\zeta}) = \boldsymbol{x}^\intercal \boldsymbol{z} - \mathsf{F}\phi(\boldsymbol{x}) - \mu(1 - \boldsymbol{e}^\intercal \boldsymbol{x}) + \boldsymbol{\zeta}^\intercal \boldsymbol{x}$, where μ and $\boldsymbol{\zeta}$ are the Lagrange multipliers corresponding to constraints $\|\boldsymbol{x}\|_1 = 1$ and $\boldsymbol{x} \geq 0$, respectively. Optimality implies that $\nabla_{\boldsymbol{x}} \mathcal{L}(\boldsymbol{\alpha}^*(\boldsymbol{z}/\mathsf{F}), \mu^*, \boldsymbol{\zeta}^*) = \boldsymbol{z} - \mathsf{F}\nabla\phi(\boldsymbol{\alpha}^*(\boldsymbol{z}/\mathsf{F})) + \mu^* \boldsymbol{e} + \boldsymbol{\zeta}^* = 0$. As discussed in Remark 3.2, the optimal value μ^* is nonnegative for $\boldsymbol{z} \in \mathsf{F}\mathcal{U}$. Suppose now there exists a component of $\boldsymbol{\alpha}^*(\boldsymbol{z}/\mathsf{F})$ such that $\alpha_i^*(\boldsymbol{z}/\mathsf{F}) = 0$. In this case, the optimality condition for i implies $z_i + \mu^* + \xi_i^* = 0$, because $\left(\nabla\phi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}))\right)_i = 0$. This is a contradiction because $z_i > 0$, $\mu^* \geq 0$ and $\zeta_i^* \geq 0$. Therefore, we have the result. ☐

Proposition 6.3. *Let* F *be a positive scalar in* $[0, 1]$ *and* $\mathbf{u} \in \mathcal{S}(\mathsf{F})$*. Suppose also that* $\Pi(\boldsymbol{a})$ *is positive definite for all* $\boldsymbol{a} > \mathbf{0}$*. Then, the optimal solution* $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F})$ *is unique. Moreover, the function* $\mathfrak{J}(\mathbf{u}, \mathsf{F})$ *is twice differentiable, and the first and second derivatives are given as follows:*

$$\nabla\mathfrak{J}(\mathbf{u}, \mathsf{F}) = \frac{\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F})}{\mathsf{F}} \text{ and } \mathit{Hess}(\mathfrak{J}(\mathbf{u}, \mathsf{F})) = \frac{1}{\mathsf{F}^2} \Pi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}))^{-1} \left(I - \frac{\boldsymbol{e}\boldsymbol{e}^\intercal \Pi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}))^{-1}}{\boldsymbol{e}^\intercal \Pi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}))^{-1} \boldsymbol{e}} \right),$$

where I *is the* $n \times n$ *identity matrix.*

Proof. We prove this result in three steps.

Step 1. In this step, we show that the optimal solution $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F})$ of the maximization problem in $\mathfrak{J}(\mathbf{u}, \mathsf{F})$ is unique. In Proposition 6.2, we show that an optimal solution is $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}) > \mathbf{0}$ for $\mathbf{u} \in \mathcal{S}(\mathsf{F})$, thus we can ignore the non-negativity constraints. Because $\boldsymbol{x} > \mathbf{0}$, Proposition 3.3 implies that the objective is twice-differentiable with hessian $-\text{Hess}(\phi(\boldsymbol{x}))$. Note that the objective function is not strictly concave over $\mathbb{R}^n_{>0}$ because $\text{Hess}(\phi(\boldsymbol{a}))\boldsymbol{a} = \mathbf{0}$ for any $\boldsymbol{a} \neq \mathbf{0}$. However, we can show that the restriction of the objective function to the feasible set $\|\boldsymbol{x}\|_1 = 1$ is strictly concave. Let $\boldsymbol{y} \in \mathbb{R}^n$ be a feasible direction satisfying $\boldsymbol{e}^\intercal \boldsymbol{y} = 0$. We obtain using (6.9) that

$$\boldsymbol{y}^\intercal \mathrm{Hess}(\phi(\boldsymbol{x}))\boldsymbol{y} = \boldsymbol{y}^\intercal \Pi(\boldsymbol{x})\boldsymbol{y} - (\boldsymbol{e}^\intercal \boldsymbol{y})^2 = \boldsymbol{y}^\intercal \Pi(\boldsymbol{x})\boldsymbol{y} > 0\,,$$

where the last inequality follows because $\Pi(\boldsymbol{x})$ is positive definite for all $\boldsymbol{x} > 0$. This implies that the optimal solution is unique.

Step 2. In this step, we show that $\mathfrak{J}(\mathbf{u}, \mathsf{F})$ is differentiable with respect to $\mathbf{u}$ and provide $\nabla\mathfrak{J}(\mathbf{u}, \mathsf{F})$. As noted by Milgrom and Segal (2002) in Corollary 3, the gradient of $\mathfrak{J}(\mathbf{u}, \mathsf{F})$ is well defined because (i) $\{\boldsymbol{x} : \|\boldsymbol{x}\|_1 = 1, \boldsymbol{x} \geq \mathbf{0}\}$ is a convex set, (ii) $\boldsymbol{x}^\intercal \mathbf{u}/\mathsf{F} - \phi(\boldsymbol{x})$ is a concave function of $\boldsymbol{x}$ and $\mathbf{u}$, and (iii) there is an optimal solution, $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F})$, such that $\nabla\mathfrak{J}(\mathbf{u}, \mathsf{F})$ exists. Then, the gradient at $\mathbf{u}$ is given by

$$\nabla\mathfrak{J}(\mathbf{u}, \mathsf{F}) = \frac{\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F})}{\mathsf{F}}\,. \tag{6.10}$$

Step 3. In this step, we show the existence of the second derivative and provide the Hessian matrix of $\mathfrak{J}(\mathbf{u}, \mathsf{F})$. In the previous step, the gradient of $\mathfrak{J}(\mathbf{u}, \mathsf{F})$ is given by $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F})/\mathsf{F}$. Therefore, the Hessian matrix consists of the derivatives of the optimal solution $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F})$ with respect to $\mathbf{u}$. We next characterize the optimal solution using the KKT conditions of the optimization problem.

Because $\boldsymbol{x} > \mathbf{0}$, we can write the Lagrangian of (6.8) as $\mathcal{L}(\boldsymbol{x}, \mu) = \boldsymbol{x}^\intercal \mathbf{u} - \mathsf{F}\phi(\boldsymbol{x}) - \mu(1 - \boldsymbol{e}^\intercal \boldsymbol{x})$. Optimality implies $\nabla_{\boldsymbol{x}}\mathcal{L}(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}), \mu^*) = \mathbf{u} - \mathsf{F}\nabla\phi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F})) + \mu^*\boldsymbol{e} = 0$ and $\nabla_\mu\mathcal{L}((\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}), \mu^*)) = \boldsymbol{e}^\intercal\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}) - 1 = 0$. Introduce an auxiliary function $Q : \mathbb{R}^{n+1} \to \mathbb{R}^{n+1}$ as follows:

$$Q_i(\boldsymbol{x}, \mu) \triangleq \mathsf{F}(h_i(\boldsymbol{x}) - \mu) = z_i\,, \qquad \forall i = 1, \ldots, n\,,$$

$$Q_{n+1}(\boldsymbol{x}, \mu) \triangleq -\mathsf{F}\sum_{i=1}^{n} x_i = z_{n+1}\,.$$

For a given $\mathbf{u}$, it follows that $Q(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}), \mu^*) = [\mathbf{u}^\intercal, -\mathsf{F}]^\intercal$. The inverse function theorem implies that if the Jacobian determinant of Q is non-zero, then Q is invertible. Specifically, if the Jacobian determinant of Q is non-zero, Q is invertible and thus the unique optimal solution is obtained by $Q^{-1}(\mathbf{u}, -\mathsf{F}) = [\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F})^\intercal, \mu^*]^\intercal$. We next argue that the Jacobian determinant of Q, denoted by

$\det(J_Q(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}), \mu^*))$, is non-zero. The Jacobian matrix of Q, $J_Q(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}))$, is given as follows:

$$J_Q(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}), \mu^*) = \mathsf{F} \begin{bmatrix} \text{Hess}(\phi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}))) & -\boldsymbol{e} \\ -\boldsymbol{e}^\intercal & 0 \end{bmatrix}.$$

Note that $J_Q(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}), \mu^*)$ is a block matrix thus its determinant is expressed in terms of the determinants of $\text{Hess}(\phi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F})))$ and $\Pi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}))$. We have that $\det(\Pi(\boldsymbol{a})) \neq 0$ for all $\boldsymbol{a} > \mathbf{0}$ because the matrix is positive definite. Moreover, it follows that $\det(\text{Hess}(\phi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F})))) = 0$ because $\text{Hess}(\phi(\boldsymbol{a}))\boldsymbol{a} = \mathbf{0}$ for any $\boldsymbol{a} \neq \mathbf{0}$. Therefore, we obtain

$$\det(J_Q(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}), \mu^*)) = \det(\text{Hess}(\phi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F})))) - \det(\Pi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}))) = -\det(\Pi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}))) \neq 0.$$

The inverse function theorem further implies that Q^{-1} is continuously differentiable.

Moreover, the theorem implies $J_{Q^{-1}}(\mathbf{u}, -\mathsf{F}) = [J_Q(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}), \mu^*)]^{-1}$, where $J_{Q^{-1}}$ is the Jacobian matrix of inverse function Q^{-1}. This system of equations implies the existence of $\text{Hess}(\mathfrak{J}(\mathbf{u}, \mathsf{F}))$ because $\text{Hess}(\mathfrak{J}(\mathbf{u}, \mathsf{F}))$ is the upper left sub-matrix of $J_{Q^{-1}}(\mathbf{u}, -\mathsf{F})$. Specifically, denoting $c_i = \dfrac{\partial \mu^*}{\partial u_i}$ and $b_i = \dfrac{\partial \alpha_i^*(\mathbf{u}/\mathsf{F})}{\partial z_{n+1}}$, we can equivalently express $J_{Q^{-1}}(\mathbf{u}, -\mathsf{F}) = [J_Q(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}), \mu^*)]^{-1}$ as follows:

$$\mathsf{F} J_{Q^{-1}}(\mathbf{u}, -\mathsf{F}) = \begin{bmatrix} \mathsf{F}^2 \text{Hess}(\mathfrak{J}(\mathbf{u}, \mathsf{F})) & b \\ c^\intercal & 0 \end{bmatrix} = \left(\begin{bmatrix} \text{Hess}(\phi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}))) & -\boldsymbol{e} \\ -\boldsymbol{e}^\intercal & 0 \end{bmatrix} \right)^{-1} = \mathsf{F} \Big(J_Q(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}), \mu^*) \Big)^{-1}.$$

Following Lemma 3 in Pringle and Rayner (1970), we obtain the following closed form expression for $\text{Hess}(\mathfrak{J}(\mathbf{u}, \mathsf{F}))$:

$$\mathsf{F}^2 \text{Hess}(\mathfrak{J}(\mathbf{u}, \mathsf{F})) = \Pi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}))^{-1} \left(I - \frac{\boldsymbol{e}\boldsymbol{e}^\intercal \Pi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}))^{-1}}{\boldsymbol{e}^\intercal \Pi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}))^{-1} \boldsymbol{e}} \right). \qquad \square$$

Proposition 6.4. *Let* F *be a positive scalar in* $[0, 1]$, $\mathbf{u}' \in \mathcal{S}(\mathsf{F})$ *and* $\mathbf{u}'' \in \mathsf{F}\mathcal{U}$ *such that* $\mathbf{u}'' > \mathbf{0}$. *Suppose* $\Pi(\boldsymbol{a})$ *is positive definite for* $\boldsymbol{a} > \mathbf{0}$. *Then* $\mathfrak{J}(\mathbf{u}', \mathsf{F})$ *is expressed via its quadratic expansion around* $\mathbf{u}''$ *as follows:*

$$\mathfrak{J}(\mathbf{u}', \mathsf{F}) = \mathfrak{J}(\mathbf{u}'', \mathsf{F}) + \nabla \mathfrak{J}(\mathbf{u}'', \mathsf{F})^\intercal (\mathbf{u}' - \mathbf{u}'') + \frac{1}{2} (\mathbf{u}' - \mathbf{u}'')^\intercal \mathit{Hess}(\mathfrak{J}(\boldsymbol{\Theta}, \mathsf{F}))(\mathbf{u}' - \mathbf{u}''),$$

where $\boldsymbol{\Theta}$ *is a point between* $\mathbf{u}'$ *and* $\mathbf{u}''$.

Proof. When $\Pi(\boldsymbol{a})$ is positive definite for all $\boldsymbol{a} > \mathbf{0}$, Proposition 6.3 provides the gradient and the

Hessian of $\mathfrak{J}$ for points in $\mathcal{S}(\mathsf{F})$. Thus, we next show that $\boldsymbol{\Theta} \in \mathcal{S}(\mathsf{F})$. We can express $\boldsymbol{\Theta}$ as a convex combination of $\mathbf{u}'$ and $\mathbf{u}''$, i.e., $\boldsymbol{\Theta} = \omega\mathbf{u}' + (1-\omega)\mathbf{u}''$ for some $\omega \in (0,1)$ because $\boldsymbol{\Theta}$ is a point between $\mathbf{u}'$ and $\mathbf{u}''$. We know that there exist $(\boldsymbol{x}'', a'')$ such that $\mathbf{u}'' = \boldsymbol{x}'' + a''\boldsymbol{e}$ where $\boldsymbol{x}'' \in \mathsf{F}\mathcal{U}$ and $\boldsymbol{x}'' > \mathbf{0}$, and $a'' \geq 0$. Thus it follows that $\boldsymbol{\Theta} = \omega\mathbf{u}' + (1-\omega)\boldsymbol{x}'' + a''\boldsymbol{e}$. Because $\mathsf{F}\mathcal{U}$ is convex (see Proposition 6.1), we have that $\omega\mathbf{u}' + (1-\omega)\boldsymbol{x}'' \in \mathsf{F}\mathcal{U}$ and $\boldsymbol{\Theta} \in \mathcal{S}(\mathsf{F})$. Therefore, following Taylor's theorem, we can express $\mathfrak{J}(\mathbf{u}', \mathsf{F})$ by the quadratic expansion around $\mathbf{u}''$ as follows:

$$\mathfrak{J}(\mathbf{u}', \mathsf{F}) = \mathfrak{J}(\mathbf{u}'', \mathsf{F}) + \nabla\mathfrak{J}(\mathbf{u}'', \mathsf{F})^{\intercal}(\mathbf{u}' - \mathbf{u}'') + \frac{1}{2}(\mathbf{u}' - \mathbf{u}'')^{\intercal}\text{Hess}(\mathfrak{J}(\boldsymbol{\Theta}, \mathsf{F}))(\mathbf{u}' - \mathbf{u}'')\,. \qquad \square$$

6.4.3 Proof of Proposition 3.6

Before the proof, we first introduce a notation for the bounds of the j-th moment for $j > 1$. We denote the minimum and the maximum of the j-th moment of agents' values by $\underline{m}_j = \min_{i\in\{1,\ldots,n\}} \mathbb{E}[v_i^j]$ and $\overline{m}_j = \max_{i\in\{1,\ldots,n\}} \mathbb{E}[v_i^j]$, respectively. We next provide the value of the constant used in (3.13). Following this definition, we provide the proof of the proposition.

Definition 6.1. *Define constants F_β and $\underline{u}_\beta$ such that $\mathsf{F}_\beta = 1 - \underline{u}_\beta/\underline{m}$ and $\underline{u}_\beta = \xi(1-\beta)$, in which scalar ξ is given by*

$$\xi \triangleq \frac{12\nu}{\underline{\beta}\bar{v}^3\underline{f}^2}\,, \tag{6.11}$$

where

$$\nu \triangleq \max\left(2\left[\frac{\bar{f}\bar{v}^2}{2} - \frac{\bar{f}}{2}\underline{m}_2 + \overline{m}\right]^2, \left(\overline{m} + \frac{\bar{f}}{2}\overline{m}_2\right)^2 + (\overline{m} + \bar{v})^2, \frac{\underline{f}^4\bar{v}^6}{36}\right). \tag{6.12}$$

We determine the value of $\underline{\beta}$ such that $\mathsf{F}_\beta \geq 1/2$ for all $\beta \geq \underline{\beta}$.

Proof of Proposition 3.6. We prove this result in two steps. We defer the proofs of lemmas used in these steps to the end of this section.

Step 1. In this step, we prove that $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}) \geq 0$ for all $\mathbf{u} \in \mathsf{F}_\beta\mathcal{U}$. For a fixed $\mathbf{u}$, for simplicity, we use $\hat{\mathbf{w}}$ to represent $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$. Recall that $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ is given in (3.11) for $\mathbf{u} \in \hat{\mathcal{U}}_\beta$, and equals to $\mathbf{w}^{\mathsf{L}}(\mathbf{v}|\mathbf{u}) = \mathbf{u}$ for $\mathbf{u} \in \mathsf{F}_\beta\mathcal{U} \setminus \hat{\mathcal{U}}_\beta$. Suppose that for an i, $u_i \leq \underline{u}_\beta$, then $\hat{\mathbf{w}} = \mathbf{w}^{\mathsf{L}}(\mathbf{v}|\mathbf{u}) = \mathbf{u}$ and $\hat{\mathbf{w}}$ is clearly nonnegative because $\mathbf{u} \in \mathsf{F}_\beta\mathcal{U}$. Suppose that $u_i \geq \underline{u}_\beta$ for all i, i.e., $\mathbf{u}$ lies in the central region

$\hat{\mathcal{U}}_\beta$. We show that the largest jump from the initial state $\mathbf{u}$ to $\hat{\mathbf{w}}$ is less than $\underline{u}_\beta$, thus $\hat{\mathbf{w}}$ is always nonnegative. The following lemma provides a bound on the largest jump.

Lemma 6.5. *For any $\mathbf{u} \in \hat{\mathcal{U}}_\beta$ and $\beta \geq \underline{\beta}$, we have $\|\hat{\mathbf{w}} - \mathbf{u}\|_2 \leq \left(\frac{1-\beta}{\underline{\beta}}\right)\sqrt{\nu}$.*

This lemma enables us to show that $\underline{u}_\beta \geq \|\hat{\mathbf{w}} - \mathbf{u}\|_2$ for all $\mathbf{u} \in \hat{\mathcal{U}}_\beta$ because it provides a uniform bound on this norm which is independent of $\mathbf{u}$ and $\mathbf{v}$. Using the fact that $\nu \geq \frac{\bar{v}^6 \underline{f}^4}{36}$ and $\xi = \frac{12\nu}{\underline{\beta}\bar{v}^3 \underline{f}^2}$, we get

$$\|\hat{\mathbf{w}} - \mathbf{u}\|_2 \leq \frac{\sqrt{\nu}}{\underline{\beta}}(1-\beta) = \frac{\nu}{\underline{\beta}\sqrt{\nu}}(1-\beta) \leq \frac{6\nu}{\underline{\beta}\bar{v}^3 \underline{f}^2}(1-\beta) \leq \frac{\xi}{2}(1-\beta) = \frac{\underline{u}_\beta}{2}\,. \tag{6.13}$$

Therefore, it follows that $\hat{w}_i \geq u_i - \|\hat{\mathbf{w}} - \mathbf{u}\|_2 \geq \underline{u}_\beta - \|\hat{\mathbf{w}} - \mathbf{u}\|_2 > 0$ for all i.

Step 2. Now we show that $\hat{\mathbf{w}} \in \mathsf{F}_\beta \mathcal{U}$. For the case that $\mathbf{u}$ is in the boundary region of $\mathsf{F}_\beta\mathcal{U}$, $\hat{\mathbf{w}} = \mathbf{w}^{\mathsf{L}}(\mathbf{v}|\mathbf{u}) = \mathbf{u}$, thus clearly $\hat{\mathbf{w}} \in \mathsf{F}_\beta\mathcal{U}$. Next, assume that $\mathbf{u}$ is in the central region $\hat{\mathcal{U}}_\beta$. Consider the signed distance function $\mathfrak{J}$ defined in (6.8). For simplicity, we use $\mathfrak{J}(\boldsymbol{x})$ to represent $\mathfrak{J}(\boldsymbol{x}, \mathsf{F}_\beta)$. Following Corollary 6.1, we have $\mathfrak{J}(\hat{\mathbf{w}}) \leq 0$ if and only if $\hat{\mathbf{w}} \in \mathsf{F}_\beta\mathcal{U}$ because $\hat{\mathbf{w}} \geq \mathbf{0}$. In order to show $\mathfrak{J}(\hat{\mathbf{w}}) \leq 0$, we consider the quadratic expansion around $\mathbf{u}$ as in Proposition 6.4 (with scalar $\mathsf{F} = \mathsf{F}_\beta$ in the proposition). We first prove that the necessary conditions for Proposition 6.4 to hold.

Lemma 6.6. *The future promise $\hat{\mathbf{w}}$ lies in $\mathcal{S}(\mathsf{F}_\beta)$ for any $\mathbf{u}$ in $\hat{\mathcal{U}}_\beta$.*

This lemma implies that the set $\mathcal{S}(\mathsf{F}_\beta)$ contains all $\boldsymbol{\Theta}$ between $\hat{\mathbf{w}}$ and $\mathbf{u}$ because it is convex. Following Proposition 6.2, we have that $\boldsymbol{\alpha}^*(\boldsymbol{\Theta}/\mathsf{F}_\beta) > \mathbf{0}$.

Lemma 6.7. *The matrix $\Pi(\boldsymbol{a})$ is positive definite for all $\boldsymbol{a} > \mathbf{0}$.*

Because both $\mathbf{u}$ and $\boldsymbol{\Theta}$ are in $\mathcal{S}(\mathsf{F}_\beta)$, Lemma 6.7 enables us to invoke Proposition 6.4. Therefore, we next analyze each item in the quadratic expansion at a time. First, $\mathbf{u} \in \mathsf{F}_\beta\mathcal{U}$ implies that $\mathfrak{J}(\mathbf{u}) \leq 0$. Thus, to show $\mathfrak{J}(\hat{\mathbf{w}}) \leq 0$, it is sufficient to show that

$$\nabla\mathfrak{J}(\mathbf{u})^{\intercal}(\hat{\mathbf{w}} - \mathbf{u}) + \frac{1}{2}(\hat{\mathbf{w}} - \mathbf{u})^{\intercal}\mathrm{Hess}(\mathfrak{J}(\boldsymbol{\Theta}))(\hat{\mathbf{w}} - \mathbf{u}) \leq 0\,. \tag{6.14}$$

In order to show this inequality holds, we find a lower bound for $-\nabla\mathfrak{J}(\mathbf{u})^{\intercal}(\hat{\mathbf{w}} - \mathbf{u})$, and an upper bound for $(\hat{\mathbf{w}} - \mathbf{u})^{\intercal}\mathrm{Hess}(\mathfrak{J}(\boldsymbol{\Theta}))(\hat{\mathbf{w}} - \mathbf{u})$. Then, we compare these bounds. We start by considering the first-order term. The following lemma provides the lower bound.

Lemma 6.8. *For any $\beta \geq \underline{\beta}$, the gradient of $\mathfrak{J}(\mathbf{u})$ satisfies*

$$\frac{(1-\beta)^2 \xi}{4\underline{\beta}\mathsf{F}_\beta^2} \leq -\nabla \mathfrak{J}\left(\mathbf{u}\right)^\intercal \left(\hat{\mathbf{w}} - \mathbf{u}\right).$$

This lemma mainly stems from the fact that $\nabla \mathfrak{J}(\mathbf{u}) = \boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)/\mathsf{F}_\beta$ (see Proposition 6.3) and the future promises lie in a plane with normal $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$ (see Proposition 3.4). Following this lemma, it is sufficient to show that $\frac{1}{2}(\hat{\mathbf{w}} - \mathbf{u})^\intercal \text{Hess}(\mathfrak{J}(\boldsymbol{\Theta}))(\hat{\mathbf{w}} - \mathbf{u}) \leq \frac{(1-\beta)^2 \xi}{4\underline{\beta}\mathsf{F}_\beta^2}$. Therefore, we next find an upper bound for the term with the Hessian in the following lemma.

Lemma 6.9. *For any $\beta \geq \underline{\beta}$, the Hessian of $\mathfrak{J}(\mathbf{u})$ satisfies*

$$\frac{1}{2}(\hat{\mathbf{w}} - \mathbf{u})^\intercal \mathit{Hess}(\mathfrak{J}(\boldsymbol{\Theta}))(\hat{\mathbf{w}} - \mathbf{u}) \leq \left(\frac{1-\beta}{\underline{\beta}}\right)^2 \frac{3\nu}{\underline{f}^2 \bar{v}^3 \mathsf{F}_\beta^2},$$

where $\boldsymbol{\Theta}$ is an arbitrary point between $\mathbf{u}$ and $\hat{\mathbf{w}}$.

In the proof of this lemma, following the min-max theorem, we first bound $(\hat{\mathbf{w}} - \mathbf{u})^\intercal \text{Hess}(\mathfrak{J}(\boldsymbol{\Theta}))(\hat{\mathbf{w}} - \mathbf{u})$ by the multiplication of the maximum eigenvalue of $\text{Hess}(\mathfrak{J}(\boldsymbol{\Theta}))$ and $\|\hat{\mathbf{w}} - \mathbf{u}\|_2^2$. Then, we replace the norm square, $\|\hat{\mathbf{w}} - \mathbf{u}\|_2^2$, with its bound provided in Lemma 6.5.

Combining the bounds provided by Lemmas 6.8 and 6.9, we obtain that (6.14) holds because $\frac{12\nu}{\underline{\beta f}^2 \bar{v}^3} \leq \xi$. Therefore, we prove $\mathfrak{J}(\hat{\mathbf{w}}) \leq 0$ implying that $\hat{\mathbf{w}} \in \mathsf{F}_\beta \mathcal{U}$ for $\mathbf{u}$ such that $u_i \geq \underline{u}_\beta$ for all i. □

6.4.4 Proof of Theorem 3.1

Proof. This result follows from Proposition 3.6 and Proposition 3.1. The set $\mathsf{F}_\beta \mathcal{U}$ characterized by the constant in (6.11) is a subset of $\mathcal{U}_\beta$ by Proposition 3.1. Thus, the maximum achievable social welfare $J_\beta = \mathsf{F}_\beta J^{\text{FB}}$, is less than J_β^*. □

6.4.5 Proofs of Lemmas 6.5, 6.6, 6.7, 6.8 and 6.9

Proof of Lemma 6.5. Proposition 3.4 states that future promises lie on a line for different $\mathbf{v}$'s when $n = 2$. The extreme points of this line correspond to $(0, \bar{v})$ and $(\bar{v}, 0)$ because $\hat{w}_i(\mathbf{v}|\mathbf{u})$ is non-increasing in v_i and non-decreasing in v_j, according to Proposition 3.9. Therefore, it follows that

$$\|\hat{\mathbf{w}} - \mathbf{u}\|_2^2 \leq \max\left(\|\hat{\mathbf{w}}((\bar{v}, 0)|\mathbf{u}) - \mathbf{u}\|_2^2, \|\hat{\mathbf{w}}((0, \bar{v})|\mathbf{u}) - \mathbf{u}\|_2^2\right).$$

Let $\boldsymbol{\gamma} = \boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)$ and suppose that $\gamma_1 \leq \gamma_2$ without loss of generality.

Step 1 (Bounding $\|\hat{\mathbf{w}}((\bar{v}, 0)|\mathbf{u}) - \mathbf{u}\|_2^2$). In this step, using (3.10) and (3.11), we bound the terms $(\hat{w}_1((\bar{v}, 0)|\mathbf{u}) - u_1)^2$ and $(\hat{w}_2((\bar{v}, 0)|\mathbf{u}) - u_2)^2$ separately. First, we obtain more compact expressions for these terms and using these expressions, we get the bounds.

$$\begin{aligned}\hat{w}_1((\bar{v}, 0)|\mathbf{u}) =& \frac{u_1}{\beta} + \frac{1-\beta}{\beta}\Bigg[\int_0^{\bar{v}} F_1(y) F_2\left(\frac{\gamma_1 y}{\gamma_2}\right) \mathrm{d}y - F_2\left(\frac{\gamma_1 \bar{v}}{\gamma_2}\right)\bar{v} \\ &+ \frac{\gamma_2}{\gamma_1}\left\{\mathbb{E}_{v_2}\left[\int_0^{v_2} F_1\left(\frac{\gamma_2 y}{\gamma_1}\right)\mathrm{d}y - F_1\left(\frac{\gamma_2 v_2}{\gamma_1}\right) v_2\right]\right\}\Bigg]\end{aligned}$$

Changing the order of integrations, we rewrite $\hat{w}_1((\bar{v}, 0)|\mathbf{u})$ as follows:

$$\hat{w}_1((\bar{v}, 0)|\mathbf{u}) = \frac{u_1}{\beta} - \frac{1-\beta}{\beta}\left[\frac{\gamma_1}{\gamma_2}\int_0^{\bar{v}} f_2\left(\frac{\gamma_1 y}{\gamma_2}\right) F_1(y) y \mathrm{d}y + \mathbb{E}[v_1]\right].$$

We are now ready to bound $(\hat{w}_1((\bar{v}, 0)|\mathbf{u}) - u_1)^2$.

$$\begin{aligned}\left(\frac{\beta}{1-\beta}\right)^2 (\hat{w}_1((\bar{v}, 0)|\mathbf{u}) - u_1)^2 &= \left[u_1 - \frac{\gamma_1}{\gamma_2}\int_0^{\bar{v}} f_2\left(\frac{\gamma_1 y}{\gamma_2}\right) F_1(y) y \mathrm{d}y - \mathbb{E}[v_1]\right]^2 \\ &\leq \underbrace{\left[\frac{\bar{f}\bar{v}^2}{2} - \frac{\bar{f}}{2}\underline{m}_2 + \overline{m}\right]^2}_{\nu_1}.\end{aligned}$$

This inequality follows from the fact that (i) γ_1/γ_2 is less than 1, (ii) the p.d.f. is less than $\bar{f}$ in its support, and (iii) we can remove u_1 because $\mathbb{E}[v_1] \geq u_1$ and they have different signs.

Next, we derive $\hat{w}_2((\bar{v},0)|\mathbf{u})$ following similar steps and bound $(\hat{w}_2((\bar{v},0)|\mathbf{u})-u_1)^2$.

$$
\begin{aligned}
(\hat{w}_2((\bar{v},0)|\mathbf{u})-u_2)^2 &= \left(\frac{1-\beta}{\beta}\right)^2\left[u_2+\frac{\gamma_1^2}{\gamma_2^2}\int_0^{\bar{v}} f_2\left(\frac{\gamma_1 y}{\gamma_2}\right) yF_1(y)\mathrm{d}y - \int_0^{\bar{v}} \bar{F}_2(y)F_1\left(\frac{\gamma_2}{\gamma_1}y\right)\mathrm{d}y\right]^2 \\
&\leq \left(\frac{1-\beta}{\beta}\right)^2\left[\max\left(u_2+\frac{\gamma_1^2}{\gamma_2^2}\int_0^{\bar{v}} f_2\left(\frac{\gamma_1 y}{\gamma_2}\right) yF_1(y)\mathrm{d}y, \int_0^{\bar{v}} \bar{F}_2(y)F_1\left(\frac{\gamma_2}{\gamma_1}y\right)\mathrm{d}y\right)\right]^2 \\
&\leq \left(\frac{1-\beta}{\beta}\right)^2\left[\max\left(\mathbb{E}[v_2]+\frac{\bar{f}\bar{v}^2}{2}-\frac{\bar{f}}{2}\mathbb{E}[v_1^2], \mathbb{E}[v_2]\right)\right]^2 \\
&\leq \left(\frac{1-\beta}{\beta}\right)^2\underbrace{\left[\overline{m}+\frac{\bar{f}\bar{v}^2}{2}-\frac{\bar{f}\underline{m}_2}{2}\right]^2}_{\nu_2}.
\end{aligned}
$$

Here, we first group the terms with same sign inside the square brackets and bound them separately. We use the fact that $\gamma_1/\gamma_2 \leq 1$ and the p.d.f.'s are less than $\bar{f}$ in its support to obtain the second inequality.

By using the bounds on $\hat{w}_i((\bar{v},0)|\mathbf{u})$ for $i=1,2$, we obtain $\|\hat{\mathbf{w}}((\bar{v},0)|\mathbf{u})-\mathbf{u}\|_2^2 \leq \left(\frac{1-\beta}{\beta}\right)^2(\nu_1+\nu_2)$.

Step 2 (Bounding $\|\hat{\mathbf{w}}((0,\bar{v})|\mathbf{u})-\mathbf{u}\|_2^2$). In this step, using (3.10) and (3.11), we bound the terms $(\hat{w}_1((0,\bar{v})|\mathbf{u})-u_1)^2$ and $(\hat{w}_2((0,\bar{v},)|\mathbf{u})-u_2)^2$ separately, following similar steps as above.

$$
\begin{aligned}
\hat{w}_1((0,\bar{v})|\mathbf{u}) =& \frac{u_1}{\beta}-\frac{1-\beta}{\beta}\Bigg\{\int_0^{\bar{v}}\bar{F}_1(y)F_2\left(\frac{\gamma_1 y}{\gamma_2}\right)\mathrm{d}y+\frac{\gamma_2}{\gamma_1}\Bigg[\int_0^{\bar{v}}F_1\left(\frac{\gamma_2 y}{\gamma_1}\right)\mathrm{d}y - F_1\left(\frac{\gamma_2\bar{v}}{\gamma_1}\right)\bar{v}- \\
&\mathbb{E}_{v_2}\left[\int_0^{v_2}F_1\left(\frac{\gamma_2 y}{\gamma_1}\right)\mathrm{d}y\right]+\mathbb{E}_{v_2}\left[F_2\left(\frac{\gamma_2 v_2}{\gamma_1}\right)v_2\right]\Bigg]\Bigg\} \\
=& \frac{u_1}{\beta}+\frac{1-\beta}{\beta}\frac{\gamma_1}{\gamma_2}\int_0^{\bar{v}}\bar{F}_1(y)yf_2\left(\frac{\gamma_1 y}{\gamma_2}\right)\mathrm{d}y\,.
\end{aligned}
$$

Because $u_1 \leq \mathbb{E}[v_1]$, $\gamma_1/\gamma_2 \leq 1$ and the p.d.f. is bounded by $\bar{f}$ in its support, we obtain the following inequality:

$$
(\hat{w}_1((0,\bar{v})|\mathbf{u})-u_1)^2 = \left(\frac{1-\beta}{\beta}\right)^2\left[u_1+\frac{\gamma_1}{\gamma_2}\int_0^{\bar{v}}\bar{F}_1(y)yf_2\left(\frac{\gamma_1 y}{\gamma_2}\right)\mathrm{d}y\right]^2 \leq \left(\frac{1-\beta}{\beta}\right)^2\underbrace{\left[\overline{m}+\frac{\bar{f}}{2}\overline{m}_2\right]^2}_{\nu_3}.
$$

Next, we derive a more compact expression for $\hat{w}_2((0,\bar{v})|\mathbf{u})$.

$$\begin{aligned}\hat{w}_2((0,\bar{v})|\mathbf{u}) =& \frac{u_2}{\beta} + \frac{1-\beta}{\beta}\left[\int_0^{\bar{v}} F_1\left(\frac{\gamma_2 y}{\gamma_1}\right)\mathrm{d}y - F_1\left(\frac{\gamma_2\bar{v}}{\gamma_1}\right)\bar{v} - \int_0^{\bar{v}} \bar{F}_2(y)F_1\left(\frac{\gamma_2 y}{\gamma_1}\right)\mathrm{d}y\right] \\ & - \frac{1-\beta}{\beta}\frac{\gamma_1}{\gamma_2}\left[\frac{\gamma_1}{\gamma_2}\int_0^{\bar{v}} \bar{F}_1(y) f_2\left(\frac{\gamma_1 y}{\gamma_2}\right) y\mathrm{d}y\right] \\ =& \frac{u_2}{\beta} + \frac{1-\beta}{\beta}\left[\int_0^{\bar{v}} F_2(y)F_1\left(\frac{\gamma_2 y}{\gamma_1}\right)\mathrm{d}y - F_1\left(\frac{\gamma_2\bar{v}}{\gamma_1}\right)\bar{v} - \int_0^{\bar{v}} \bar{F}_1\left(\frac{\gamma_2 y}{\gamma_1}\right) f_2(y)\, y\mathrm{d}y\right].\end{aligned}$$

Using a similar approach to previous steps, we now bound $(\hat{w}_2((0,\bar{v})|\mathbf{u}) - u_2)^2$.

$$\begin{aligned}&(\hat{w}_2((0,\bar{v})|\mathbf{u}) - u_2)^2 \\ \le& \left(\frac{1-\beta}{\beta}\right)^2\left[\max\left(u_2 + \int_0^{\bar{v}} F_2(y)F_1\left(\frac{\gamma_2 y}{\gamma_1}\right)\mathrm{d}y, \int_0^{\bar{v}} y f_2(y)\bar{F}_1\left(\frac{\gamma_2 y}{\gamma_1}\right)\mathrm{d}y + F_1\left(\frac{\gamma_2\bar{v}}{\gamma_1}\right)\bar{v}\right)\right]^2 \\ \le& \left(\frac{1-\beta}{\beta}\right)^2 \underbrace{(\overline{m}+\bar{v})^2}_{\nu_4}.\end{aligned}$$

It follows that $\|\hat{\mathbf{w}}((0,\bar{v})|\mathbf{u}) - \mathbf{u}\|_2^2 \le \left(\frac{1-\beta}{\beta}\right)^2(\nu_3+\nu_4)$.

Combining the bounds on $\|\hat{\mathbf{w}}((\bar{v},0)|\mathbf{u}) - \mathbf{u}\|_2^2$ and $\|\hat{\mathbf{w}}((0,\bar{v})|\mathbf{u}) - \mathbf{u}\|_2^2$, we obtain the following bound on the largest jump, i.e., $\|\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}) - \mathbf{u}\|_2^2 \le \left(\frac{1-\beta}{\beta}\right)^2 \max(\nu_1+\nu_2, \nu_3+\nu_4) \le \left(\frac{1-\beta}{\underline{\beta}}\right)^2 \nu$.

□

Proof of Lemma 6.6. First, when $n=2$, it is obvious that

$$\mathcal{S}(\mathsf{F}_\beta) = \{\boldsymbol{z} : |z_1 - z_2| < \mathbf{1}\{z_1 \ge z_2\}\mathsf{F}_\beta\mathbb{E}[v_1] + \mathbf{1}\{z_2 \ge z_1\}\mathsf{F}_\beta\mathbb{E}[v_2], \boldsymbol{z} > 0\}.$$

Without loss of generality, assume $\hat{w}_1 \ge \hat{w}_2$. Because $\mathbf{u} \in \hat{\mathcal{U}}_\beta$, we know that $u_1 < \mathsf{F}_\beta\mathbb{E}[v_1]$ and $u_2 \ge \underline{u}_\beta$, which further implies

$$\hat{w}_1 - \mathsf{F}_\beta\mathbb{E}[v_1] \le u_1 - \mathsf{F}_\beta\mathbb{E}[v_1] + \|\hat{\mathbf{w}} - \mathbf{u}\|_2 < \|\hat{\mathbf{w}} - \mathbf{u}\|_2, \text{ and } \hat{w}_2 - \underline{u}_\beta \ge u_2 - \underline{u}_\beta - \|\hat{\mathbf{w}} - \mathbf{u}\|_2 \ge -\|\hat{\mathbf{w}} - \mathbf{u}\|_2.$$

Combining these two inequalities, we obtain $\hat{w}_1 - \hat{w}_2 - \mathsf{F}_\beta\mathbb{E}[v_1] + \underline{u}_\beta < 2\|\hat{\mathbf{w}} - \mathbf{u}\|_2$. Equation (6.13) further implies that $\underline{u}_\beta \ge 2\|\hat{\mathbf{w}} - \mathbf{u}\|_2$, and, therefore, $\hat{w}_1 - \hat{w}_2 < \mathsf{F}_\beta\mathbb{E}[v_1]$. □

Proof of Lemma 6.7. The Hessian matrix of $\phi(\boldsymbol{a})$ is provided in Proposition 3.3 for $\boldsymbol{a} > \mathbf{0}$. Denote

$z \triangleq \int_0^{\bar{v}\min(a_1,a_2)} x^2 f_1\left(\frac{x}{a_1}\right) f_2\left(\frac{x}{a_2}\right) \mathrm{d}x$. Then, $\Pi(\boldsymbol{a})$ is given as follows:

$$\Pi(\boldsymbol{a}) = \begin{bmatrix} \frac{z}{a_1^3 a_2} + 1 & -\frac{z}{a_1^2 a_2^2} + 1 \\ -\frac{z}{a_1^2 a_2^2} + 1 & \frac{z}{a_1 a_2^3} + 1 \end{bmatrix} .$$

Following Sylvester's criterion, a matrix is positive definite if its leading principal minors are all positive. Because $z/a_1^3 a_2 + 1 > 0$, it is sufficient to show $\det(\Pi(\boldsymbol{a})) > 0$. Thus, we next derive this determinant and show that it is positive. Because the p.d.f. $f_i(\cdot)$ is lower bounded in the support $z \geq \underline{f}^2\Big(\bar{v}\min(a_1,a_2)\Big)^3/3$. Thus, we have that

$$\det(\Pi(\boldsymbol{a})) = \frac{z}{a_1^3 a_2^3} \geq \frac{\underline{f}^2\Big(\bar{v}\min(a_1,a_2)\Big)^3}{3a_1^3 a_2^3} = \frac{\bar{v}^3 \underline{f}^2}{3\max(a_1^3, a_2^3)} > 0 . \qquad \square$$

Proof of Lemma 6.8. Lemma 6.7 implies we can invoke Proposition 6.3 to show that the gradient of $\mathfrak{J}$ is well defined and $\nabla\mathfrak{J}(\mathbf{u}) = \boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)/\mathsf{F}_\beta$ as given in (6.10). Moreover, using Proposition 3.4, we can alternatively express $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)^\intercal(\mathbf{u} - \hat{\mathbf{w}})$ as follows:

$$-\nabla\mathfrak{J}(\mathbf{u})^\intercal(\hat{\mathbf{w}} - \mathbf{u}) = \frac{\left(\boldsymbol{\alpha}^*\left(\frac{\mathbf{u}}{\mathsf{F}_\beta}\right)\right)^\intercal(\mathbf{u} - \hat{\mathbf{w}})}{\mathsf{F}_\beta} = \frac{\left(\boldsymbol{\alpha}^*\left(\frac{\mathbf{u}}{\mathsf{F}_\beta}\right)\right)^\intercal\left(\nabla\phi\left(\boldsymbol{\alpha}^*\left(\frac{\mathbf{u}}{\mathsf{F}_\beta}\right)\right) - \mathbf{u}\right)}{\mathsf{F}_\beta}\left(\frac{1-\beta}{\beta}\right) .$$

The first-order conditions for the optimization problem $\sup_{\|\boldsymbol{x}\|_1 = 1, \boldsymbol{x} \geq \mathbf{0}} \left\{\frac{\boldsymbol{x}^\intercal\mathbf{u}}{\mathsf{F}_\beta} - \phi(\boldsymbol{x})\right\}$ enable us to express $\mathbf{u}$ as $\mathbf{u} = \mathsf{F}_\beta\nabla\phi\left(\boldsymbol{\alpha}^*\left(\mathbf{u}/\mathsf{F}_\beta\right)\right) - \mu\boldsymbol{e}$, where μ is the dual variable corresponding to constraint $\|\boldsymbol{x}\|_1 = 1$ and $\boldsymbol{e}$ is the 2 dimensional vector of ones. Therefore, we get the following equation.

$$-\nabla\mathfrak{J}(\mathbf{u})^\intercal(\mathbf{w} - \mathbf{u}) = \frac{\left(\boldsymbol{\alpha}^*\left(\frac{\mathbf{u}}{\mathsf{F}_\beta}\right)\right)^\intercal((1-\mathsf{F}_\beta)\nabla\phi\left(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta)\right) + \mu\boldsymbol{e})}{\mathsf{F}_\beta}\left(\frac{1-\beta}{\beta}\right) .$$

Note that $\boldsymbol{x}^\intercal\nabla\phi(\boldsymbol{x}) = \phi(\boldsymbol{x})$, and the dual variable μ is nonnegative because $\mathbf{u} \in \hat{\mathcal{U}}_\beta$ and the term $\mathsf{F}_\beta\nabla\phi(\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta))$ corresponds to a state in the efficient frontier of $\hat{\mathcal{U}}_\beta$. Moreover $\mathsf{F}_\beta \geq 1/2$ and $\phi(\boldsymbol{x}) \geq \underline{m}/2$ for all $\boldsymbol{x} \in \mathbb{R}_+^2$ by definition. Thus, using these observations we obtain

$$-\nabla\mathfrak{J}(\mathbf{u})^\intercal(\hat{\mathbf{w}} - \mathbf{u}) \geq \frac{(1-\mathsf{F}_\beta)}{\mathsf{F}_\beta}\frac{(1-\beta)}{\beta}\phi\left(\boldsymbol{\alpha}^*\left(\frac{\mathbf{u}}{\mathsf{F}_\beta}\right)\right) \geq \frac{(1-\beta)(1-\mathsf{F}_\beta)\underline{m}}{4\underline{\beta}\mathsf{F}_\beta^2} = \frac{(1-\beta)^2\xi}{4\underline{\beta}\mathsf{F}_\beta^2} . \qquad \square$$

Proof of Lemma 6.9. Lemma 6.7 enables us to use Proposition 6.3 to obtain that

$$\text{Hess}\left(\mathfrak{J}(\boldsymbol{\Theta})\right) = \frac{1}{\mathsf{F}_\beta^2}\Pi(\boldsymbol{\alpha}^*(\boldsymbol{\Theta}/\mathsf{F}_\beta))^{-1}\left(I - \frac{\boldsymbol{e}\boldsymbol{e}^\intercal\Pi(\boldsymbol{\alpha}^*(\boldsymbol{\Theta}/\mathsf{F}_\beta))^{-1}}{\boldsymbol{e}^\intercal\Pi(\boldsymbol{\alpha}^*(\boldsymbol{\Theta}/\mathsf{F}_\beta))^{-1}\boldsymbol{e}}\right).$$

Denote $\boldsymbol{\lambda} = \boldsymbol{\alpha}^*(\boldsymbol{\Theta}/\mathsf{F}_\beta)$ and $z \triangleq \int_0^{\bar{v}\min(\lambda_1,\lambda_2)} x^2 f_1\left(\frac{x}{\lambda_1}\right) f_2\left(\frac{x}{\lambda_2}\right) \mathrm{d}x$, then $\Pi(\boldsymbol{\lambda}) = \text{Hess}(\phi(\boldsymbol{\lambda})) + \boldsymbol{e}\boldsymbol{e}^\intercal$ by definition. Therefore, we can evaluate the matrix multiplication operations using the definitions and we obtain

$$\text{Hess}(\mathfrak{J}(\boldsymbol{\Theta})) = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \frac{1}{\mathsf{F}_\beta^2}\frac{\lambda_1^3\lambda_2^3}{z}.$$

The largest eigenvalue of matrix $\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$ is 2. Because the p.d.f.'s $f_i(\cdot)$ are lower bounded in the support, we have that $z \geq \underline{f}^2\left(\bar{v}\min(\lambda_1,\lambda_2)\right)^3/3$. These observations and the fact that $\max(\lambda_1,\lambda_2) \leq 1$ imply that the largest eigenvalue of $\text{Hess}(\mathfrak{J}(\boldsymbol{\Theta}))$ is bounded by $\frac{6}{\underline{f}^2\bar{v}^3\mathsf{F}_\beta^2}$. Finally, using Lemma 6.5, we can also bound $\|\hat{\mathbf{w}} - \mathbf{u}\|_2^2$ and obtain the following inequality:

$$\frac{1}{2}(\hat{\mathbf{w}} - \mathbf{u})^\intercal\text{Hess}(\mathfrak{J}(\boldsymbol{\Theta}))(\hat{\mathbf{w}} - \mathbf{u}) \leq \frac{3}{\underline{f}^2\bar{v}^3\mathsf{F}_\beta^2}\|\hat{\mathbf{w}} - \mathbf{u}\|_2^2 \leq \left(\frac{1-\beta}{\underline{\beta}}\right)^2\frac{3\nu}{\underline{f}^2\bar{v}^3\mathsf{F}_\beta^2}.$$

□

6.4.6 Proof of Theorem 3.2

Proof. In this proof, we consider a two-agent setting where only the values of the first agent are private. First, we provide the optimization problem which maximizes the total discounted expected social welfare for this setting. We next bound the optimal value for this problem using the linear programming approach to approximate dynamic program. We conclude by using the envelope theorem to provide a bound on the optimal rate of convergence when the discount rate is close to one.

Step 1. Consider a relaxation where the valuations of the second agent are publicly observable for all periods. Because the values of the second agent are public, the feasibility constraint (FA) should bind since the principal should allocate to the second agent whenever the resource is not allocated to the first. As a result, the principal can internalize the value of the second agent as his cost and we can consider an equivalent, simpler single-agent problem allocation problem of a costly resource.

In particular, we assume that there is a single agent with value v and the principal has a cost c for allocating the resource. The principal's cost is private from the perspective of the agent. First best is equal to $J^{\text{FB}} = \mathbb{E}_{v,c}[\max(v-c,0)]$ because it is socially optimal to allocate whenever the agent's value is above the principal's cost. We denote by J^*_β the optimal objective value of the principal. An optimal mechanism can be recursively characterized using dynamic programming. Let $J_\beta(u)$ be the optimal expected payoff-to-go of the principal when the promised utility of the agent is $u \in [0, \mathbb{E}[v]]$. We have

$$\begin{aligned} J_\beta(u) = \max_{p,w} \ & \mathbb{E}[(1-\beta)(v-c)p(v,c) + \beta J_\beta(w(v,c))] \\ \text{st. } & u_i = \mathbb{E}[(1-\beta)vp(v,c) + \beta w(v,c)]\,, && \text{(PK(u))} \\ & (1-\beta)vP(v) + \beta W(v) \geq (1-\beta)vP(v') + \beta W(v'), \quad v, v'\,, && \text{(IC)} \\ & 0 \leq w(v,c) \leq \mathbb{E}[v], \quad \forall v, c\,, && \text{(BC)} \\ & 0 \leq p(v,c) \leq 1 \quad \forall v, c\,, && \text{(FA)} \\ & P(v) = \mathbb{E}_c[p(v,c)] \text{ and } W(v) = \mathbb{E}_c[w(v,c)], \quad \forall v\,. && \text{(INT)} \end{aligned}$$

Given an optimal value function J_β, the optimal objective value is given by $J^*_\beta = \max_{u\in[0,\mathbb{E}[v]]} J_\beta(u)$.

Step 2. In this step, we find a family of upper bounds parametrized by b for the optimal value function $J^*_\beta(\cdot)$ using the linear programming approach to approximate dynamic programming (see De Farias and Van Roy, 2003). Define the set of feasible mechanisms $\mathcal{M}(u)$, i.e., $(p,w) \in \mathcal{M}(u)$ if and only if (p,w) satisfy (PK(u)), (IC), (BC), (FA), (INT) constraints. We impose a quadratic parametric form for the value functions given by $\hat{J}_\beta(u) = a - \frac{b}{2}(u-u^*)^2$ where $u^* = \mathbb{E}_{v,c}[v\mathbf{1}\{v \geq c\}]$ and $b < \max(1/(\bar{f}\mathbb{E}[v^2]), 1/u^*)$. The value of b is fixed, while we optimize over the value of a. This leads to the upper bound:

$$\begin{aligned} J_\beta(\tilde{u}) \leq \min_a \ & \hat{J}_\beta(\tilde{u}) \\ \text{st. } & \hat{J}_\beta(u) \geq \mathbb{E}_{v,c}[(1-\beta)(v-c)p(v,c) + \beta \hat{J}_\beta(w(v,c))], \quad \forall u \in [0,\mathbb{E}[v]], (p,w) \in \mathcal{M}(u)\,, \\ & \hat{J}_\beta(u) = a - \frac{b}{2}(u-u^*)^2, \quad \forall u \in [0, \mathbb{E}[v]]\,. \end{aligned}$$

Using $\hat{J}_\beta(u) = a - \frac{b}{2}(u-u^*)^2$ and rearranging the inequality constraint, we obtain the following:

$$J_\beta(\tilde{u}) \le \min_a a - \frac{b}{2}(\tilde{u}-u^*)^2$$

$$\text{st. } a \ge \max_{\substack{u\in[0,\mathbb{E}[v]],\\ (p,w)\in\mathcal{M}(u)}} \mathbb{E}_{v,c}\left[(v-c)p(v,c) - \frac{\beta}{(1-\beta)}\left(\frac{b}{2}\left(w(v,c)-u^*\right)^2\right)\right] + \frac{1}{1-\beta}\frac{b}{2}(u-u^*)^2\,.$$

Note that we focus on the upper bound for J_β^*, therefore when we take the supremum over $\tilde{u}$ at both sides of the inequality we obtain that $\tilde{u} = u^*$ at the right-hand side. Additionally, at the optimal a the inequality should be binding. Therefore, we obtain that:

$$J_\beta^* \le \max_{\substack{u\in[0,\mathbb{E}[v]],\\ (p,w)\in\mathcal{M}(u)}} \mathbb{E}_{v,c}\left[(v-c)p(v,c) - \frac{\beta}{(1-\beta)}\left(\frac{b}{2}\left(w(v,c)-u^*\right)^2\right)\right] + \frac{1}{1-\beta}\frac{b}{2}(u-u^*)^2\,.$$

We next describe how to eliminate the promise function $w(v,c)$ from the optimization problem. Using Jensen's inequality on c we lower bound the quadratic term in the objective involving $w(v,c)$ as follows:

$$\mathbb{E}_{v,c}\left[\left(w(v,c)-u^*\right)^2\right] \ge \mathbb{E}_v\left[\left(W(v)-u^*\right)^2\right]\,,$$

where we used the interim representation $W(v) = \mathbb{E}_c[w(v,c)]$. The envelope formula and the promise keeping constraint implies that we can equivalently write the incentive compatibility constraint as

$$\beta W(v) = u + (1-\beta)\left(\int_0^v P(y)\mathrm{d}y - vP(v) - \mathbb{E}_{\tilde{v}}\left[\int_0^{\tilde{v}} P(y)\mathrm{d}y\right]\right)\,, \tag{6.15}$$

together with $P(v)$ being non-decreasing in v, which we refer to as the (ND) constraint. Relaxing the (BC) constraint on the promised function we obtain an upper bound in which we can eliminate the promised function from the optimization program. This leads to:

$$J_\beta^* \le \max_{\substack{u\in[0,\mathbb{E}[v]],\\ p\in\mathcal{P}_{\mathrm{ND}}}} \mathbb{E}_{v,c}\left[(v-c)p(v,c) - \frac{\beta}{(1-\beta)}\left(\frac{b}{2}\left(W(v)-u^*\right)^2\right)\right] + \frac{1}{1-\beta}\frac{b}{2}(u-u^*)^2\,.$$

where $W(v)$ as given in (6.15) and $\mathcal{P}_{\mathrm{ND}}$ is the set of allocations satisfying (ND) and (FA) constraints

Step 3. In this step, we focus on the optimization problem over u, and show that for any given p the optimal u is such that $u = \mathbb{E}_{v,c}[vp(v,c)]$. The objective is differentiable over u with derivative equal to

$$\frac{b}{\beta}\left(\mathbb{E}_{v,c}[vp(v,c)] - u\right) .$$

Therefore, because $b > 0$ the objective function is concave in u for any given p, and the first-order conditions are sufficient for optimality. This yields that $u = \mathbb{E}_{v,c}[vp(v,c)]$ satisfies the first-order conditions for any given p. We conclude by checking the boundary conditions. Specifically, the derivative with respect to u evaluated at $u = 0$ is $b\mathbb{E}_{v,c}[vp(v,c)]/\beta \geq 0$ and at $u = \mathbb{E}[v]$ is $b\mathbb{E}_{v,c}[(p(v,c)-1)v]/\beta \leq 0$.

Using this result, we hereafter focus on the following optimization problem.

$$V(\beta) \triangleq \max_{p \in \mathcal{P}_{\mathrm{ND}}} G(p;\beta)\,,$$

where

$$G(p;\beta) \triangleq \mathbb{E}_{v,c}\left[(v-c)p(v,c)\right] + \frac{b}{2}\left(\mathbb{E}_{v,c}[vp(v,c)] - u^*\right)^2 - \frac{b}{2}\left(\frac{1-\beta}{\beta}\right)\mathrm{Var}_v\left[vP(v) - \int_0^v P(y)\mathrm{d}y\right] .$$

Step 4. In this step, we find the optimal allocation function p that maximizes $G(p;1)$. Observe that the function $G(p;1)$ is not a concave function of p. Therefore, we first show that any stationary p is in a threshold form with a weight on v, i.e., $p_r(v,c) = \mathbf{1}\{vr \geq c\}$ for some $r \geq 0$. We consider the Gateaux derivative of $G(\cdot,1)$ at an arbitrary p for some direction h which is given as follows:

$$\frac{\mathrm{d}}{\mathrm{d}\delta}G(p+\delta h;1)\Big|_{\delta=0} = \mathbb{E}_{v,c}[h(v,c)g(v,c;p)]\,,$$

where

$$g(v,c;p) = v(1 + b(\mathbb{E}_{\tilde{v},\tilde{c}}[\tilde{v}p(\tilde{v},\tilde{c})] - u^*)) - c\,.$$

For p to be a local extrema, it should satisfy the following conditions. When $g(v,c;p) > 0$ the allocation should satisfy $p(v,c) = 1$, when $g(v,c;p) < 0$ the allocation should satisfy $p(v,c) = 0$, and when $g(v,c;p) = 0$ the allocation can take any value $p(v,c) \in [0,1]$. Note that for $b < 1/u^*$, $g(v,c;p)$ is strictly increasing in v and strictly decreasing in c. Moreover, $g(v,c;p)$ is a linear function of v and c and $g(0,0;p) = 0$. Therefore, the set of points $\{(v,c) : g(v,c;p) = 0\}$ is a line going through the origin $(0,0)$. Let $vr = c$ for some $r \geq 0$ be such line. Because g is monotone, we have that $g(v,c;p) > 0$ for $vr > c$ and $g(v,c;p) < 0$ for $vr < c$. These observations imply that $p_r(v,c) = \mathbf{1}\{vr \geq c\}$ is a local extrema.

We can solve for the value of the optimal r by using the fact that $g(v, vr; p_r) = 0$ for all v. Clearly, $r = 1$ is a solution to this equation. We next show that $r = 1$ is the unique solution because $g(v, vr; p_r)$ is strictly decreasing in r. Using Leibniz rule we obtain that

$$\frac{\mathrm{d}g(v, vr; p_r)}{\mathrm{d}r} = v\left(b\int_0^{\min(\bar{v}, \bar{v}/r)} x^2 f(xr) f(x)\mathrm{d}x - 1\right) < 0\,,$$

where the inequality follows from the fact that $\int_0^{\min(\bar{v},\bar{v}/r)} x^2 f(xr)f(x)\mathrm{d}x \le \bar{f}\int_0^{\bar{v}} x^2 f(x)\mathrm{d}x = \bar{f}\mathbb{E}[v^2]$ and $b < 1/(\bar{f}\mathbb{E}[v^2])$. Thus, only $r = 1$ satisfies that $g(v, vr; p_r) = 0$ and hence it is the unique solution. Letting $r = 0$ and $r \to \infty$ we get that the boundary solutions $p(v, c) = 1$ and $p(v, c) = 0$ are dominated. Therefore, it follows that $V(1) = J^{\text{FB}}$ with the essentially unique optimal solution equal to $p^*(v, c) = \mathbf{1}\{v \ge c\}$ (i.e., solutions can differ from p^* in a set of measure zero).

Step 5. Note that the upper bound for J^*_β is given by $V(\beta)$. The next lemma shows that the required conditions to invoke the envelope theorem stated in Corollary 4 from Milgrom and Segal (2002) are satisfied. Because the approach is topological in nature, we endow the space of stage mechanism with a topology. See the proof of Proposition 6.5 for details.

Lemma 6.10. *The following statements hold:*

1. *$\mathcal{P}_{\text{ND}}$ is a nonempty weak-* compact set.*
2. *$G(p, \beta)$ is weak-* continuous in p for $p \in \mathcal{P}_{\text{ND}}$.*
3. *$dG(p, \beta)/d\beta$ is jointly weak-* continuous in β and p for $\beta > 0$.*

Therefore, it follows that $V(\beta)$ is left differentiable at $\beta = 1$ with left derivative equal to

$$V'(1_-) = \min_{\tilde{p}\in\mathcal{P}^*_{\text{ND}}(1)} \frac{\mathrm{d}G(p, \beta)}{\mathrm{d}\beta}\Big|_{p=\tilde{p},\beta=1}\,,$$

where $\mathcal{P}^*_{\text{ND}}(\beta) \triangleq \{p \in \mathcal{P}_{\text{ND}} : V(\beta) = G(p; \beta)\}$. When $\beta = 1$, the optimal allocation is essentially unique, i.e., $\mathcal{P}^*_{\text{ND}}(1)$ is essentially singleton. Thus, using that $P(v) = F(v)$ when $p(v, c) = \mathbf{1}\{v \ge c\}$ it follows that

$$V'(1_-) = \frac{\mathrm{d}G(p; \beta)}{\mathrm{d}\beta}\Big|_{p(v,c)=\mathbf{1}\{v\ge c\},\beta=1} = \frac{b}{2}\mathrm{Var}_v\left[vF(v) - \int_0^v F(y)\mathrm{d}y\right] \ge \eta > 0\,.$$

There exists a positive constant η because $b > 0$ and the variance is positive. We prove latter by contradiction. Because the density is positive, the variance is zero only if the term inside the variance is zero everywhere, i.e., $vF(v) - \int_0^v F(y)\mathrm{d}y = 0$ for all v. Taking derivatives we obtain that this condition is satisfied only when the density is zero. A contradiction.

For β values around 1, we can use the differentiability of $V(\beta)$ to obtain

$$V(\beta) = V(1) - (1-\beta)V'(1_-) + o(1-\beta) \leq J^{\text{FB}} - (1-\beta)\eta + o(1-\beta)\,.$$

Noting that $V(\beta) = \max_{p \in \mathcal{P}_{\text{ND}}} G(p;\beta) \geq J^*_\beta$, we obtain the following and conclude the proof:

$$\frac{J^{\text{FB}} - J^*_\beta}{J^{\text{FB}}} \geq \frac{J^{\text{FB}} - V(\beta)}{J^{\text{FB}}} \geq \frac{(1-\beta)\eta}{J^{\text{FB}}} - o(1-\beta)\,. \qquad \square$$

6.4.7 Proof of Lemmas 6.10 and 6.11

Proof of Lemma 6.10. We prove each part at a time.

1. We assume that the allocation function p lies in the space $L^\infty \triangleq L^\infty([0,\bar{v}], (\mathbb{R}^2, \|\cdot\|_1))$ of essentially bounded vector functions from $[0,\bar{v}] \times [0,\bar{c}]$ to $(\mathbb{R}^2, \|\cdot\|_1)$. Define the set $\mathcal{P}_{\text{ND}} \triangleq \{p \in L^\infty : 0 \leq p(v,c) \leq 1 \forall v, c \text{ and } \mathbb{E}_c[p(v,c)] \text{ is nondecreasing}\}$. In Lemma 6.21 Item 2, we show that $\{p \in L^\infty : 0 \leq p(v) \leq 1 \forall v\}$ is weak-* compact. Note that the intersection of a weak-* compact set and a weak-* closed set is weak-* compact. Therefore, it is sufficient to show that the set of (single dimensional) nondecreasing functions is (strongly) closed and convex. $\mathcal{C} \triangleq \{\rho \in L^\infty : \mathbb{E}_c[\rho(v,c)] \text{ is nondecreasing in } v\}$ is weak-* closed. This follows by the Hahn-Banach separation theorem (see Aliprantis and Border, 2006, Theorem 5.98 in p. 214) because the set of (single dimensional) nondecreasing functions is (strongly) closed and convex. In Lemma 6.21 item 3, we show that the interim operator is weak-* continuous. The set $\mathcal{C}$ is obtained by evaluating the inverse of interim operator at the set of (single dimension) nondecreasing functions. Therefore, the set $\mathcal{C}$ is weak-* closed and $\mathcal{P}_{\text{ND}}$ is weak-* compact.

2. We argue about each term in $G(p;\beta)$ at a time. The first term $\mathbb{E}_{v,c}[p(v,c)(v-c)]$ is weak-* continuous (see Lemma 6.21 item 4). Similarly, the second term is weak-* continuous because $\mathbb{E}_{v,c}[p(v,c)v]$ is also weak-* continuous and squaring preserves. Finally, we focus on the third

term. Consider the function $\zeta : L^\infty \to \mathbb{R}$ given by

$$\zeta(p) \triangleq \mathrm{Var}_v \left[vP(v) - \int_0^v P(y)\mathrm{d}y\right]$$

$$= \mathbb{E}_v \left[\left(\int_0^v P(y)\mathrm{d}y - P(v)v - \mathbb{E}_{\tilde{v}}\left[\int_0^{\tilde{v}} P(y)\mathrm{d}y\right] + \mathbb{E}_{\tilde{v}}\left[\tilde{v}P(\tilde{v})\right]\right)^2\right].$$

In order to show that $\zeta(\cdot)$ is weak-* continuous we need to show that for any sequence of $p^m \in \mathcal{P}_{\text{ND}}$ and $p \in \mathcal{P}_{\text{ND}}$ such that $p^m \rightharpoonup p$ it follows $\zeta(p^m) \rightharpoonup \zeta(p)$. Note that $\zeta(p)$ is (strongly) continuous and weak-* lower semicontinuous, but not necessarily weak-* continuous. However, we will show that when we restrict attention to the set $\mathcal{P}_{\text{ND}}$, the function is weak-* continuous. A key step in the proof involves showing that when we restrict attention to the set $\mathcal{P}_{\text{ND}}$, weak-* convergence implies convergence almost everywhere.

Lemma 6.11. *Let P^m be a sequence of interim allocations such that $P^m \rightharpoonup P$ with P, P^m in the space $L^{\infty,1} \triangleq L^{\infty,1}([0,\bar{v}], (\mathbb{R}, |\cdot|))$ of essentially bounded vector functions from $[0,\bar{v}]$ to $(\mathbb{R}, |\cdot|)$. Suppose that P^m and P are nondecreasing, then $P^m(v) \to P(v)$ almost everywhere in v.*

Let $P^m(v) = \mathbb{E}_c[p^m(v,c)]$ and $P(v) = \mathbb{E}_c[p(v,c)]$. Because the interim operator is weak-* continuous, $P^m \rightharpoonup P$. Note also that P^m is a nondecreasing interim allocation function. Therefore, Lemma 6.11 implies $P^m \to P$ almost everywhere. Define $\tilde{\zeta} : L^{\infty,1} \to \mathbb{R}$ as $\tilde{\zeta}(P) = \mathrm{Var}_v \left[vP(v) - \int_0^v P(y)\mathrm{d}y\right]$. Because the interim allocation function P and the support of v is bounded, the dominated convergence theorem implies that the terms inside the variance are continuous. Because the summation of continuous functions is continuous and the squaring within the variance preserves continuity, we conclude that the integrand is continuous (and bounded). Using the dominated convergence theorem one more time, it follows that, $\tilde{\zeta}(P)$ is strongly continuous. Recalling that $\zeta(p) = \tilde{\zeta}(P)$ for $\mathbb{E}_{v,c}[p(v,c)] = P(v)$, it follows $\zeta(p^m) \rightharpoonup \zeta(p)$. Therefore, $\zeta(p)$ is weak-* continuous, and thus $G(p;\beta)$ is weak-* continuous in p because the sum of weak-* continuous terms is weak-* continuous.

3. Note that the derivative of $G(p;\beta)$ with respect to β is given as follows:

$$\frac{\mathrm{d}G(p;\beta)}{\mathrm{d}\beta} = \frac{b}{2\beta^2}\zeta(p).$$

By the previous item, weak-* continuity in p immediately holds. Moreover, this function is a continuous function of β for $\beta > 0$. Therefore the result follows because the product is continuous. $\square$

Proof of Lemma 6.11. Let $L^{1,1} \triangleq L^{1,1}([0,\bar{v}], (\mathbb{R}, |\cdot|))$ denote the pre-dual of $L^{\infty,1}$, i.e., a measurable function $\phi : [0,\bar{v}] \to \mathbb{R}$ lies in $L^{1,1}$ if $\int_0^{\bar{v}} |\phi(v)| \mathrm{d}v < \infty$. Recall that $P^m \rightharpoonup P$ is equivalent to $\int_0^{\bar{v}} \phi(v) P^m(v) \mathrm{d}v \to \int_0^{\bar{v}} \phi(v) P(v) \mathrm{d}v$ for all $\phi \in L^{1,1}$. Because P is monotone, it is continuous almost everywhere. Additionally, because $L^{\infty,1}$ is an equivalent class we can assume without loss that P is right-continuous. Let z be a point of continuity of P. For z, we are going to consider two functions which has jumps around z from right and left directions.

Step 1: We first consider the left hand side. Let

$$\phi_\varepsilon(y) = \begin{cases} \frac{1}{\varepsilon}\,, & z - \varepsilon \le y \le z\,, \\ 0\,, & \text{otherwise,} \end{cases}$$

and

$$\psi_\varepsilon(y) = \int_0^y \phi_\varepsilon(x) \mathrm{d}x = \begin{cases} 0\,, & 0 \le y \le z - \varepsilon\,, \\ \frac{y - z + \varepsilon}{\varepsilon}\,, & z - \varepsilon \le y \le z\,, \\ 1\,, & z \le y \le \bar{v}\,. \end{cases}$$

Observe that $\phi_\varepsilon \in L^1$ because $\int_0^{\bar{v}} \phi_\varepsilon(x) \mathrm{d}x = 1$. Since $\phi_\varepsilon \in L^1$, it follows that $\int_0^1 \phi_\varepsilon(v) P^m(v) \mathrm{d}v \to \int_0^1 \phi_\varepsilon(v) P(v) \mathrm{d}v$. Note that ψ_ε is continuous so by using integration by parts, we obtain the following.

$$\int_0^{\bar{v}} \phi_\varepsilon(v) P(v) \mathrm{d}v = \int_0^{\bar{v}} P(v) \mathrm{d}\psi_\varepsilon(v) = P(\bar{v}_+)\psi_\varepsilon(\bar{v}) - P(0)\psi_\varepsilon(0) - \int_0^{\bar{v}} \psi_\varepsilon(v) \mathrm{d}P(v)$$
$$= P(\bar{v}_+) - \int_0^{\bar{v}} \psi_\varepsilon(v) \mathrm{d}P(v) = P(\bar{v}) - \int_0^{\bar{v}} \psi_\varepsilon(v) \mathrm{d}P(v)\,,$$

where we denote by $P(\bar{v}_+)$ the right limit at $\bar{v}$. Here, the first equality is obtained by changing the integrand and the second equality follows from integration by parts. The third equality follows because $\psi_\varepsilon(\bar{v}) = 1$ and $\psi_\varepsilon(0) = 0$. Finally, the last equality holds because P is right continuous.

Note that $\psi_\varepsilon \to \mathbf{1}\{y \ge z\}$ pointwise as $\varepsilon \to 0$ and ϕ_ε is bounded. Thus, the bounded convergence theorem implies the following.

$$\int_0^{\bar{v}} \phi_\varepsilon(v) P(v) \mathrm{d}v = P(\bar{v}) - \int_0^{\bar{v}} \psi_\varepsilon(v) \mathrm{d}P(v) \to P(\bar{v}) - \int_{[z,\bar{v}]} \mathrm{d}P(v) = P(z_-) = P(z)\,. \tag{6.16}$$

Here, the last equality follows because z is a point of continuity.

Next, we consider P^m.

$$\int_0^{\bar{v}} \phi_\varepsilon(v) P^m(v) \mathrm{d}v = P^m(\bar{v}_+) - \int_0^{\bar{v}} \psi_\varepsilon(v) \mathrm{d}P^m(v) \le P^m(\bar{v}_+) - \int_{[z,\bar{v}]} \mathrm{d}P^m(v) = P^m(z_-)\,. \quad (6.17)$$

Here, the inequality holds because $\psi_\varepsilon(v) \ge \mathbf{1}\{v \ge z\}$.

Now, we combine these observations to show the following.

$$\liminf_{m\to\infty} P^m(z_-) \ge \liminf_{m\to\infty} \int_0^{\bar{v}} \phi_\varepsilon(v) P^m(v) \mathrm{d}v = \int_0^{\bar{v}} \phi_\varepsilon(v) P(v) \mathrm{d}v.$$

Here, the first inequality follows from (6.17) and the equality follows because $P^m \rightharpoonup P$. When we take the limit of ε to 0, and using (6.16) we obtain

$$\liminf_{m\to\infty} P^m(z_-) \ge P(z)\,.$$

Step 2: Now, we redefine ϕ_ε such that its jump point is at the right hand side of z. Let

$$\phi_\varepsilon(y) = \begin{cases} \frac{1}{\varepsilon}\,, & z \le y \le z + \varepsilon\,, \\ 0\,, & \text{otherwise.} \end{cases}$$

Similarly, we define ψ_ε.

$$\psi_\varepsilon(y) = \begin{cases} 0\,, & 0 \le y \le z\,, \\ \frac{y-z}{\varepsilon}\,, & z \le y \le z + \varepsilon\,, \\ 1\,, & z + \varepsilon \le y \le \bar{v}\,. \end{cases}$$

Following a similar argument (with $\psi_\varepsilon(v) \to \mathbf{1}\{y > z\}$), it follows that

$$\int_0^{\bar{v}} \phi_\varepsilon(v) P(v) \mathrm{d}v = P(\bar{v}) - \int_0^{\bar{v}} \psi_\varepsilon(v) \mathrm{d}P(v) \to P(z_+) = P(z)\,. \quad (6.18)$$

Again, the last equality follows because z is a continuity point. Next, we consider P^m.

$$\int_0^{\bar{v}} \phi_\varepsilon(v) P^m(v) \mathrm{d}v = P^m(\bar{v}_+) - \int_0^{\bar{v}} \psi_\varepsilon(v) \mathrm{d}P^m(v) \ge P^m(\bar{v}_+) - \int_{(z,\bar{v}]} \mathrm{d}P^m(v) = P^m(z_+)\,. \quad (6.19)$$

The inequality here follows from the fact that $\psi_\varepsilon(v) \leq \mathbf{1}\{y > z\}$. Combining (6.18) and (6.19), we obtain

$$\limsup_{m\to\infty} P^m(z_+) \leq \limsup_{m\to\infty} \int_0^{\bar{v}} \phi_\varepsilon(v) P^m(v)\mathrm{d}v = \int_0^{\bar{v}} \phi_\varepsilon(v)P(v)\mathrm{d}v\,.$$

After taking ε to its limit, we obtain

$$\limsup_{m\to\infty} P^m(z_+) \leq P(z)\,.$$

Finally, we consider the fact that P^m are nondecreasing to obtain

$$P(z) \leq \liminf_{m\to\infty} P^m(z_-) \leq \liminf_{m\to\infty} P^m(z_+) \leq \limsup_{m\to\infty} P^m(z_+) \leq P(z)$$

which implies $\lim_{m\to\infty} P^m(z_+) = \lim_{m\to\infty} P^m(z_-) = P(z)$ almost everywhere. $\square$

6.5 Appendix for Chapter 3 Section 3.4

6.5.1 Proof of Proposition 3.7

Proof. Throughout this proof we assume that, for each agent i, values lie in the discrete set $\mathcal{V}$ and $f_i(x) \triangleq \mathbb{P}\{v_i = x\}$ is the probability of observing $x \in \mathcal{V}$.[1] We first prove the if part (i.e., condition (3.15) implies securing efficient utility levels) and then the only if part (i.e., securing efficient utility levels implies condition (3.15)).

Step 1 (if part). In this step, we assume that $\beta^{\tau_\beta} \to 1$ and $\tau_\beta \to \infty$ as β approaches 1. We consider that agent i follows an approximately truthful reporting strategy dubbed truthful when possible (TWP) and denoted by σ_i^{TWP}. In this strategy, agent i reports truthfully if there is available budget, otherwise selects a random report from the set of values with available budget. We denote by $V_i(\boldsymbol{\pi}^{\text{JS}}, \sigma_i^{\text{TWP}}, \boldsymbol{\sigma}_{-i})$ the utility of agent i when he follows the TWP strategy and the other agent uses a strategy $\boldsymbol{\sigma}_{-i}$ under the budget-based mechanism $\boldsymbol{\pi}^{\text{JS}}$ proposed in Section 3.4. Moreover, we denote by $\tilde{\mathbf{1}}\{v_i \geq v_j\}$ the indicator function which takes value 1 (0) if $v_i > v_j$ ($v_i < v_j$), and 0.5 if $v_i = v_j$. In this step, we show that the utility of agent i converges to the efficient utility level when he employs the TWP strategy regardless of the competitor's strategy, i.e., $\lim_{\beta\uparrow 1} \inf_{\boldsymbol{\sigma}_{-i}} V_i(\boldsymbol{\pi}^{\text{JS}}, \sigma_i^{\text{TWP}}, \boldsymbol{\sigma}_{-i}) = \mathbb{E}[v_i \tilde{\mathbf{1}}\{v_i \geq v_j\}]$. In particular, we provide a lower bound for $V_i(\boldsymbol{\pi}^{\text{JS}}, \sigma_i^{\text{TWP}}, \boldsymbol{\sigma}_{-i})$ whose error term vanishes as β converges to 1.

Let $V_i^{(k)}(\boldsymbol{\pi}^{\text{JS}}, \sigma_i^{\text{TWP}}, \boldsymbol{\sigma}_{-i}) = \beta^{-(k-1)\tau_\beta} \mathbb{E}\left[\sum_{t=(k-1)\tau_\beta+1}^{k\tau_\beta} \beta^{t-1} v_{i,t} \pi_{i,t}^{\text{JS}}\right]$ be the utility of agent i in the k-th cycle (discounted to the beginning of the cycle). Here, $\pi_{i,t}^{\text{JS}} = \tilde{\mathbf{1}}\{\hat{v}_{i,t} \geq \hat{v}_{-i,t}\}$ represents the probability that agent i obtains the resource at time t under the budget-based mechanism when reports of agents i and $-i$ are $\hat{v}_{i,t}$ and $\hat{v}_{-i,t}$, respectively. Note that we denote by $\hat{v}_{i,t}$ the report of agent i in period t, and in this notation we do not impose a relation between the report $\hat{v}_{i,t}$ and the value $v_{i,t}$. We proceed by lower bounding the utility of the first cycle. Because this bound applies uniformly over the competitor's strategy, we obtain the final lower bound by summing over all cycles, i.e., using that $V_i(\boldsymbol{\pi}^{\text{JS}}, \sigma_i^{\text{TWP}}, \boldsymbol{\sigma}_{-i}) = (1-\beta)\sum_{k=1}^{\infty} \beta^{(k-1)\tau_\beta} V_i^{(k)}(\boldsymbol{\pi}^{\text{JS}}, \sigma_i^{\text{TWP}}, \boldsymbol{\sigma}_{-i})$.

Let T be the stopping time corresponding to the first time that the budget of some type of agent

[1] Although we use $f(\cdot)$ to denote the p.d.f. of the continuous values, in this section we use $f(\cdot)$ to denote the p.m.f. of the discrete values.

i depletes. This is given by:

$$T = \min_{x \in \mathcal{V}} \inf \left\{ t \geq 1 : \sum_{j=1}^{t} \mathbf{1}\{v_{i,j} = x\} = \tau_\beta f_i(x) \right\} .$$

Denote by $\pi_{i,t}^{\text{TRUTH}} = \tilde{\mathbf{1}}\{v_{i,t} \geq \hat{v}_{-i,t}\}$ the allocation when agent i reports truthfully at period t. We can decompose the utility of the first cycle as follows:

$$\begin{aligned} V_i^{(1)}(\boldsymbol{\pi}^{\text{JS}}, \sigma_i^{\text{TWP}}, \boldsymbol{\sigma}_{-i}) &= \mathbb{E}\left[\sum_{t=1}^{T} \beta^{t-1} v_{i,t}\pi_{i,t}^{\text{TRUTH}} + \sum_{t=T+1}^{\tau_\beta} \beta^{t-1} v_{i,t}\pi_{i,t}^{\text{JS}}\right] , \\ &= \underbrace{\mathbb{E}\left[\sum_{t=1}^{\tau_\beta} \beta^{t-1} v_{i,t}\pi_{i,t}^{\text{TRUTH}}\right]}_{(I)} - \underbrace{\mathbb{E}\left[\sum_{t=T+1}^{\tau_\beta} \beta^{t-1} v_{i,t}(\pi_{i,t}^{\text{TRUTH}} - \pi_{i,t}^{\text{JS}})\right]}_{(II)} . \end{aligned}$$

Here, $\pi_{i,t}^{\text{TRUTH}}$ represents the allocation when agent i reports truthfully. We next bound each term at a time. First, the lower bound for (I):

$$\begin{aligned} (I) &= \sum_{t=1}^{\tau_\beta} \beta^{t-1} \mathbb{E}\left[v_{i,t}\tilde{\mathbf{1}}\{v_{i,t} \geq \hat{v}_{-i,t}\}\right] \overset{(i)}{=} \sum_{t=1}^{\tau_\beta} \beta^{t-1} \mathbb{E}_{\hat{v}_{-i,t}}\left[\mathbb{E}[v_i\tilde{\mathbf{1}}\{v_i \geq \hat{v}_{-i,t}\}]\right] \\ &\overset{(ii)}{=} \sum_{t=1}^{\tau_\beta} \beta^{t-1} \sum_{x \in \mathcal{V}} \mathbb{P}(\hat{v}_{-i,t} = x)\mathbb{E}[v_i\tilde{\mathbf{1}}\{v_i \geq x\}] \overset{(iii)}{=} \sum_{x \in \mathcal{V}} \mathbb{E}\left[\sum_{t=1}^{\tau_\beta} \beta^{t-1}\mathbf{1}\{\hat{v}_{-i,t} = x\}\right] \mathbb{E}[v_i\tilde{\mathbf{1}}\{v_i \geq x\}] \\ &\overset{(iv)}{\geq} \sum_{x \in \mathcal{V}} \beta^{\tau_\beta - 1}\tau_\beta f_{-i}(x)\mathbb{E}[v_i\tilde{\mathbf{1}}\{v_i \geq x\}] \overset{(v)}{=} \beta^{\tau_\beta - 1}\tau_\beta \mathbb{E}[v_i\tilde{\mathbf{1}}\{v_i \geq v_j\}] . \end{aligned}$$

Because the strategy of the competitor is non-anticipative, and values are independent across agents, and identically distributed across periods, we obtain equality (i). Expanding the expectation yields (ii). Further changing the order of summations and writing the probability as the expectation, we obtain (iii). We next replace β^{t-1} with $\beta^{\tau_\beta - 1}$ for all $t \leq \tau_\beta$, because $\beta \in (0,1)$. Moreover, the budget-based mechanism implies that regardless the distribution of $\hat{v}_{-i,t}$, we have $\sum_{t=1}^{\tau_\beta} \mathbf{1}\{\hat{v}_{-i,t} = x\} = \tau_\beta f_{-i}(x)$. Therefore, the lower bound (iv) follows. Finally, the last equation, (v), follows from definition.

Next we identify an upper bound for (II). Note that $\bar{v} \geq v_{i,t}(\pi_{i,t}^{\text{TRUTH}} - \pi_{i,t}^{\text{JS}})$ because allocations lie in $[0,1]$ and $\beta \leq 1$, we obtain the following upper bound:

$$(II) \leq \bar{v}\mathbb{E}\left[(\tau_\beta - T)^+\right] ,$$

where $x^+ = \max(x, 0)$ is the positive part of a number $x \in \mathbb{R}$. Consider a coupling in which agents are allowed to report even after the end of the cycle and agent i is allowed to report a type even when its budget is depleted. Let $T(x)$ be the stopping time corresponding to the first time that budget of type $x \in \mathcal{V}$ of agent i depletes. Because $T = \min_{x \in \mathcal{V}} T(x)$ and $T(x)$ is distributed as a negative binomial random variable with mean τ_β and variance $\tau_\beta(1 - f_i(x))/f_i(x)$, we obtain that

$$\begin{aligned}
\mathbb{E}\left[(\tau_\beta - T)^+\right] &= \mathbb{E}\left[\max_{x \in \mathcal{V}} (\tau_\beta - T(x))^+\right] \leq \left(\mathbb{E}\left[\max_{x \in \mathcal{V}} (T(x) - \tau_\beta)^2\right]\right)^{1/2} \\
&\leq \left(\sum_{x \in \mathcal{V}} \mathrm{Var}(T(x))\right)^{1/2} = \tau_\beta^{1/2} \left(\sum_{x \in \mathcal{V}} \frac{1 - f_i(x)}{f_i(x)}\right)^{1/2},
\end{aligned}$$

where the first inequality follows from Jensen's inequality and $x^+ \leq x^2$ for all $x \in \mathbb{R}$, and the second inequality holds because the maximum of a set of non-negative numbers is less than or equal to the sum of the elements in the set and using that $\mathbb{E}[T(x)] = \tau_\beta$. This implies that $(II) \leq \tau_\beta^{1/2}\Delta$, where $\Delta = \bar{v} \left(\sum_{x \in \mathcal{V}} (1 - f_i(x))/(f_i(x))\right)^{1/2}$.

Putting everything together we obtain that

$$\begin{aligned}
V_i(\boldsymbol{\pi}^{\text{JS}}, \sigma_i^{\text{TWP}}, \boldsymbol{\sigma}_{-i}) &= (1 - \beta) \sum_{k=1}^{\infty} \beta^{(k-1)\tau_\beta} V_i^{(k)}(\boldsymbol{\pi}^{\text{JS}}, \sigma_i^{\text{TWP}}, \boldsymbol{\sigma}_{-i}) \\
&\geq (1 - \beta) \sum_{k=1}^{\infty} \beta^{\tau_\beta(k-1)} \left(\beta^{\tau_\beta - 1} \tau_\beta \mathbb{E}[v_i \tilde{\mathbf{1}}\{v_i \geq v_j\}] - \tau_\beta^{1/2}\Delta\right) \\
&= \mathbb{E}[v_i \tilde{\mathbf{1}}\{v_i \geq v_j\}] \frac{(1 - \beta)\beta^{\tau_\beta}\tau_\beta}{(1 - \beta^{\tau_\beta})} - \frac{(1 - \beta){\tau_\beta}^{1/2}}{(1 - \beta^{\tau_\beta})}\Delta.
\end{aligned}$$

Lemma 6.12, items (a) and (b) show the coefficients of the first term and second term converge to one and zero, respectively, as β approaches one and β^{τ_β} goes to infinity. This concludes this step.

Lemma 6.12. *It follows that:*

(a) $1 \geq \dfrac{(1 - \beta)\beta^{\tau_\beta}\tau_\beta}{(1 - \beta^{\tau_\beta})} \geq \beta^{\tau_\beta + 1}$.

(b) $\dfrac{(1 - \beta)\tau_\beta^{1/2}}{(1 - \beta^{\tau_\beta})} \to 0$ *as* $\beta \to 1$ *and* $\beta^{\tau_\beta} \to \infty$.

(c) $\dfrac{(1 - \beta)\beta^{\tau_\beta}\sqrt{\tau_\beta}}{(1 - \beta^{\tau_\beta})} \geq \beta^{\tau_\beta}\beta^{1/2}(1 - \beta)^{1/2}$.

Step 2 (only if). We prove the contrapositive, i.e., if either $\limsup_{\beta\to 1}\beta^{\tau_\beta} < 1$ or $\limsup_{\beta\to 1}\tau_\beta < \infty$, then there is no strategy for agent i that can secure efficient utility levels. We prove each case at a time. In both cases we consider a setting in which agents' values can be either high $v_H = 1$ or low $v_L = 0$ with equal probabilities, and the competitor (agent $-i$) follows a blocking strategy. The blocking strategy is one such that the competitor reports v_H in the first $\tau_\beta/2$ periods and v_L in the remaining periods. More formally, the blocking strategy is given as follows:

$$\sigma^{\text{BLOCK}}_{-i,t} = \begin{cases} v_H & t \le \tau_\beta/2 \\ v_L & t > \tau_\beta/2 \end{cases} .$$

Step 2 (only if, $\limsup_{\beta\to 1}\beta^{\tau_\beta} < 1$). We show that when the competitor (agent $-i$) follows the blocking strategy, there is no strategy that secures an efficient utility level when $\limsup_{\beta\to 1}\beta^{\tau_\beta} < 1$. We proceed by upper bounding the utility of the first cycle $V_i^{(1)}(\boldsymbol{\pi}^{\text{JS}}, \sigma_i, \boldsymbol{\sigma}^{\text{BLOCK}}_{-i})$ over all strategies σ_i for agent i. Consider a relaxation in which the budget constraints of agent i are relaxed, i.e., agent i is allowed to report any type without restrictions. This clearly provides an upper bound to the expected utility of the agent. Because the probability that the resource is allocated is increasing with the report, in the absence of budget constraints, agent i would report v_H in every time period. Therefore,

$$V_i^{(1)}(\boldsymbol{\pi}^{\text{JS}}, \sigma_i, \boldsymbol{\sigma}^{\text{BLOCK}}_{-i}) \le \sum_{t=1}^{\tau_\beta}\beta^{t-1}\frac{\mathbb{E}[v]}{2} + \sum_{t=\tau_\beta/2+1}^{\tau_\beta}\beta^{t-1}\mathbb{E}[v] = \frac{\mathbb{E}[v]}{2}\frac{(1-\beta^{\tau_\beta/2})(1+2\beta^{\tau_\beta/2})}{(1-\beta)} ,$$

where we used that the probability of obtaining the resource is $1/2$ in the first $\tau_\beta/2$ periods and one in the remaining periods. By summing over all cycles we get the following upper bound for the total discounted expected utility of agent i:

$$V_i(\boldsymbol{\pi}^{\text{JS}}, \sigma_i, \boldsymbol{\sigma}^{\text{BLOCK}}_{-i}) \le \frac{\mathbb{E}[v]}{2}\frac{(1-\beta^{\tau_\beta/2})(1+2\beta^{\tau_\beta/2})}{1-\beta^{\tau_\beta}} = \frac{\mathbb{E}[v]}{2}\frac{1+2\beta^{\tau_\beta/2}}{1+\beta^{\tau_\beta/2}} .$$

Because the function $g(x) = (1+2x)/(1+x)$ is increasing and continuous for $x \ge 0$, taking limits and using $\limsup_{\beta\to 1}\beta^{\tau_\beta} < 1$, we obtain that

$$\lim_{\beta\to 1}\sup_{\sigma_i} V_i(\boldsymbol{\pi}^{\text{JS}}, \sigma_i, \boldsymbol{\sigma}^{\text{BLOCK}}_{-i}) < \frac{3}{4}\mathbb{E}[v] = \frac{3}{8} .$$

In this case, the first-best total discounted expected utility is $\mathbb{E}[v_i \tilde{\mathbf{1}}\{v_i \geq v_j\}] = 3/8 v_H + 1/8 v_L = 3/8$. The result follows.

Step 2 (only if, $\limsup_{\beta \to 1} \tau_\beta < \infty$). In this case we provide an upper bound for the term $V_i^{(1)}(\boldsymbol{\pi}^{\text{JS}}, \sigma_i, \boldsymbol{\sigma}_{-i}^{\text{BLOCK}})$ using an information relaxation technique, which is a different approach from the previous case. Specifically, we consider a perfect information relaxation in which agent i knows the entire vector of values at time zero. Let $\mathbf{v}_i = (v_{i,1}, \ldots, v_{i,\tau_\beta}) \in \{v_L, v_H\}^{\tau_\beta}$ be a realization of agent i's values for the first cycle and let $V_i^{(1)}(\mathbf{v}_i)$ be the optimal utility of agent i with perfect information of $\mathbf{v}_i$. Because every non-anticipative policy is feasible in the perfect information relaxation, we have that $V_i^{(1)}(\boldsymbol{\pi}^{\text{JS}}, \sigma_i, \boldsymbol{\sigma}_{-i}^{\text{BLOCK}}) \leq \mathbb{E}_{\mathbf{v}_i}[V_i^{(1)}(\mathbf{v}_i)]$.

Because $\beta \leq 1$, we can upper bound the optimal utility with perfect information, for every realization $\mathbf{v}_i$ as follows

$$V_i^{(1)}(\mathbf{v}_i) \leq \max_{\sigma_t \in \{0,1\}} \sum_{t=1}^{\tau_\beta/2} v_{i,t} \frac{\sigma_t}{2} + \sum_{t=\tau_\beta/2+1}^{\tau_\beta} v_{i,t} \frac{\sigma_t + 1}{2}$$

$$\text{st.} \sum_{t=1}^{\tau_\beta} \sigma_t = \tau_\beta/2 \,,$$

where $\sigma_t = \mathbf{1}\{\hat{v}_{i,t} = v_H\} \in \{0,1\}$ indicates whether agent i reports v_H in period t. In the objective we used that the competitor reports v_H in the first $\tau_\beta/2$ periods and v_L in the remaining to determine the allocation probability.

Note that the objective of the previous optimization problem can be written as

$$\frac{1}{2} \sum_{t=1}^{\tau_\beta} v_{i,t} \sigma_t + \frac{1}{2} \sum_{t=\tau_\beta/2+1}^{\tau_\beta} v_{i,t} \,.$$

Because the second term is independent of the decision variables, the agent with perfect information should report high ($\sigma_t = 1$) in every period in which the value is high. If the total number of periods in which the value is high is less than the budget, then the agent should report high in periods in which the value is low until the budget is depleted. Let $H(\mathbf{v}_i) = \sum_{i=1}^{\tau_\beta} \mathbf{1}\{v_{i,t} = v_H\}$ denote the total number of high values, v_H in the first cycle. Using this notation, the optimal value of the first term is:

$$\frac{1}{2}\sum_{t=1}^{\tau_\beta} v_{i,t}\sigma_t = \frac{1}{2}\left(v_H \min\left(H(\mathbf{v}_1), \frac{\tau_\beta}{2}\right) + v_L \left(\frac{\tau_\beta}{2} - H(\mathbf{v}_1)\right)^+\right)$$

$$= \frac{1}{2} v_H \frac{\tau_\beta}{2} - \frac{1}{2}(v_H - v_L)\left(\frac{\tau_\beta}{2} - H(\mathbf{v}_1)\right)^+ .$$

Taking expectations with respect to $\mathbf{v}_1$ and using the fact that $v_H = 1$ and $v_L = 0$ we obtain:

$$\mathbb{E}_{\mathbf{v}_i}[V_i^{(1)}(\mathbf{v}_i)] \leq \frac{3}{8}\tau_\beta - \frac{1}{2}\mathbb{E}_{\mathbf{v}_1}\left[\left(\frac{\tau_\beta}{2} - H(\mathbf{v}_1)\right)^+\right] .$$

Because the random variable $\frac{\tau_\beta}{2} - H(\mathbf{v}_1)$ has zero mean, following Marcinkiewicz-Zygmund inequality (Marcinkiewicz and Zygmund, 1937), we find a lower bound for the negative term.

$$\frac{1}{2}\mathbb{E}_{\mathbf{v}_1}\left[\left(\frac{\tau_\beta}{2} - H(\mathbf{v}_1)\right)^+\right] = \frac{1}{4}\mathbb{E}_{\mathbf{v}_1}\left[\left|\frac{\tau_\beta}{2} - H(\mathbf{v}_1)\right|\right] \geq K\sqrt{\tau_\beta} ,$$

where K is a positive constant independent of β and τ_β. Therefore, we obtain that

$$\sup_{\sigma_i} V_i^{(1)}(\boldsymbol{\pi}^{\text{JS}}, \sigma_i, \boldsymbol{\sigma}_{-i}^{\text{BLOCK}}) \leq \mathbb{E}_{\mathbf{v}_i}[V_i^{(1)}(\mathbf{v}_i)] \leq \frac{3}{8}\tau_\beta - K\sqrt{\tau_\beta} .$$

Summing over all cycles, it follows that

$$\sup_{\sigma_i} V_i(\boldsymbol{\pi}^{\text{JS}}, \sigma_i, \boldsymbol{\sigma}_{-i}^{\text{BLOCK}}) \leq \frac{(1-\beta)}{(1-\beta^{\tau_\beta})}\left(\frac{3}{8}\tau_\beta - K\sqrt{\tau_\beta}\right) .$$

Let $\bar{\tau} = \limsup_{\beta \to 1} \tau_\beta$, with $\bar{\tau} < \infty$. Because $(1-\beta)/(1-\beta^\tau)$ converges to $1/\tau$ as $\beta \to 1$ uniformly over all $\tau \geq 1$, it follows that

$$\limsup_{\beta \to 1} \sup_{\sigma_i} V_i(\boldsymbol{\pi}^{\text{JS}}, \sigma_i, \boldsymbol{\sigma}_{-i}^{\text{BLOCK}}) \leq \frac{3}{8} - \frac{K}{\sqrt{\bar{\tau}}} .$$

The result follows because $\bar{\tau} < \infty$ and $\mathbb{E}[v_i \tilde{\mathbf{1}}\{v_i \geq v_j\}] = 3/8$. □

Proof of Lemma 6.12. We consider each item separately.

(a) The upper bound is trivial. We next derive the lower bound. We can lower bound β^{τ_β} by using the fact that $\log\beta \geq 1 - 1/\beta$ as follows:

$$\beta^{\tau_\beta} = \exp\{\tau_\beta \log\beta\} \geq \exp\left\{\tau_\beta\left(1 - \frac{1}{\beta}\right)\right\} = \exp\left\{-\frac{(1-\beta)\tau_\beta}{\beta}\right\} \geq 1 - \frac{(1-\beta)\tau_\beta}{\beta}.$$

Note that $\exp(-x) \geq 1 - x$, thus the last inequality follows. Using the lower bound β^{τ_β}, we obtain the following:

$$\frac{(1-\beta)\beta^{\tau_\beta}\tau_\beta}{1-\beta^{\tau_\beta}} \geq \frac{(1-\beta)\beta^{\tau_\beta}\tau_\beta}{\frac{(1-\beta)\tau_\beta}{\beta}} = \beta^{\tau_\beta+1}.$$

(b) Because $\log\beta \leq -(1-\beta)$, we have $0 \leq (1-\beta)\tau_\beta \leq -\tau_\beta \log\beta = \log\beta^{\tau_\beta}$. Since $\beta^{\tau_\beta} \to 1$, we have $(1-\beta)\tau_\beta \to 0$. We next bound $1 - \beta^{\tau_\beta}$.

$$1 - \beta^{\tau_\beta} = 1 - \exp\{\tau_\beta \log\beta\} \overset{(i)}{\geq} 1 - \exp\{-(1-\beta)\tau_\beta\} \overset{(ii)}{\geq} (1-\beta)\tau_\beta\left(1 - \frac{(1-\beta)\tau_\beta}{2}\right).$$

First, (i) follows from $\log\beta \leq -(1-\beta)$. We obtain (ii) by using the fact that $\exp(-x) \leq 1 - x + x^2/2 = 1 - x(1 - x/2)$. Because the right-hand side is positive for β large enough, we conclude that

$$\frac{(1-\beta)\tau_\beta^{1/2}}{(1-\beta^{\tau_\beta})} \leq \tau_\beta^{-1/2}\left(1 - \frac{(1-\beta)\tau_\beta}{2}\right)^{-1},$$

which converges to zero because $(1-\beta)\tau_\beta \to 0$ and $\tau_\beta \to \infty$.

(c) From item (a) we know that $\beta^{\tau_\beta} \geq \exp\{-(1-\beta)\tau_\beta/\beta\}$. Note that $\exp(-x) \geq 1 - x \geq 1 - x^{1/2}$ for all $x \in [0,1]$, and if $x \geq 1$ then $0 \geq 1 - x^{1/2}$ and $\exp(-x) \geq 0$. These two facts imply that $\exp(-x) \geq 1 - x^{1/2}$ for all $x \in \mathbb{R}_+$. Using this fact, we obtain the following

$$\beta^{\tau_\beta} \geq 1 - \sqrt{\tau_\beta}\left(\frac{1-\beta}{\beta}\right)^{1/2}.$$

The result follows from using this lower bound in the denominator of the statement. □

6.5.2 Proof of Theorem 3.3

Proof. Let $J^{\text{JS}(k)}_{\beta,\tau_\beta}(\boldsymbol{\sigma}) = \beta^{-(k-1)\tau_\beta}\mathbb{E}\left[\sum_{t=(k-1)\tau_\beta+1}^{k\tau_\beta} \beta^{t-1}\sum_{i=1}^{n} v_{i,t}\pi^{\text{JS}}_{i,t}\right]$ be the social welfare in the k-th cycle (discounted to the beginning of the cycle). Here, $\pi^{\text{JS}}_{i,t} = \tilde{\mathbf{1}}\{\hat{v}_{i,t} \geq \hat{v}_{-i,t}\}$ represents the probability that agent i obtains the resource at time t under the budget-based mechanism when reports of agents i and $-i$ are $\hat{v}_{i,t}$ and $\hat{v}_{-i,t}$, respectively. We proceed by lower bounding the social welfare of the first cycle.

We consider a setting in which agents' values can be either high v_H or low v_L with probabilities $f_i(v_H)$ and $f_i(v_L)$, respectively. Intuitively, each time that an agent misreports his type, there is a positive probability that the outcome is not efficient (i.e., the resource is not allocated to the agent with the highest value). Let $M_t = \{\hat{v}_{i,t} \neq v_{i,t} \text{ for some } i\}$ be the event that some agent misreports his type, and $\bar{M}_t$ its complement. Therefore, the expected social welfare in the first cycle $J^{\text{JS}(1)}_{\beta,\tau_\beta\beta}(\boldsymbol{\sigma})$ can be written as the difference of the efficient social welfare and the efficiency loss as follows:

$$\begin{aligned} J^{\text{JS}(1)}_{\beta,\tau_\beta}(\boldsymbol{\sigma}) &= \mathbb{E}\left[\sum_{t=1}^{\tau_\beta} \beta^{t-1} \sum_{i=1}^{n} v_{i,t}\pi^{\text{JS}}_{i,t}\right] \\ &= \mathbb{E}\left[\sum_{t=1}^{\tau_\beta} \beta^{t-1} \left(\mathbf{1}_{\bar{M}_t} \max_i(v_{i,t}) + \mathbf{1}_{M_t} \sum_{i=1}^{n} v_{i,t}\pi^{\text{JS}}_{i,t}\right)\right] \\ &= \mathbb{E}\left[\sum_{t=1}^{\tau_\beta} \beta^{t-1} \max_i(v_{i,t})\right] - \underbrace{\mathbb{E}\left[\sum_{t=1}^{\tau_\beta} \beta^{t-1}\mathbf{1}_{M_t} \left(\max_i(v_{i,t}) - \sum_{i=1}^{n} v_{i,t}\pi^{\text{JS}}_{i,t}\right)\right]}_{(M)}, \end{aligned}$$

where the second equality follows because the welfare is $\max_i(v_{i,t})$ when all agents report truthfully. We next control the term (M), which gives the expected efficiency loss caused by the time periods in which some agent misreports his type. We lower bound (M) in terms of the expected total number of periods in which a particular agent misreports.

Lemma 6.13. *Let $M_{i,t} = \{v_{i,t} \neq \hat{v}_{i,t}\}$ be the event that agent i misreports his type in period t. There exists some $\varepsilon_1 > 0$ such that $\mathbb{E}\left[\left(\max_i(v_{i,t}) - \sum_{i=1}^{n} v_{i,t}\pi^{\text{JS}}_{i,t}\right)\mathbf{1}_{M_{i,t}}\right] \geq \varepsilon_1 \mathbb{P}\{M_{i,t}\}$.*

Using Lemma 6.13 and conditioning on the event $M_{i,t} \subseteq M_t$, we can bound the total discounted efficiency loss for the first cycle for a fixed agent i and a value x with $f_i(x) \in (0,1)$:

$$\begin{aligned} (M) &\geq \beta^{\tau_\beta}\mathbb{E}\left[\sum_{t=1}^{\tau_\beta}\left(\max_i(v_{i,t}) - \sum_{i=1}^{n} v_{i,t}\pi^{\text{JS}}_{i,t}\right)\mathbf{1}_{M_{i,t}}\right] \geq \beta^{\tau_\beta}\varepsilon_1\mathbb{E}\left[\sum_{t=1}^{\tau_\beta}\mathbf{1}_{M_{i,t}}\right] \\ &\geq \beta^{\tau_\beta}\varepsilon_1\mathbb{E}\left[\sum_{t=1}^{\tau_\beta}\mathbf{1}\{\hat{v}_{i,t} \neq x,\, v_{i,t} = x\}\right], \end{aligned}$$

where the first inequality follows because $\beta^{t-1} \leq \beta^{\tau_\beta}$ together with the fact that the welfare of the allocation can never be greater than first best for every realization, and the last because $\{\hat{v}_{i,t} \neq x,\, v_{i,t} = x\} \subseteq M_{i,t}$.

Note that the agent is forced to lie whenever the total number of realizations of type x exceeds his budget $\tau_\beta f_i(x)$. Hence, a straightforward lower bound for the expected number of times when agent i misreports when his value is x is given as follows:

$$\mathbb{E}\left[\sum_{t=1}^{\tau_\beta} \mathbf{1}\{\hat{v}_{i,t} \neq x,\, v_{i,t} = x\}\right] \geq \mathbb{E}\left[\left(\sum_{t=1}^{\tau_\beta} \mathbf{1}\{v_{i,t} = x\} - \tau_\beta f_i(x)\right)^+\right].$$

Let $X_t = \mathbf{1}\{v_{i,t} = x\} - f_i(x)$. Because the random variable $\sum_{t=1}^{\tau_\beta} X_t$ is mean zero we obtain that

$$\mathbb{E}\left[\left(\sum_{t=1}^{\tau_\beta} X_t\right)^+\right] = \frac{1}{2}\mathbb{E}\left[\left|\sum_{t=1}^{\tau_\beta} X_t\right|\right] \geq K\mathbb{E}\left[\left(\sum_{t=1}^{\tau_\beta} X_t^2\right)^{1/2}\right],$$

where the inequality follows from Marcinkiewicz-Zygmund inequality and the positive constant K is independent of τ_β. Note that $X_t^2 \geq \min(f_i(x)^2, (1-f_i(x))^2)$. Define $\varepsilon_2 = K\min(f_1(x), (1-f_1(x)))$. Therefore, $(M) \geq \varepsilon_1\varepsilon_2\beta^{\tau_\beta}\sqrt{\tau_\beta}$.

Putting everything together and summing over all cycles we obtain that

$$\begin{aligned} J^{\text{JS}}_{\beta,\tau_\beta}(\boldsymbol{\sigma}) = (1-\beta)\sum_{k=1}^{\infty} \beta^{(k-1)\tau_\beta} J^{\text{JS}(k)}_{\beta,\tau_\beta}(\boldsymbol{\sigma}) &\leq J^{\text{FB}} - \varepsilon_1\varepsilon_2\frac{(1-\beta)\beta^{\tau_\beta}\sqrt{\tau_\beta}}{(1-\beta^{\tau_\beta})} \\ &\leq J^{\text{FB}} - \varepsilon_1\varepsilon_2\beta^{\tau_\beta}\beta^{1/2}(1-\beta)^{1/2}. \end{aligned}$$

where the last inequality follows from Lemma 6.12, item (c). ☐

Proof of Lemma 6.13. Let $M_{i,t}(x)$ be the event that agent i misreports when his value is x. Let $\bar{M}_{i,t}(x)$ be the complement of $M_{i,t}(x)$. We have: Without loss of generality, let $i = 1$.

$$\mathbf{1}_{M_{1,t}} = \mathbf{1}_{M_{1,t}(v_H)} + \mathbf{1}_{M_{1,t}(v_L)} \geq \mathbf{1}_{M_{1,t}(v_H)}\big(\mathbf{1}_{M_{2,t}(v_L)} + \mathbf{1}_{\bar{M}_{2,t}(v_L)}\big) + \mathbf{1}_{M_{1,t}(v_L)}\big(\mathbf{1}_{\bar{M}_{2,t}(v_H)} + \mathbf{1}_{\bar{M}_{2,t}(v_H)}\big).$$

We first partition the event $M_{1,t}$ and then we consider conditional expectations with respect to the value of the second agent and him being truthful or not. For each case, we derive a positive lower bound and take the minimum of those bounds to characterize $\varepsilon_1 \triangleq 0.5(v_H - v_L)\min(f_2(v_L), f_2(v_H))$.

Case 1: $M_{1,t}(v_H) \cap M_{2,t}(v_L)$. Agent 1 does not report truthfully when his value is v_H, while agent 2 does not report truthfully when his value is v_L.

$$\mathbb{E}\left[\left(\max_i(v_{i,t}) - \sum_{i=1}^{n} v_{i,t}\pi^{\text{JS}}_{i,t}\right)\mathbf{1}_{M_{1,t}(v_H)}\mathbf{1}_{M_{2,t}(v_L)}\right] = (v_H - v_L)\mathbb{E}[\mathbf{1}_{M_{1,t}(v_H)}\mathbf{1}_{M_{2,t}(v_L)}].$$

Case 2: $M_{1,t}(v_H) \cap \bar{M}_{2,t}(v_L)$**.** Agent 1 does not report truthfully when his value is v_H, while agent 2 reports truthfully when his value is v_L.

$$\mathbb{E}\left[\left(\max_i(v_{i,t}) - \sum_{i=1}^{n} v_{i,t}\pi_{i,t}^{\text{JS}}\right) \mathbf{1}_{M_{1,t}(v_H)}\mathbf{1}_{\bar{M}_{2,t}(v_L)}\right] = 0.5(v_H - v_L)\mathbb{E}[\mathbf{1}_{M_{1,t}(v_H)}\mathbf{1}_{\bar{M}_{2,t}(v_L)}]\,.$$

Case 3: $M_{1,t}(v_L) \cap M_{2,t}(v_H)$**.** Agent 1 does not report truthfully when his value is v_L, while agent 2 does not report truthfully when his value is v_H.

$$\mathbb{E}\left[\left(\max_i(v_{i,t}) - \sum_{i=1}^{n} v_{i,t}\pi_{i,t}^{\text{JS}}\right) \mathbf{1}_{M_{1,t}(v_L)}\mathbf{1}_{M_{2,t}(v_H)}\right] = (v_H - v_L)\mathbb{E}[\mathbf{1}_{M_{1,t}(v_L)}\mathbf{1}_{M_{2,t}(v_H)}]\,.$$

Case 4: $M_{1,t}(v_L) \cap \bar{M}_{2,t}(v_H)$**.** Agent 1 does not report truthfully when his value is v_L, while agent 2 reports truthfully when his value is v_H.

$$\mathbb{E}\left[\left(\max_i(v_{i,t}) - \sum_{i=1}^{n} v_{i,t}\pi_{i,t}^{\text{JS}}\right) \mathbf{1}_{M_{1,t}(v_L)}\mathbf{1}_{\bar{M}_{2,t}(v_H)}\right] = 0.5(v_H - v_L)\mathbb{E}[\mathbf{1}_{M_{1,t}(v_L)}\mathbf{1}_{\bar{M}_{2,t}(v_H)}]\,.$$

Considering all these cases, it follows that

$$\begin{aligned}\mathbb{E}\left[\left(\max_i(v_{i,t}) - \sum_{i=1}^{n} v_{i,t}\pi_{i,t}^{\text{JS}}\right) \mathbf{1}_{M_{1,t}}\right] &\geq 0.5(v_H - v_L)\Big[f_2(v_L)\mathbb{E}[\mathbf{1}_{M_{1,t}(v_H)}] + f_2(v_H)\mathbb{E}[\mathbf{1}_{M_{1,t}(v_L)}]\Big] \\ &\geq 0.5(v_H - v_L)\min(f_2(v_L), f_2(v_H))\mathbb{E}[\mathbf{1}_{M_{1,t}}] = \varepsilon_1\mathbb{P}(M_{1,t})\,. \square\end{aligned}$$

6.6 Appendix for Chapter 3 Section 3.5

6.6.1 Proof of Proposition 3.8

Proof. Agent i's allocation is given in (3.9) as $\hat{p}_i(\mathbf{v}|\mathbf{u}) = \mathbf{1}\{\alpha_i^*(\mathbf{u}/\mathsf{F}_\beta)v_i \geq \max_{j\neq i} \alpha_j^*(\mathbf{u}/\mathsf{F}_\beta)v_j\}$ where $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta) = \underset{\|\boldsymbol{x}\|_1=1,\boldsymbol{x}\geq\mathbf{0}}{\operatorname{argmax}} \left\{\frac{\mathbf{u}^\intercal\boldsymbol{x}}{\mathsf{F}_\beta} - \phi(\boldsymbol{x})\right\}$.

Step 1 ($\hat{p}_i(\mathbf{v}|\mathbf{u})$ is non-decreasing in v_i). Let $v \geq v'$, and fix $\mathbf{v}_{-i} \in [0,\bar{v}]^{n-1}$ and $\mathbf{u}$. If $\alpha_i^*(\mathbf{u}/\mathsf{F}_\beta)v' \geq \max_{j\neq i} \alpha_j^*(\mathbf{u}/\mathsf{F}_\beta)v_j$ then $\alpha_i^*(\mathbf{u}/\mathsf{F}_\beta)v \geq \max_{j\neq i} \alpha_j^*(\mathbf{u}/\mathsf{F}_\beta)v_j$ because $\boldsymbol{\alpha}^*(\mathbf{u}/\mathsf{F}_\beta) \geq \mathbf{0}$. Therefore, whenever $\mathbf{1}\{\alpha_i^*(\mathbf{u}/\mathsf{F}_\beta)v' \geq \max_{j\neq i} \alpha_j^*(\mathbf{u}/\mathsf{F}_\beta)v_j\} = 1$, it follows that $\mathbf{1}\{\alpha_i^*(\mathbf{u}/\mathsf{F}_\beta)v \geq \max_{j\neq i} \alpha_j^*(\mathbf{u}/\mathsf{F}_\beta)v_j\} = 1$. Hence, $\hat{p}_i(v,\mathbf{v}_{-i}|\mathbf{u}) \geq \hat{p}_i(v',\mathbf{v}_{-i}|\mathbf{u})$.

Step 2 ($\hat{p}_i(\mathbf{v}|\mathbf{u})$ is non-decreasing in u_i). Consider the following optimization problem,

$$\sup_{\|\boldsymbol{x}\|_1=1,\boldsymbol{x}\geq\mathbf{0}} \left\{\frac{\mathbf{u}^\intercal\boldsymbol{x}}{\mathsf{F}_\beta} - \phi(\boldsymbol{x})\right\} .$$

Change variables with $\boldsymbol{a} = \boldsymbol{x}/x_i$, such that $a_i = 1$, and, for any $\boldsymbol{x}$ satisfying $\|\boldsymbol{x}\|_1 = 1$, $x_i = 1/\left(1+\sum_{j\neq i} a_j\right)$. Consequently, the optimization problem above is equivalent to $\sup_{\boldsymbol{a}_{-i}\geq\mathbf{0}} \Xi(u_i,\boldsymbol{a}_{-i})$, in which

$$\Xi(u_i,\boldsymbol{a}_{-i}) \triangleq \frac{1}{1+\sum_{j\neq i} a_j}\left(\frac{u_i + \mathbf{u}_{-i}^\intercal\boldsymbol{a}_{-i}}{\mathsf{F}_\beta} - \phi\Big((\boldsymbol{a}_{-i},1)\Big)\right) ,$$

where we use that $\phi(\cdot)$ is homogeneous of degree one. We have

$$\frac{\partial^2\Xi}{\partial u_i \partial a_j} = -\frac{1}{\mathsf{F}_\beta} < 0 .$$

Therefore, function Ξ is submodular in u_i and a_j, which implies that as u_i increases, the optimal a_j^* decreases. As a result, $\hat{p}_i(\mathbf{v}|\mathbf{u}) = \mathbf{1}\{v_i \geq \max_{j\neq i} a_j^* v_j\}$ increases with u_i. □

6.6.2 Proof of Proposition 3.9

Proof. In this proof, we first show that agent i's interim allocation function is a non-decreasing function of his value, v_i. Using this observation, we show that agent i's interim future promise function is a non-increasing function of his value v_i. Finally, using the definition of $\hat{w}_i(\mathbf{v}|\mathbf{u})$, we conclude the proof.

The interim allocation $\hat{P}_i(v_i|\mathbf{u}) = \prod_{j\neq i} F_j\left(\frac{\alpha_i^*(\mathbf{u}/\mathsf{F}_\beta)v_i}{\alpha_j^*(\mathbf{u}/\mathsf{F}_\beta)}\right)$ is a non-decreasing function of v_i because it is a product of non-decreasing functions of v_i. The interim future promise function $\hat{W}_i(v_i|\mathbf{u})$ is given in (3.10) as

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

We next show that $\hat{P}_i(v_i|\mathbf{u})$ being non-decreasing implies that $\hat{W}_i(v_i|\mathbf{u})$ is a non-increasing function of v_i. Let $v \geq v'$. Then, it follows that $\hat{P}_i(v|\mathbf{u}) \geq \hat{P}_i(y|\mathbf{u})$ for all $v \geq y$. Using this fact, we obtain the following inequality.

$$v\hat{P}_i(v|\mathbf{u}) - v'\hat{P}_i(v'|\mathbf{u}) \geq v\hat{P}_i(v|\mathbf{u}) - v'\hat{P}_i(v|\mathbf{u}) = \int_{v'}^{v}\hat{P}_i(v|\mathbf{u})\mathrm{d}y \geq \int_{v'}^{v}\hat{P}_i(y|\mathbf{u})\mathrm{d}y$$

Rearranging the terms in the inequality $v\hat{P}_i(v|\mathbf{u}) - v'\hat{P}_i(v'|\mathbf{u}) \geq \int_{v'}^{v}\hat{P}_i(y|\mathbf{u})\mathrm{d}y$, we obtain

$$\int_0^{v'}\hat{P}_i(y|\mathbf{u})\mathrm{d}y - v'\hat{P}_i(v'|\mathbf{u}) \geq \int_0^{v}\hat{P}_i(y|\mathbf{u})\mathrm{d}y - v\hat{P}_i(v|\mathbf{u})\,.$$

This observation implies that $\int_0^{v_i}\hat{P}_i(y|\mathbf{u})\mathrm{d}y - v_i\hat{P}_i(v_i|\mathbf{u})$, and, hence, $\hat{W}_i(v_i|\mathbf{u})$ is a non-increasing function of v_i.

Because $\hat{W}_i(v_i|\mathbf{u})$ is non-increasing in v_i, the future promise function $\hat{w}_i(\mathbf{v}|\mathbf{u})$ given in (3.11) as

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

is non-increasing in v_i, and non-decreasing in v_j, for all $j \neq i$. ☐

6.7 Appendix for Chapter 4 Section 4.3

6.7.1 Proof of Proposition 4.1

Before proving the proposition, we provide the values of constants used to define our mechanism. Following this definitions, we provide the proof of the proposition.

Definition 6.2. *The functions $K_1(k)$ and $K_2(k)$ are defined as follows:*

$$K_1(k) \triangleq \big(3c_{\scriptscriptstyle S}(k)c_{\scriptscriptstyle N}(k)\big)^{\frac{1}{2}}$$

$$K_2(k) \triangleq \frac{k3^{k+1}(c_{\scriptscriptstyle N}(k))^2(c_{\scriptscriptstyle S}(k))^{k+3}}{\mathbb{E}[v]c_{\scriptscriptstyle H}(k)}\,,$$

where

$$c_{\scriptscriptstyle N}(k) \triangleq \sqrt{k}(\mathbb{E}[v] + \bar{v})\,,$$

$$c_{\scriptscriptstyle S}(k) \triangleq k\bar{f}\mathbb{E}[v^2]\,,$$

$$c_{\scriptscriptstyle H}(k) \triangleq \frac{(\underline{f}\bar{v})^{k+2}(k-1)}{k(k-1)\bar{f}^2\bar{v} + \bar{f}^{k+2}\bar{v}^{k+2} + (k-1)\bar{v}^2\bar{f}^2}\frac{1}{k+1}\,.$$

Definition 6.3. *We define the sequence $\left\{\mathsf{F}_\beta^{(k,n)}\right\}_{k=1}^n$ and the constant $\underline{u}_\beta^{(n)}$ as follows:*

$$\mathsf{F}_\beta^{(k,n)} \triangleq 1 - \frac{n(k-1)\underline{u}_\beta^{(n)}}{\mathbb{E}[v]} \quad and \quad \underline{u}_\beta^{(n)} \triangleq \xi^{(n)}(1-\beta)^{\frac{1}{n+4}}\,,$$

where $\xi^{(n)} \triangleq \max\left[\frac{K_1(n)}{\underline{\beta}^{1/2}}, \max_{k=2\ldots n}\left(\frac{\mathbb{E}[v]K_2(k)}{n(k-1)\underline{\beta}^2}\right)^{\frac{1}{k+4}}\right]$, and $K_1(k)$ and $K_2(k)$ are defined in Definition 6.2. Furthermore, we determine $\underline{\beta}$ such that for all $\beta \geq \underline{\beta}$, $\mathsf{F}_\beta^{(k,n)} \in [0.5, 1]$ for all $k = 2, \ldots, n$, and $n\underline{u}_\beta^{(n)} \leq \mathbb{E}[v]$.

Proof of Proposition 4.1. We prove this result in three steps. We defer the proofs of lemmas used in these steps to the end of this appendix.

Following Proposition 3.1, it is sufficient to show that $\Omega_\beta \in B_\beta(\Omega_\beta)$. Specifically, we need to show that for an arbitrary $\mathbf{u} \in \Omega_\beta$, there exists $(\mathbf{p}, \mathbf{w})$ satisfying (IC), (FA), (PK($\mathbf{u}$)) and $\mathbf{w}(\mathbf{v}) \in \Omega_\beta$ for all $\mathbf{v} \in [0, \bar{v}]^n$. Fix $\mathbf{u} \in \Omega_\beta$ and consider our mechanism $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$. By construction of $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$, (IC), (FA), (PK($\mathbf{u}$)) are satisfied. We will analyze the last condition.

Step 1. First, we show that $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}) \geq \mathbf{0}$. Note that when $k \leq 1$ we have $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}) = \mathbf{u}$, and the result holds trivially. Therefore, the proof focuses on situations with $k > 1$. For simplicity, we use $\mathbf{w}^{(k)}$ to represent the future promises $\mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u})$. For an inactive agent i (i.e., $u_i < \underline{u}_\beta^{(n)}$), $\hat{w}_i(\mathbf{v}|\mathbf{u}) = u_i \geq 0$, thus we focus on $\mathbf{w}^{(k)}$, and the bound on the jump from $\mathbf{u}^{(k)}$ to $\mathbf{w}^{(k)}$ which is provided by the following lemma.

Lemma 6.14. *We have that* $\|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2 \leq \left(\frac{1-\beta}{\underline{\beta}}\right) \frac{(K_1(k))^2}{3\underline{u}_\beta^{(n)}} s(\mathbf{u})$.

Following this lemma, we can argue that $\|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2 < \underline{u}_\beta^{(n)}$ which implies $\mathbf{w}^{(k)} > \mathbf{0}$ and $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}) \geq \mathbf{0}$. In particular, using the fact that $(1-\beta) < (1-\beta)^{\frac{2}{n+4}}$ for $n \geq 1$, $\left(K_1(k)\right)^2/\underline{\beta} \leq \left(\xi^{(n)}\right)^2$, and $s(\mathbf{u}) \leq 1$ we have that

$$\|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2 \leq \left(\frac{1-\beta}{\underline{\beta}}\right) \frac{(K_1(k))^2}{3\underline{u}_\beta^{(n)}} \leq \frac{(1-\beta)^{\frac{2}{n+4}} \left(\xi^{(n)}\right)^2}{3\underline{u}_\beta^{(n)}} = \frac{\left(\underline{u}_\beta^{(n)}\right)^2}{3\underline{u}_\beta^{(n)}} \leq \frac{\underline{u}_\beta^{(n)}}{3}. \tag{6.20}$$

Step 2. In this step, we show that $\mathbf{w}^{(k)} \in s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathcal{U}^{(k)}$. We consider the indicator function $\mathfrak{J}$ defined in (6.8) with scalar $s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}$. For simplicity, we define $\mathfrak{J}(\boldsymbol{x}) = \mathfrak{J}\left(\boldsymbol{x}, s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\right)$ in this proof. Corollary 6.1 states that $\mathfrak{J}\left(\mathbf{w}^{(k)}\right) \leq 0$ if and only if $\mathbf{w}^{(k)} \in s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathcal{U}^{(k)}$, because $\mathbf{w}^{(k)} \geq \mathbf{0}$ as shown earlier. In order to show $\mathfrak{J}\left(\mathbf{w}^{(k)}\right) \leq 0$, we consider the quadratic expansion of $\mathfrak{J}\left(\mathbf{w}^{(k)}\right)$ around $\mathbf{u}^{(k)}$ provided in Proposition 6.4. We next prove that the necessary conditions for Proposition 6.4 to hold.

Lemma 6.15. *The future promise* $\mathbf{w}^{(k)}$ *is in* $\mathcal{S}\left(s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\right)$.

Following this lemma, we obtain that the set $\mathcal{S}\left(s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\right)$ contains all $\boldsymbol{\Theta}^{(k)}$ between $\mathbf{w}^{(k)}$ and $\mathbf{u}^{(k)}$ because it is convex. Therefore, Proposition 6.2 implies that $\boldsymbol{\alpha}^{(k)}\left(\boldsymbol{\Theta}^{(k)}/s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\right) > \mathbf{0}$.

Lemma 6.16. *The matrix* $\Pi(\boldsymbol{a})$ *is positive definite for all* $\boldsymbol{a} > \mathbf{0}$, *and its minimum eigenvalue is given by* $c_{H}(k)\left(\min_i a_i\right)^{k+1}$.

This lemma enables us to invoke Proposition 6.4. Thus, we can analyze the terms in the quadratic expansion. We have that $\mathfrak{J}\left(\mathbf{u}^{(k)}\right) \leq 0$ because $\mathbf{u}^{(k)} \in s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathcal{U}^{(k)}$. In order to show, $\mathfrak{J}\left(\mathbf{w}^{(k)}\right) \leq 0$, it is sufficient to show the following inequality holds.

$$\nabla \mathfrak{J}\left(\mathbf{u}^{(k)}\right)^{\intercal}(\mathbf{w}^{(k)}-\mathbf{u}^{(k)})+\frac{1}{2}\left(\mathbf{w}^{(k)}-\mathbf{u}^{(k)}\right)^{\intercal}\mathrm{Hess}\left(\mathfrak{J}\left(\mathbf{\Theta}^{(k)}\right)\right)\left(\mathbf{w}^{(k)}-\mathbf{u}^{(k)}\right)\leq 0\,.$$

We first consider the first-order term. The following lemma provides a lower bound for the term $-\nabla \mathfrak{J}\left(\mathbf{u}^{(k)}\right)^{\intercal}(\mathbf{w}^{(k)}-\mathbf{u}^{(k)})$.

Lemma 6.17. *The gradient of $\mathfrak{J}(\boldsymbol{x})$ evaluated at $\mathbf{u}^{(k)}$ satisfies*

$$\frac{(1-\beta)\left(1-\mathsf{F}_{\beta}^{(k,n)}\right)\mathbb{E}[v]}{2k\left(\mathsf{F}_{\beta}^{(k,n)}\right)^{2}}\leq -\nabla \mathfrak{J}\left(\mathbf{u}^{(k)}\right)^{\intercal}(\mathbf{w}^{(k)}-\mathbf{u}^{(k)})\,.$$

We next consider the second-order term.

Lemma 6.18. *We have that*

$$\frac{1}{2}\left(\mathbf{w}^{(k)}-\mathbf{u}^{(k)}\right)^{\intercal}\mathit{Hess}\left(\mathfrak{J}\left(\mathbf{\Theta}^{(k)}\right)\right)\left(\mathbf{w}^{(k)}-\mathbf{u}^{(k)}\right)\leq \frac{1}{2}\frac{1}{\left(\mathsf{F}_{\beta}^{(k,n)}\right)^{2}}\frac{1}{\left(\underline{u}_{\beta}^{(n)}\right)^{k+3}}\left(\frac{1-\beta}{\underline{\beta}}\right)^{2}\frac{\mathbb{E}[v]K_{2}(k)}{k}\,.$$

For $n \geq k$, it follows that $(1-\beta)^{\frac{k+4}{n+4}} \geq (1-\beta)$, and $\xi^{(n)}\left(n(k-1)\beta^{2}\right)^{\frac{1}{k+4}} \geq (\mathbb{E}[v]K_{2}(k))^{\frac{1}{k+4}}$, thus we have that $(1-\mathsf{F}_{\beta}^{(k,n)})\left(\underline{u}_{\beta}^{(n)}\right)^{k+3}\underline{\beta}^{2} \geq (1-\beta)K_{2}(k)$. Finally, we use this observation to replace $(1-\beta)K_{2}(k)$ in the upper bound provided in Lemma 6.18 and show that this upper bound is smaller than the lower bound on the term with gradient (negative term) given in Lemma 6.14. Consequently, it follows that $\mathbf{w}^{(k)} \in s(\mathbf{u})\mathsf{F}_{\beta}^{(k,n)}\mathcal{U}^{(k)}$.

Note that $\mathbf{w}^{(k)} \in s(\mathbf{u})\mathsf{F}_{\beta}^{(k,n)}\mathcal{U}^{(k)}$ does not necessarily guarantee that $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ lies in the set Ω_{β} because $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ could involve fewer active agents than $\mathbf{u}$. That is, in the next round the number of active agents may decrease to $k' < k$. Therefore, in the following step, we prove that $\mathbf{w}^{(k')}(\mathbf{v}|\mathbf{u}) \in s(\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}))\mathsf{F}_{\beta}^{(k',n)}\mathcal{U}^{(k')}$.

Step 3. For simplicity, we use vector $\mathbf{w}^{(k')}$ to represent $\mathbf{w}^{(k')}(\mathbf{v}|\mathbf{u})$. The support function of the set $s(\mathbf{u})\mathsf{F}_{\beta}^{(k,n)}\mathcal{U}^{(k)}$ is $s(\mathbf{u})\mathsf{F}_{\beta}^{(k,n)}\phi^{(k)}(\cdot)$. According to Schneider (2013, p. 44), the support function satisfies that $\boldsymbol{z} \in s(\mathbf{u})\mathsf{F}_{\beta}^{(k,n)}\mathcal{U}^{(k)}$ if and only if $\boldsymbol{z}^{\intercal}\boldsymbol{x}^{(k)} \leq s(\mathbf{u})\mathsf{F}_{\beta}^{(k,n)}\phi^{(k)}(\boldsymbol{x}^{(k)})$ for all $\boldsymbol{x}^{(k)} \in \mathbb{R}^{k}$. Therefore, the fact that $\mathbf{w}^{(k)} \in s(\mathbf{u})\mathsf{F}_{\beta}^{(k,n)}\mathcal{U}^{(k)}$ implies that $(\mathbf{w}^{(k)})^{\intercal}\boldsymbol{x}^{(k)} \leq s(\mathbf{u})\mathsf{F}_{\beta}^{(k,n)}\phi^{(k)}(\boldsymbol{x}^{(k)})$ for all

$\boldsymbol{x}^{(k)} \in \mathbb{R}^k$.

To prove that $\mathbf{w}^{(k')} \in s(\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}))\mathsf{F}_\beta^{(k',n)}\mathcal{U}^{(k')}$, we show that $(\mathbf{w}^{(k')})^\intercal \boldsymbol{x}^{(k')}$ is less than or equal to $s(\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}))\mathsf{F}_\beta^{(k',n)}\phi^{(k')}(\boldsymbol{x}^{(k')})$ for an arbitrary $\boldsymbol{x}^{(k')} \in \mathbb{R}^{k'}$. Fix $\boldsymbol{x}^{(k')} \in \mathbb{R}^{k'}$. Let $\boldsymbol{x} \in \mathbb{R}^k$ be such that $x_i = x_i^{(k')}$ for all $i \leq k'$ and $x_i = 0$ otherwise. Therefore, we obtain $\boldsymbol{x}^\intercal \mathbf{w}^{(k)} = (\boldsymbol{x}^{(k')})^\intercal \mathbf{w}^{(k')}$, and $\phi^{(k)}(\boldsymbol{x}) = \phi^{(k')}(\boldsymbol{x}^{(k')})$. Now, using the property of the support functions, we obtain $(\mathbf{w}^{(k')})^\intercal \boldsymbol{x}^{(k')} \leq s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\phi^{(k')}(\boldsymbol{x}^{(k')})$.

We next prove that $s(\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}))\mathsf{F}_\beta^{(k',n)} \geq s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}$. Using the definitions, we get the following inequality for $\mathsf{F}_\beta^{(k,n)}$ and $\mathsf{F}_\beta^{(k',n)}$:

$$\mathsf{F}_\beta^{(k,n)} = \mathsf{F}_\beta^{(k,n)} - \mathsf{F}_\beta^{(k',n)} + \mathsf{F}_\beta^{(k,n)} = \mathsf{F}_\beta^{(k',n)} - \frac{n(k-k')\underline{u}_\beta^{(n)}}{\mathbb{E}[v]} \leq \mathsf{F}_\beta^{(k',n)}\left(1 - \frac{n(k-k')\underline{u}_\beta^{(n)}}{\mathbb{E}[v]}\right).$$

Recall that for an inactive agent i (i.e., $u_i < \underline{u}_\beta^{(n)}$) our future promise function $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$ is such that $\hat{w}_i(\mathbf{v}|\mathbf{u}) = u_i$. If an agent is inactive for a given initial state $\mathbf{u}$, then he remains to be inactive for a given future promise $\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})$. Therefore, we can bound $s(\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}))$ in terms of $s(\mathbf{u})$ as follows:

$$s(\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u})) = s(\mathbf{u}) - \frac{\sum_{i=k'+1}^{k} \hat{w}_i(\mathbf{v}|\mathbf{u})}{\mathbb{E}[v]} \overset{(a)}{\geq} s(\mathbf{u})\left(1 - \frac{(k-k')\underline{u}_\beta^{(n)}}{s(\mathbf{u})\mathbb{E}[v]}\right) \overset{(b)}{\geq} s(\mathbf{u})\left(1 - \frac{n(k-k')\underline{u}_\beta^{(n)}}{\mathbb{E}[v]}\right).$$

where (a) follows from the fact that $\underline{u}_\beta^{(n)} \geq \hat{w}_i(\mathbf{v}|\mathbf{u})$ for the inactive agents and (b) follows because $s(\mathbf{u}) \geq 1/n$ (since $\mathbb{E}[v] \geq n\underline{u}_\beta^{(n)}$ and $k > 1$). Therefore, it follows that $s(\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}))\mathsf{F}_\beta^{(k',n)} \geq s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}$.

Finally, this observation implies that $(\mathbf{w}^{(k')})^\intercal \boldsymbol{x}^{(k')} \leq s(\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}))\mathsf{F}_\beta^{(k',n)}\phi^{(k')}(\boldsymbol{x}^{(k')})$, and $\mathbf{w}^{(k')}(\mathbf{v}|\mathbf{u})$ is in $s(\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}))\mathsf{F}_\beta^{(k',n)}\mathcal{U}^{(k')}$. Hence, we obtain that

$$\hat{\mathbf{w}}(\mathbf{v}|\mathbf{u}) \in \left\{\mathbf{u} \in \mathbb{R}^n \mid \mathbf{u}^{(k')} \geq \underline{u}_\beta^{(n)},\ \mathbf{u}^{(\bar{k}')} < \underline{u}_\beta^{(n)},\ \text{ and }\ \mathbf{u}^{(k')} \in s(\mathbf{u})\mathsf{F}_\beta^{(k',n)}\mathcal{U}^{(k')}\right\} \subseteq \Omega_\beta\,. \qquad \square$$

6.7.2 Proof of Theorem 4.1

Proof. The set Ω_β is a subset of $\mathcal{U}_\beta^{(n)}$ by Proposition 4.1. Thus, the maximum achievable social welfare $J_\beta = \mathsf{F}_\beta^{(n,n)}J^{\text{FB}}$, is less than J_β^*. $\square$

6.7.3 Proofs of Lemmas 6.14, 6.15, 6.16, 6.17, 6.18, and 6.19

In the following proofs, we consider $\mathbf{u} \in \Omega_\beta$ with k active agents and the k dimensional subvector $\mathbf{u}^{(k)}$ represent the promised utilities for the active agents. We use $\boldsymbol{\Theta}^{(k)}$ to represent an arbitrary k dimensional point between $\mathbf{u}^{(k)}$ and $\mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u})$. Here the set Ω_β is characterized by constants $\underline{u}_\beta^{(n)}$ and $\left\{\mathsf{F}_\beta^{(k,n)}\right\}_{k=1}^n$ following Definition 6.3. Further, define vectors

$$\boldsymbol{\gamma} = \boldsymbol{\alpha}^{(k)}\left(\frac{\mathbf{u}^{(k)}}{s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}}\right) \text{ and } \boldsymbol{\lambda} = \boldsymbol{\alpha}^{(k)}\left(\frac{\boldsymbol{\Theta}^{(k)}}{s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}}\right). \tag{6.21}$$

Proof of Lemma 6.14. For simplicity, we use the short notation $\mathbf{w}^{(k)}$ to represent the future promises $\mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u})$. Because $\|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2 \leq \sqrt{k}\|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_\infty$, we focus on an upper bound for $\|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_\infty = \max_i |w_i^{(k)} - u_i^{(k)}|$. We start with the following expansion.

$$\begin{aligned} w_i^{(k)} - u_i^{(k)} = &\left(\frac{1-\beta}{\beta}\right)\left[u_i^{(k)} + \int_0^{v_i} P_i^{(k)}(y|\mathbf{u})\mathrm{d}y - P_i^{(k)}(v_i|\mathbf{u})v_i - \int_0^{\bar{v}} \bar{F}(y)P_i^{(k)}(y|\mathbf{u})\mathrm{d}y\right] \\ &- \frac{1}{k-1}\sum_{j\neq i}\frac{\gamma_j}{\gamma_i}\Big(W_j^{(k)}(v_j|\mathbf{u}) - \mathbb{E}_{\hat{v}_j}\left[W_j^{(k)}(\hat{v}_j|\mathbf{u})\right]\Big), \end{aligned}$$

where $P_i^{(k)}(y|\mathbf{u}) = s(\mathbf{u})\prod_{j\neq i} F\left(\frac{\gamma_i y}{\gamma_j}\right)$ and $W_j^{(k)}(v_j|\mathbf{u}) - \mathbb{E}_{\hat{v}_j}\left[W_j^{(k)}(\hat{v}_j|\mathbf{u})\right]$ is given by

$$\left(\frac{1-\beta}{\beta}\right)\left\{\int_0^{v_j} P_j^{(k)}(y|\mathbf{u})\mathrm{d}y - P_j^{(k)}(v_j|\mathbf{u})v_j - \mathbb{E}_{\hat{v}_j}\left[\int_0^{\hat{v}_j} P_j^{(k)}(y|\mathbf{u})\mathrm{d}y\right] + \mathbb{E}_{\hat{v}_j}\left[P_j^{(k)}(\hat{v}_j|\mathbf{u})\hat{v}_j\right]\right\}.$$

We prove this result in three steps. In the first two steps, we consider two cases: $|w_i^{(k)} - u_i^{(k)}| = w_i^{(k)} - u_i^{(k)}$ and $|w_i^{(k)} - u_i^{(k)}| = u_i^{(k)} - w_i^{(k)}$, and get a bound on the norm which depends on $1/\min_i \gamma_i$. In the last step, we get an upper bound for $1/\min_i \gamma_i$ that depends on $\underline{u}_\beta^{(n)}$.

Step 1 (Bounding $w_i^{(k)} - u_i^{(k)}$). Here, we remove negative terms to obtain the following bound.

$$\begin{aligned} &w_i^{(k)} - u_i^{(k)} \\ &\leq \left(\frac{1-\beta}{\beta}\right)\left[u_i^{(k)} + \int_0^{v_i} P_i^{(k)}(y|\mathbf{u})\mathrm{d}y + \frac{1}{k-1}\sum_{j\neq i}\frac{\gamma_j}{\gamma_i}\left(P_j^{(k)}(v_j|\mathbf{u})v_j + \mathbb{E}_{v_j}\left[\int_0^{v_j} P_j^{(k)}(y|\mathbf{u})\mathrm{d}y\right]\right)\right]. \end{aligned}$$

Note that $P_i^{(k)}(v|\mathbf{u}) \leq s(\mathbf{u})$ for all v by definition, and $u_i^{(k)} \leq s(\mathbf{u})\mathbb{E}[v]$. Moreover, using the fact that the sum of γ_i's adds up to 1, and v_i and v_j are less than $\bar{v}$. we get the following bound.

$$w_i^{(k)} - u_i^{(k)} \leq \left(\frac{1-\beta}{\beta}\right) s(\mathbf{u})(\mathbb{E}[v] + \bar{v}) \left[1 + \frac{(1-\gamma_i)}{(k-1)\gamma_i}\right] \leq s(\mathbf{u}) \left(\frac{1-\beta}{\beta}\right) \left(\frac{\mathbb{E}[v] + \bar{v}}{\min\limits_i \gamma_i}\right) .$$

Step 2 (Bounding $u_i^{(k)} - w_i^{(k)}$). Similar to the previous step, we obtain

$$\begin{aligned} u_i^{(k)} - w_i^{(k)} \leq \left(\frac{1-\beta}{\beta}\right) &\Bigg[P_i^{(k)}(v_i|\mathbf{u})v_i + \int_0^{\bar{v}} \bar{F}(y) P_i^{(k)}(y|\mathbf{u})\mathrm{d}y \\ &+ \frac{1}{k-1} \sum_{j \neq i} \frac{\gamma_j}{\gamma_i} \left(\int_0^{v_j} P_j^{(k)}(y|\mathbf{u})\mathrm{d}y + \mathbb{E}_{\hat{v}_j}[P_j^{(k)}(\hat{v}_j|\mathbf{u})\hat{v}_j]\right)\Bigg] , \\ &\leq s(\mathbf{u}) \left(\frac{1-\beta}{\beta}\right) \left(\frac{\mathbb{E}[v] + \bar{v}}{\min\limits_i \gamma_i}\right) . \end{aligned}$$

Summarizing the previous two steps, we obtain $\|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_\infty \leq s(\mathbf{u}) \left(\frac{1-\beta}{\beta}\right) \left(\frac{\mathbb{E}[v] + \bar{v}}{\min\limits_i \gamma_i}\right)$. Recognizing $\|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2 \leq \sqrt{k}\|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_\infty$, we obtain

$$\|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2 \leq s(\mathbf{u}) \left(\frac{1-\beta}{\beta}\right) \left(\frac{\sqrt{k}(\mathbb{E}[v] + \bar{v})}{\min\limits_i \gamma_i}\right) \leq s(\mathbf{u}) \left(\frac{1-\beta}{\underline{\beta}}\right) \left(\frac{c_{\mathrm{N}}(k)}{\min\limits_i \gamma_i}\right) .$$

Step 3 (Bounding $1/\min_i \gamma_i$) In this step, we find an upper bound for $1/\min_i \gamma_i$ as follows:

$$u_i^{(k)} \overset{(a)}{\leq} s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathbb{E}[v_i \mathbf{1}\{\gamma_i v_i \geq \max_{j\neq i} \gamma_j v_j\}] \overset{(b)}{\leq} \gamma_i k \bar{f} \mathbb{E}[v^2] = c_{\mathrm{S}}(k)\gamma_i \, .$$

As discussed in Remark 3.2, $u_i^{(k)} = s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}(\nabla\phi^{(k)}(\boldsymbol{\gamma}))_i - \mu$ for some nonnegative scalar μ. Using the expression derived in Proposition 3.3 for $(\nabla\phi^{(k)}(\boldsymbol{\gamma}))_i$, we get (a). Because the p.d.f. $f(\cdot)$ is bounded in the support by $\bar{f}$ and the c.d.f. $F(\cdot)$ is always less than one, $\mathbb{E}[v_i \mathbf{1}\{\gamma_i v_i \geq \max_{j\neq i} \gamma_j v_j\}]$ is bounded by $\gamma_i \bar{f}\mathbb{E}\left[v^2\right] / \max_{j\neq i} \gamma_j$. Together with this observation, using the fact that $\max_{j\neq i} \gamma_j \geq 1/k$, and $s(\mathbf{u}) \leq 1$ and $\mathsf{F}_\beta^{(k,n)} \leq 1$, we obtain (b). Taking the minimum over i's in both sides, we get

$c_{\mathrm{S}}(k)\min_i \gamma_i \geq \min_i u_i^{(k)} \geq \underline{u}_\beta^{(n)}$. Consequently, we have that

$$\frac{\|\mathbf{w}^{(k)}-\mathbf{u}^{(k)}\|_2}{s(\mathbf{u})} \leq \left(\frac{1-\beta}{\underline{\beta}}\right)\left(\frac{c_{\mathrm{N}}(k)}{\min\limits_i \gamma_i}\right) \leq \left(\frac{1-\beta}{\underline{\beta}}\right)\left(\frac{c_{\mathrm{S}}(k)c_{\mathrm{N}}(k)}{\underline{u}_\beta^{(n)}}\right) = \left(\frac{1-\beta}{\underline{\beta}}\right)\frac{(K_1(k))^2}{3\underline{u}_\beta^{(n)}}. \qquad \square$$

Proof of Lemma 6.15. Equation (6.20) implies that $\|\mathbf{w}^{(k)}-\mathbf{u}^{(k)}\|_2 \leq \underline{u}_\beta^{(n)}/3$. Because $\mathbf{u}^{(k)} \geq \underline{u}_\beta^{(n)}$, we must have $\mathbf{w}^{(k)} > 0$. If the point $\mathbf{w}^{(k)} \in s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathcal{U}^{(k)}$, the result holds immediately. Now suppose $\mathbf{w}^{(k)} \notin s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathcal{U}^{(k)}$, and, therefore, there must exist a component j such that $w_j^{(k)} > u_j^{(k)}$. Select $\bar{a}$ such that $\bar{a} = \inf\{a : \mathbf{w}^{(k)} - a\boldsymbol{e} \leq \mathbf{u}^{(k)}\}$, therefore $\mathbf{w}^{(k)} - \bar{a}\boldsymbol{e} \in \mathcal{S}(s(\mathbf{u})\mathsf{F}_\beta^{(k,n)})$ if $\mathbf{w}^{(k)} - \bar{a}\boldsymbol{e} > \mathbf{0}$. The definition of $\bar{a}$ implies that there must exist a component i such that $w_i^{(k)} - \bar{a} = u_i^{(k)} > 0$. We next focus on components $j \neq i$. Using the fact that $\bar{a} = w_i^{(k)} - u_i^{(k)}$, we have $w_j^{(k)} - \bar{a} = w_j^{(k)} - w_i^{(k)} + u_i^{(k)}$. Thus, we only need to show that $w_j^{(k)} - w_i^{(k)} + u_i^{(k)} > 0$. Because $u_i^{(k)} + \|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2 \geq w_i^{(k)}$ and $w_j^{(k)} \geq u_j^{(k)} - \|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2 \geq \underline{u}_\beta^{(k)} - \|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2$, we have

$$w_j^{(k)} - w_i^{(k)} + u_i^{(k)} \geq \underline{u}_\beta^{(n)} - 2\|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2 > 0\,,$$

where the last inequality follows because $\underline{u}_\beta^{(n)} \geq 3\|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2$. $\square$

Definition 6.4. *We define the weighted dot product $\langle\cdot,\cdot\rangle_{\boldsymbol{x}}$ and the induced norm as follows. Let $\mathbf{z}$ and $\tilde{\mathbf{z}}$ be arbitrary vectors in $\mathbb{R}^k$.*

$$\langle \mathbf{z}, \tilde{\mathbf{z}}\rangle_{\boldsymbol{x}} = \sum_{i=1}^{k} z_i \tilde{z}_i x_i \ \textit{and}\ \|\mathbf{z}\|_{\boldsymbol{x}}^2 = \langle \mathbf{z}, \mathbf{z}\rangle_{\boldsymbol{x}} = \sum_{i=1}^{k} z_i^2 x_i\,.$$

Proof of Lemma 6.16. Denote $M = \mathrm{Hess}\left(\phi^{(k)}(\boldsymbol{a})\right)$. Proposition 3.3 shows that M is well defined for $\boldsymbol{a} > \mathbf{0}$. Thus, $\Pi(\boldsymbol{a}) = M + \boldsymbol{e}\boldsymbol{e}^\intercal$ is also well defined. We identify a lower bound on the minimum eigenvalue of the matrix $M + \boldsymbol{e}\boldsymbol{e}^\intercal$ which is positive and thus implying that $M + \boldsymbol{e}\boldsymbol{e}^\intercal$ is positive definite. We show this result in three steps.

Step 1. In this step, we provide a lower bound for the minimum eigenvalue of $M + \boldsymbol{e}\boldsymbol{e}^\intercal$. Define the following functions:

$$\omega(y) \triangleq \frac{\prod\limits_t F\left(\frac{y}{a_t}\right)}{y^2} \text{ and } \Phi \triangleq \int_0^{\bar{v}} \omega(y)\mathrm{d}y\,,$$

$$\eta(y) \triangleq y^2 \frac{f(y)}{F(y)} \mathbf{1}\left\{y \leq \bar{v}\right\} \text{ and } g_i(y) \triangleq \frac{1}{a_i}\eta\left(\frac{y}{a_i}\right) .$$

Let G be a matrix whose entries are given by $G_{ij} = \langle a_i g_i, a_j g_j \rangle_\omega \triangleq \int_0^{\bar{v}} a_i g_i(y) a_j g_j(y) \omega(y) \mathrm{d}y$.

Using the weighted norm notation in Definition 6.4, we can equivalently express M as follows:

$$M = \operatorname{diag}\left(\sum_j a_j^2 \langle g_i, g_j \rangle_\omega\right) - G \,.$$

Here, the first term defines a diagonal matrix whose i^{th} diagonal component is $\sum_j a_j^2 \langle g_i, g_j \rangle_\omega$. Define $M(y) = \operatorname{diag}\left(\sum_j a_j^2 g_i(y) g_j(y)\right) - G(y)$ to alternatively represent M, i.e., $M = \int_0^{\bar{v}} M(y)\omega(y)\mathrm{d}y$.

Next, consider the following optimization problem, whose objective value is the minimum eigenvalue of the matrix $(M + \boldsymbol{e}\boldsymbol{e}^\intercal)$,

$$\Psi \triangleq \min_{\|\boldsymbol{x}\|_2 \geq 1} \boldsymbol{x}^\intercal (M + \boldsymbol{e}\boldsymbol{e}^\intercal)\boldsymbol{x} = \min_{\|\boldsymbol{x}\|_2 \geq 1} \int_0^{\bar{v}} (\boldsymbol{x}^\intercal M(y) \boldsymbol{x}) \omega(y) \mathrm{d}y + (\boldsymbol{x}^\intercal \boldsymbol{e})^2 \ . \tag{6.22}$$

A straightforward lower bound for Ψ is obtained by taking the minimization to inside the integral. To do that, we also need to take the second term, $(\boldsymbol{x}^\intercal \boldsymbol{e})^2$, to inside the integral by properly adjusting the weight as follows:

$$\tilde{\boldsymbol{e}} \triangleq \frac{\boldsymbol{e}}{\sqrt{\Phi}} \,.$$

Therefore, we have

$$\Psi \geq \int_0^{\bar{v}} \left[\min_{\|\boldsymbol{x}\|_2 \geq 1} (\boldsymbol{x}^\intercal M(y) \boldsymbol{x}) + (\boldsymbol{x}^\intercal \tilde{\boldsymbol{e}})^2\right] \omega(y) \mathrm{d}y \,. \tag{6.23}$$

Step 2. Let $\triangle(y)$ represent $\min\limits_{\|\boldsymbol{x}\|_2\geq 1}(\boldsymbol{x}^\intercal M(y)\boldsymbol{x})+(\boldsymbol{x}^\intercal\tilde{\boldsymbol{e}})^2$. Then, the inequality (6.23) is represented as $\Psi \geq \int_0^{\bar{v}} \triangle(y)\omega(y)\mathrm{d}y$. We continue by analyzing $\triangle(y)$.

$$\begin{aligned}
\triangle(y) &= \min_{\|\boldsymbol{x}\|_2\geq 1}(\boldsymbol{x}^\intercal M(y)\boldsymbol{x})+(\boldsymbol{x}^\intercal\tilde{\boldsymbol{e}})^2 \\
&= \min_{\|\boldsymbol{x}\|_2\geq 1}\left(\sum_j a_j^2 g_j(y)\right)\left(\sum_j x_j^2 g_j(y)\right)-\left(\sum_j x_j g_j(y)a_j\right)^2+\left(\sum_j x_j\frac{1}{\sqrt{\Phi}}\right)^2 \\
&= \min_{\|\boldsymbol{x}\|_2\geq 1}\|\boldsymbol{a}\|_g^2\|\boldsymbol{x}\|_g^2-\langle\boldsymbol{x},\boldsymbol{a}\rangle_g^2+\left(\sum_j x_j g_j\frac{1}{g_j\sqrt{\Phi}}\right)^2 .
\end{aligned}$$

Denote $\boldsymbol{e}_g \triangleq \left(\frac{1}{\sqrt{\Phi}g_i}\right)_i$. Then, $\triangle(y)$ is given as follows:

$$\triangle(y) = \min_{\|\boldsymbol{x}\|_2\geq 1}\|\boldsymbol{a}\|_g^2\|\boldsymbol{x}\|_g^2-\langle\boldsymbol{x},\boldsymbol{a}\rangle_g^2+\langle\boldsymbol{x},\boldsymbol{e}_g\rangle_g^2 . \tag{6.24}$$

We next provide a lower bound for $\triangle(y)$ by using the expression (6.24). We first consider a relaxation obtained by replacing the constraint $\|\boldsymbol{x}\|_2\geq 1$ with $\|\boldsymbol{x}\|_g^2 \geq \min\limits_i g_i$ as follows:

$$\begin{aligned}
\triangle(y) &\geq \min_{\|\boldsymbol{x}\|_g^2\geq \min\limits_i g_i}\|\boldsymbol{a}\|_g^2\|\boldsymbol{x}\|_g^2-\langle\boldsymbol{x},\boldsymbol{a}\rangle_g^2+\langle\boldsymbol{x},\boldsymbol{e}_g\rangle_g^2 = \min_i g_i\|\boldsymbol{a}\|_g^2+\min_{\|\boldsymbol{x}\|_g^2\geq \min\limits_i g_i}\langle\boldsymbol{x},\boldsymbol{e}_g\rangle_g^2-\langle\boldsymbol{x},\boldsymbol{a}\rangle_g^2 , \\
&= \min_i g_i\left[\|\boldsymbol{a}\|_g^2+\min_{\|\boldsymbol{x}\|_g^2\geq 1}\langle\boldsymbol{x},\boldsymbol{e}_g\rangle_g^2-\langle\boldsymbol{x},\boldsymbol{a}\rangle_g^2\right] .
\end{aligned}$$

In order to express this lower bound on $\triangle(y)$, there is a need for a closed-form solution for the optimization problem over the differences of weighted dot products. The following lemma provides a closed-form solution for it.

Lemma 6.19. *The optimal value of the optimization problem* $\min\limits_{\|\boldsymbol{x}\|_g^2\geq 1}\langle\boldsymbol{x},\boldsymbol{e}_g\rangle_g^2-\langle\boldsymbol{x},\boldsymbol{a}\rangle_g^2$ *is given by*

$$\frac{1}{2}\left(\|\boldsymbol{e}_g\|_g^2-\|\boldsymbol{a}\|_g^2-\sqrt{\left(\|\boldsymbol{e}_g\|_g^2+\|\boldsymbol{a}\|_g^2\right)^2-(2\langle\boldsymbol{e}_g,\boldsymbol{a}\rangle_g)^2}\right) .$$

We defer the proof of this lemma to the end of this section. Using the closed-form solution given in the lemma, the lower bound on $\triangle(y)$ is updated as follows:

$$\triangle(y) \geq \frac{\min_i g_i}{2}\left[\|\boldsymbol{e}_g\|_g^2 + \|\boldsymbol{a}\|_g^2 - \sqrt{\left(\|\boldsymbol{e}_g\|_g^2 + \|\boldsymbol{a}\|_g^2\right)^2 - (2\langle \boldsymbol{e}_g, \boldsymbol{a}\rangle_g)^2}\right]$$

$$\geq \frac{\min_i g_i}{2}\left(\|\boldsymbol{e}_g\|_g^2 + \|\boldsymbol{a}\|_g^2\right)\frac{(2\langle \boldsymbol{e}_g, \boldsymbol{a}\rangle_g)^2}{2\left(\|\boldsymbol{e}_g\|_g^2 + \|\boldsymbol{a}\|_g^2\right)^2} = \min_i g_i \frac{\langle \boldsymbol{e}_g, \boldsymbol{a}\rangle_g^2}{\|\boldsymbol{e}_g\|_g^2 + \|\boldsymbol{a}\|_g^2}\,.$$

We obtain the second inequality by taking the common multiplier $\|\boldsymbol{e}_g\|_g^2 + \|\boldsymbol{a}\|_g^2$ outside the square brackets because $1 - \sqrt{1-y^2} \geq y^2/2$. We now consider each term separately to simplify the lower bound.

$$\langle \boldsymbol{e}_g, \boldsymbol{a}\rangle_g = \sum_j a_j \frac{1}{g_j\sqrt{\Phi}} g_j = \frac{1}{\sqrt{\Phi}}\,,\ \|\boldsymbol{e}_g\|_g^2 = \sum_j \frac{1}{g_j^2\Phi} g_j = \frac{1}{\Phi}\sum_j \frac{1}{g_j}\,,\ \text{and } \|\boldsymbol{a}\|_g^2 = \sum_j a_j^2 g_j\,.$$

Using these terms, we define the following lower bound.

$$\triangle(y) \geq \frac{\min_i g_i(y)}{\sum_j \frac{1}{g_j(y)} + \Phi\sum_j a_j^2 g_j(y)} = \frac{\min_i \frac{1}{a_i}\left(\frac{y}{a_i}\right)^2 \frac{f\left(\frac{y}{a_i}\right)}{F\left(\frac{y}{a_i}\right)}\mathbf{1}\{y \leq a_i\bar{v}\}}{\sum_j \frac{1}{g_j(y)} + \Phi\sum_j a_j^2 g_j(y)} \geq \frac{y\min_i \frac{1}{a_i^2}\frac{\underline{f}}{\bar{f}}\mathbf{1}\{y \leq a_i\bar{v}\}}{\frac{k}{\underline{g}(y)} + \Phi\bar{g}(y)}$$

$$\geq \frac{y\frac{\underline{f}}{\bar{f}}\mathbf{1}\{y \leq \min_i a_i\bar{v}\}}{\frac{k}{\underline{g}(y)} + \Phi\bar{g}(y)}\,,$$

where $\underline{g}$ and $\bar{g}$ are lower and upper bounds for any g_i, respectively, and given by

$$\underline{g}(y) \triangleq \frac{y\underline{f}\mathbf{1}\{y \leq \min_i a_i\bar{v}\}}{\bar{f}} \text{ and } \bar{g}(y) \triangleq \frac{\bar{v}^2\bar{f}\mathbf{1}\{y \leq \bar{v}\}}{y\underline{f}}\,.$$

Replacing $\underline{g}(y)$ and $\bar{g}(y)$ in the right-hand side, we obtain the following lower bound for $\triangle(y)$:

$$\triangle(y) \geq \frac{y^2\underline{f}^2\mathbf{1}\{y \leq \min_i a_i\}}{k\bar{f}^2 + \Phi\bar{v}^2\bar{f}^2\mathbf{1}\{y \leq \bar{v}\}\mathbf{1}\{y \leq \min_i a_i\}}\,. \tag{6.25}$$

Step 3. In this step, we find an upper bound for Φ to obtain a lower bound of $\triangle(y)$ which we denote by $\Gamma(y)$. Then, we integrate $\Gamma(y)\omega(y)$ to bound Ψ. In particular, we use the upper bound $\bar{f}$

on the p.d.f. $f(\cdot)$ in its support to get the first inequality.

$$\Phi \leq \int_0^{\bar{v}\min_i a_i} \frac{\prod_t \bar{f}\left(\frac{y}{a_t}\right)}{y^2} \mathrm{d}y + \int_{\bar{v}\min_i a_i}^{\bar{v}} \frac{1}{y^2} \mathrm{d}y \leq \frac{(\bar{f})^k}{(\min_i a_i)^k} \frac{(\bar{v}\min_i a_i)^{k-1}}{(k-1)} + \frac{1}{\bar{v}\min_i a_i}$$

$$= \frac{(\bar{f}\bar{v})^k + (k-1)}{\bar{v}(\min_i a_i)(k-1)}\,.$$

When we replace the upper bound on Φ in the right-hand side of (6.25), we get another lower bound for $\triangle(y)$. Define the piece-wise function

$$\Gamma(y) \triangleq \frac{y^2 \underline{f}^2 \mathbf{1}\{y \leq \min_i a_i\}}{k\bar{f}^2 + \frac{(\bar{f}\bar{v})^k + (k-1)}{\bar{v}(\min_i a_i)(k-1)} \bar{v}^2 \bar{f}^2 \mathbf{1}\{y \leq \bar{v}\}\mathbf{1}\{y \leq \min_i a_i\}} = \begin{cases} \frac{y^2 (k+1) c_{\mathrm{H}}(k)(\min_i a_i)}{\underline{f}^k \bar{v}^{k+1}} & \text{if } y \leq \min_i a_i \bar{v} \\ 0 & \text{otherwise}\,. \end{cases}$$

Summarizing the previous steps, we obtain $\triangle(y) \geq \Gamma(y)$.

Recall that in (6.23), we have the following inequality $\Psi \geq \int_0^{\bar{v}} \triangle(y)\omega(y)\mathrm{d}y$. Because $\Gamma(y)$ is less than $\triangle(y)$, we can lower bound Ψ as follows.

$$\Psi \geq \int_0^{\bar{v}} \Gamma(y)\omega(y)\mathrm{d}y = \frac{(k+1)c_{\mathrm{H}}(k)(\min_i a_i)}{\underline{f}^k \bar{v}^{k+1}} \int_0^{\bar{v}\min_i a_i} \prod_t F\left(\frac{y}{a_t}\right) \mathrm{d}y$$

$$\geq \frac{(k+1)c_{\mathrm{H}}(k)(\min_i a_i)}{\bar{v}^{k+1}} \frac{1}{\prod_t a_t} \int_0^{\bar{v}\min_i a_i} y^k \mathrm{d}y = \frac{c_{\mathrm{H}}(k)(\min_i a_i)^{k+2}}{\prod_t a_t} \geq c_{\mathrm{H}}(k) \left(\min_i a_i\right)^{k+1}\,.$$

This lower bound on Ψ, at the same time, is a lower bound on the minimum eigenvalue of the matrix $(M + \boldsymbol{e}\boldsymbol{e}^{\intercal})$. □

Proof of Lemma 6.17. Proposition 6.3 provides $\nabla\mathfrak{J}\left(\mathbf{u}^{(k)}\right)$ because Lemma 6.16 implies that $\Pi(\boldsymbol{a})$ is positive definite for all $\boldsymbol{a} > \mathbf{0}$ and $\mathbf{u}^{(k)} \in \mathcal{S}(s(\mathbf{u})\mathsf{F}_\beta^{(k,n)})$. Following Proposition 6.6 which extends Proposition 3.4 to account the scaling factor, we obtain the following equations.

$$-\nabla\mathfrak{J}\left(\mathbf{u}^{(k)}\right)^{\intercal} (\mathbf{w}^{(k)} - \mathbf{u}^{(k)}) = \frac{\boldsymbol{\gamma}^{\intercal}(\mathbf{u}^{(k)} - \mathbf{w}^{(k)})}{s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}} = \frac{\boldsymbol{\gamma}^{\intercal}\left(s(\mathbf{u})\nabla\phi^{(k)}(\boldsymbol{\gamma}) - \mathbf{u}^{(k)}\right)}{s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}} \left(\frac{1-\beta}{\beta}\right)\,.$$

Using the fact that $\mathbf{u}^{(k)} + \mu\boldsymbol{e} = s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\nabla\phi^{(k)}(\boldsymbol{\gamma})$ with $\mu \geq 0$ (see Remark 3.2), it follows that $\mathbf{u}^{(k)} \leq s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\nabla\phi^{(k)}(\boldsymbol{\gamma})$ and $\boldsymbol{x}^{\intercal}\nabla\phi^{(k)}(\boldsymbol{x}) = \phi^{(k)}(\boldsymbol{x})$. Moreover $\mathsf{F}_\beta^{(k,n)} \geq 1/2$, $1 \geq \beta$ and

$\phi^{(k)}(\boldsymbol{x}) \geq \mathbb{E}[v]/k$ for all $\boldsymbol{x} \in \mathbb{R}^k_+$ by definition. Thus, we obtain

$$-\nabla \mathfrak{J}\left(\mathbf{u}^{(k)}\right)^\intercal \left(\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\right) \geq \frac{(1-\beta)\left(1 - \mathsf{F}_\beta^{(k,n)}\right)\mathbb{E}[v]}{2k\left(\mathsf{F}_\beta^{(k,n)}\right)^2}. \qquad \square$$

Proof of Lemma 6.18. We can invoke Proposition 6.3 to get $\text{Hess}(\mathfrak{J}(\boldsymbol{\Theta}^{(k)}))$ because $\boldsymbol{\Theta}^{(k)} \in \mathcal{S}(s(\mathbf{u})\mathsf{F}_\beta^{(k,n)})$ and Lemma 6.16 shows that $\Pi(\boldsymbol{a})$ is positive definite for all $\boldsymbol{a} > \mathbf{0}$.

$$\text{Hess}\left(\mathfrak{J}\left(\boldsymbol{\Theta}^{(k)}\right)\right) = \frac{1}{\left(s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\right)^2}\Pi(\boldsymbol{\lambda})^{-1}\left(I - \frac{\boldsymbol{e}\boldsymbol{e}^\intercal\Pi(\boldsymbol{\lambda})^{-1}}{\boldsymbol{e}^\intercal\Pi(\boldsymbol{\lambda})^{-1}\boldsymbol{e}}\right). \tag{6.26}$$

We prove this result in two steps. For simplicity, we use A to represent $\left(s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\right)^2 \text{Hess}\left(\mathfrak{J}\left(\boldsymbol{\Theta}^{(k)}\right)\right)$ and H to represent $\Pi(\boldsymbol{\lambda})^{-1}$.

Step 1. In this step, we find an upper bound for $\sup_{\|\mathbf{x}\|_2=1} \mathbf{x}^\intercal A\mathbf{x}$ which depends on $\min_i \lambda_i$. Using (6.26), we can rewrite $\mathbf{x}^\intercal A\mathbf{x}$ as follows:

$$\mathbf{x}^\intercal A\mathbf{x} = \frac{(\boldsymbol{e}^\intercal H\boldsymbol{e})\mathbf{x}^\intercal H\mathbf{x} - \mathbf{x}^\intercal H\boldsymbol{e}\boldsymbol{e}^\intercal H\mathbf{x}}{\boldsymbol{e}^\intercal H\boldsymbol{e}}.$$

Note that H is symmetric, thus eliminating the negative term $-\mathbf{x}^\intercal H\boldsymbol{e}\boldsymbol{e}^\intercal H\mathbf{x}$ yields the following bound,

$$\mathbf{x}^\intercal A\mathbf{x} \leq \frac{(\boldsymbol{e}^\intercal H\boldsymbol{e})\mathbf{x}^\intercal H\mathbf{x}}{\boldsymbol{e}^\intercal H\boldsymbol{e}} = \mathbf{x}^\intercal H\mathbf{x} = \varepsilon_H.$$

Here, ε_H is the largest eigenvalue of H, and in the last equality we use that $\|\mathbf{x}\|_2^2 = 1$. In Lemma 6.16, we further show that the minimum eigenvalue of $\Pi(\boldsymbol{\lambda})$ is $c_{\text{H}}(k)(\min_i \lambda_i)^{k+1}$. Therefore the largest eigenvalue of H is $\frac{1}{c_{\text{H}}(k)(\min_i \lambda_i)^{k+1}}$. Therefore, we have that

$$\frac{1}{2}\left(\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\right)^\intercal \text{Hess}\left(\mathfrak{J}\left(\boldsymbol{\Theta}^{(k)}\right)\right)\left(\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\right) \leq \frac{1}{2}\frac{1}{\left(s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\right)^2}\frac{\|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2^2}{c_{\text{H}}(k)(\min_i \lambda_i)^{k+1}}. \tag{6.27}$$

Step 2. In this step, we find an upper bound for $1/\min_i \lambda_i$ that depends on $\underline{u}_\beta^{(n)}$. Optimality of $\boldsymbol{\lambda}$ implies that $\Theta_i^{(k)} = s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}(\nabla\phi^{(k)}(\boldsymbol{\lambda}))_i - \mu$. If we knew $\boldsymbol{\Theta}^{(k)}$ is in $s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathcal{U}^{(k)}$, then it would follow μ is nonnegative and $\boldsymbol{\Theta}^{(k)} \leq s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\nabla\phi^{(k)}(\boldsymbol{\lambda})$. However $\boldsymbol{\Theta}^{(k)}$ could lie outside of $s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathcal{U}^{(k)}$. In that case, μ becomes negative. Thus, it could be the case that $\Theta_i^{(k)} > s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}(\nabla\phi^{(k)}(\boldsymbol{\lambda}))_i$ for all i. Despite this fact, we know $-\mu$ is less than $\|\mathbf{u}^{(k)} - \mathbf{w}^{(k)}\|_2$ because $\boldsymbol{\Theta}^{(k)}$ is between $\mathbf{u}^{(k)}$ and $\mathbf{w}^{(k)}$. Therefore, we have that

$$\Theta_i^{(k)} \leq s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}(\nabla\phi^{(k)}(\boldsymbol{\lambda}))_i + \|\mathbf{u}^{(k)} - \mathbf{w}^{(k)}\|_2 = s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathbb{E}[v_i\mathbf{1}\{\lambda_i v_i \geq \max_{j\neq i}\lambda_j v_j\}] + \|\mathbf{u}^{(k)} - \mathbf{w}^{(k)}\|_2\,.$$

Because the p.d.f. $f(\cdot)$ is bounded by $\bar{f}$ and the c.d.f. is always less than one, $\mathbb{E}[v_i\mathbf{1}\{\lambda_i v_i \geq \max_{j\neq i}\lambda_j v_j\}]$ can be bounded by $\lambda_i\bar{f}\mathbb{E}[v^2]/\max_{j\neq i}\lambda_j$. Moreover, we have that $\max_{j\neq i}\lambda_j \geq 1/k$, and $s(\mathbf{u})\mathsf{F}_\beta^{(k,n)} \leq 1$. Thus, it follows that $c_{\mathrm{S}}(k)\min_i\lambda_i \geq \min_i \Theta_i^{(k)} - \|\mathbf{u}^{(k)} - \mathbf{w}^{(k)}\|_2$.

We next derive a lower bound for $\min_i \Theta_i^{(k)}$ in terms of $\underline{u}_\beta^{(n)}$. As shown in the first step of Proposition 4.1, $(1-\beta)\left(K_1(k)\right)^2$ is less than $\left(\underline{u}_\beta^{(n)}\right)^2\beta$ and we know $\beta \geq \underline{\beta}$. Together with these observations, Lemma 6.14 implies that $\underline{u}_\beta^{(n)} \geq 3\|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2$. Recognizing the fact that $\boldsymbol{\Theta}^{(k)}$ is between $\mathbf{w}^{(k)}$ and $\mathbf{u}^{(k)}$, we have $\Theta_i^{(k)} \geq u_i^{(k)} - \|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2$ for all i. Because $u_i^{(k)} \geq \underline{u}_\beta^{(n)}$, it follows $\Theta_i^{(k)} \geq \underline{u}_\beta^{(n)} - \|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2$. Using the fact that $\underline{u}_\beta^{(n)} \geq 3\|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2$ we get $\Theta_i^{(k)} \geq \underline{u}_\beta^{(n)} - \underline{u}_\beta^{(n)}/3 = 2\underline{u}_\beta^{(n)}/3$ for all i, and so $\min_i \Theta_i \geq 2\underline{u}_\beta^{(n)}/3$. Combining these bounds we obtain that $\min_i \lambda_i \geq \underline{u}_\beta^{(n)}/(3c_{\mathrm{S}}(k))$.

Finally, using Lemma 6.14 to bound $\|\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\|_2^2$ and $\min_i \lambda_i \geq \underline{u}_\beta^{(n)}/(3c_{\mathrm{S}}(k))$, we can bound the right-hand side of (6.27), and have that

$$\frac{1}{2}\left(\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\right)^{\mathsf{T}}\mathrm{Hess}\left(\mathfrak{J}\left(\boldsymbol{\Theta}^{(k)}\right)\right)\left(\mathbf{w}^{(k)} - \mathbf{u}^{(k)}\right) \leq \frac{1}{2}\frac{1}{\left(\mathsf{F}_\beta^{(k,n)}\right)^2}\frac{1}{\left(\underline{u}_\beta^{(n)}\right)^{k+3}}\left(\frac{1-\beta}{\underline{\beta}}\right)^2\frac{\mathbb{E}[v]K_2(k)}{k}\,.$$

□

Proof of Lemma 6.19. The optimization problem in the statement of the lemma can be reformulated by the following change of variables and parameters. Let ${e_g}_i\sqrt{g_i} = \mathsf{v}_i$ and $a_i\sqrt{g_i} = \mathsf{w}_i$, and $x_i\sqrt{g_i} = \mathsf{z}_i$ for all i. Then it follows that

$$\min_{\|\boldsymbol{x}\|_g^2 \geq 1}\langle\boldsymbol{x}, \boldsymbol{e}_g\rangle_g^2 - \langle\boldsymbol{x}, \boldsymbol{a}\rangle_g^2 = \min_{\|\mathsf{z}\|_2^2 \geq 1}\langle\mathsf{z}, \mathsf{v}\rangle^2 - \langle\mathsf{z}, \mathsf{w}\rangle^2 = \min_{\|\mathsf{z}\|_2^2 \geq 1}\mathsf{z}^{\mathsf{T}}(\mathsf{v}\mathsf{v}^{\mathsf{T}} - \mathsf{w}\mathsf{w}^{\mathsf{T}})\mathsf{z}\,.$$

Following the min-max theorem, the optimal value of this problem is given by the minimum eigenvalue of the matrix $\mathsf{v}\mathsf{v}^{\mathsf{T}} - \mathsf{w}\mathsf{w}^{\mathsf{T}}$. To find the eigenvalues, we express v and w with two orthonormal unit vectors $\mathsf{v} = c_1 u_1$ and $\mathsf{w} = c_2 u_1 + c_3 u_2$, respectively. This implies that $c_1 = \|\mathsf{v}\|_2$ and $u_1 = \mathsf{v}/\|\mathsf{v}\|_2$, which further implies that $c_2 = u_1^\intercal \mathsf{w}$ and $c_3 = u_2^\intercal \mathsf{w}$. Finally, we have $u_2 = \dfrac{\mathsf{w} - c_2 u_1}{\|\mathsf{w} - c_2 u_1\|_2}$. In an orthonormal basis starting with u_1 and u_2, the matrix $\mathsf{v}\mathsf{v}^\intercal - \mathsf{w}\mathsf{w}^\intercal$ has first two rows and columns

$$\begin{bmatrix} c_1^2 - c_2^2 & -c_2 c_3 \\ -c_2 c_3 & -c_3^2 \end{bmatrix},$$

and everything else is 0. The minimum eigenvalue emerges as follows:

$$\begin{aligned} &\frac{1}{2}\left[c_1^2 - c_2^2 - c_3^2 - \sqrt{((c_3^2 + c_1^2 + c_2^2)^2 - 4(c_1 c_2)^2)}\right] \\ &= \frac{1}{2}\left[\|\mathsf{v}\|_2^2 - c_2^2 - c_3^2 - \sqrt{((\|\mathsf{v}\|_2^2 + c_3^2 + c_2^2)^2 - 4(\mathsf{v}^\intercal \mathsf{w})^2)}\right]. \end{aligned}$$

Therefore, we have $\|\mathsf{v}\|_2^2 = \|\boldsymbol{e}_g\|_g^2$ and $\mathsf{v}^\intercal \mathsf{w} = \langle \boldsymbol{a}, \boldsymbol{e}_g \rangle_g$. To complete the proof, we show that $c_2^2 + c_3^2 = \|\mathsf{w}\|_2^2 = \|\boldsymbol{a}\|_g^2$. In fact,

$$\begin{aligned} c_2^2 + c_3^2 &= \frac{(\mathsf{v}^\intercal \mathsf{w})^2}{\|\mathsf{v}\|_2^2} + \frac{\|\mathsf{v}\|_2^4 \|\mathsf{w}\|_2^4 + (\mathsf{w}^\intercal \mathsf{v})^4 - 2\|\mathsf{v}\|_2^2 \|\mathsf{w}\|_2^2 (\mathsf{w}^\intercal \mathsf{v})^2}{\left\| \|\mathsf{v}\|_2^2 \mathsf{w} - (\mathsf{v}^\intercal \mathsf{w})\mathsf{v} \right\|_2^2} \\ &= \frac{\langle \boldsymbol{e}_g, \boldsymbol{a} \rangle_g^2}{\|\boldsymbol{e}_g\|_g^2} + \frac{\|\boldsymbol{e}_g\|_g^4 \|\boldsymbol{a}\|_g^4 + \langle \boldsymbol{e}_g, \boldsymbol{a} \rangle_g^4 - 2\|\boldsymbol{e}_g\|_g^2 \|\boldsymbol{a}\|_g^2 \langle \boldsymbol{e}_g, \boldsymbol{a} \rangle_g^2}{\|\boldsymbol{e}_g\|_g^2 \left(\|\boldsymbol{e}_g\|_g^2 \|\boldsymbol{a}\|_g^2 - \langle \boldsymbol{e}_g, \boldsymbol{a} \rangle_g^2\right)} \\ &= \frac{\|\boldsymbol{e}_g\|_g^4 \|\boldsymbol{a}\|_g^4 - \|\boldsymbol{e}_g\|_g^2 \|\boldsymbol{a}\|_g^2 \langle \boldsymbol{e}_g, \boldsymbol{a} \rangle_g^2}{\|\boldsymbol{e}_g\|_g^2 \left(\|\boldsymbol{e}_g\|_g^2 \|\boldsymbol{a}\|_g^2 - \langle \boldsymbol{e}_g, \boldsymbol{a} \rangle_g^2\right)} = \frac{\|\boldsymbol{e}_g\|_g^2 \|\boldsymbol{a}\|_g^2 \left(\|\boldsymbol{e}_g\|_g^2 \|\boldsymbol{a}\|_g^2 - \langle \boldsymbol{e}_g, \boldsymbol{a} \rangle_g^2\right)}{\|\boldsymbol{e}_g\|_g^2 \left(\|\boldsymbol{e}_g\|_g^2 \|\boldsymbol{a}\|_g^2 - \langle \boldsymbol{e}_g, \boldsymbol{a} \rangle_g^2\right)} = \|\boldsymbol{a}\|_g^2. \end{aligned}$$

□

6.8 Computing the Optimal Achievable Set $\mathcal{U}_\beta$

In this section, we discuss how to numerically characterize the optimal achievable set $\mathcal{U}_\beta$. Because our approach is topological in nature, we need to endow the space of mechanism with a topology. We assume that $\mathbf{p}$ and $\mathbf{w}$ lie in the space $L^\infty \triangleq L^\infty([0,\bar{v}]^n, (\mathbb{R}^n, \|\cdot\|_1))$ of essentially bounded vector functions (see Appendix 6.8.2 for details). This is required to show that the recursive approach delineated in Proposition 6.5 and Corollary 6.2 converges to the set of achievable utilities $\mathcal{U}_\beta$. Since the space L^∞ is an equivalence class of functions which agree almost everywhere, all point-wise inequalities are understood to hold almost everywhere (a.e.). To simplify the exposition we dropped the "almost everywhere" qualifier in the rest of the appendix.

First, we show that the set $\mathcal{U}_\beta$ can be obtained through repeatedly applying operator B_β.

Proposition 6.5. *Let $\mathcal{U}^0 = \left[0, \mathbb{E}[v]\right]^n$ and $\mathcal{U}^k = B_\beta(\mathcal{U}^{k-1})$ for all $k \geq 1$. Then, we have $\mathcal{U}^k \subseteq \mathcal{U}^{k-1}$ and $\lim_{k\to\infty} \mathcal{U}^k = \mathcal{U}_\beta$.*

The set $\mathcal{U}^0$ can be understood as the achievable utilities of a relaxation in which the resource can be simultaneously allocated to all agents in every period. It is not hard to see that the operator B_β is monotone, and $B_\beta(\mathcal{U}_0) \subseteq \mathcal{U}_0$. This implies that the sequence $\mathcal{U}^k$ is monotone and converges to a set $\mathcal{U}^\infty$. We prove the result by showing that $\mathcal{U}_\beta \subseteq \mathcal{U}^\infty$ and $\mathcal{U}^\infty \subseteq \mathcal{U}_\beta$. For the first part, notice that because $\mathcal{U}_\beta \subseteq \mathcal{U}^0$ and $\mathcal{U}_\beta$ is a fixed-point of B_β, monotonicity of the operator B_β implies that $\mathcal{U}_\beta \subseteq \mathcal{U}^k$ for all k, and therefore its limit. For the second part, it suffices to show, following Proposition 3.1, that $\mathcal{U}^\infty$ is self-generating. This last step of the proof is technical and relies on the topological properties of the space of stage mechanisms. Specifically, because $\mathcal{U}^\infty \subseteq \mathcal{U}^k$, for each $\mathbf{u} \in \mathcal{U}^\infty$ we can construct a sequence of stage mechanisms $(\mathbf{p}^k, \mathbf{w}^k)$ satisfying (IC), (FA), (PK($\mathbf{u}$)), and $\mathbf{w}^k(\mathbf{v}) \in \mathcal{U}^{k-1}$ for all $\mathbf{v}$. Because the space of stage mechanisms satisfying these conditions is compact (under an appropriate topology), we can construct a stage mechanism $(\mathbf{p}, \mathbf{w})$ satisfying (IC), (FA), (PK($\mathbf{u}$)), and $\mathbf{w}(\mathbf{v}) \in \mathcal{U}^\infty$ for all $\mathbf{v}$ by taking limits.

6.8.1 Support Function Representation

The iterative procedure described in Proposition 6.5 is of theoretical interest, but does not directly yield a straightforward numerical procedure due to the difficulty of determining set $B_\beta(\mathcal{A})$ from a set $\mathcal{A}$. In this section, we present an equivalent, support function representation of the set $\mathcal{U}_\beta$. This support function representation is not only amenable to numerical procedures, but also provides a foundation for the proof of our asymptotic results.

First, we present some basic properties of the set $\mathcal{U}_\beta$, which enable us to obtain its support function representation. The next result shows that the set of achievable utilities is convex and compact. The result follows from Proposition 6.5, and the fact that $\mathcal{U}^0$ is convex and compact, and that the operator B_β preserves convexity and compactness.

Lemma 6.20. *The set of achievable utilities, $\mathcal{U}_\beta$, is convex and compact, and satisfies $\mathcal{U}_\beta = \mathbb{R}^n_+ \cap$ hyp$(\mathcal{U}_\beta)$ where hyp$(\mathcal{U}_\beta) = \{\mathbf{u} \in \mathbb{R}^n \mid \mathbf{u} \leq \bar{\mathbf{u}}, \exists \bar{\mathbf{u}} \in \mathcal{U}_\beta\}$.*

Similar to (3.5), for any $\boldsymbol{\alpha} \in \mathbb{R}^n_+$ with $\|\boldsymbol{\alpha}\|_1 = 1$, define the support function of the convex perfect information achievable set $\mathcal{U}_\beta$ as follows:

$$\phi_\beta(\boldsymbol{\alpha}) \triangleq \sup_{\mathbf{u} \in \mathcal{U}_\beta} \boldsymbol{\alpha}^\intercal \mathbf{u} \,. \tag{6.28}$$

Furthermore, we focus on $\boldsymbol{\alpha} \geq 0$, because all nonnegative Pareto dominated points lie in $\mathcal{U}_\beta$, following Lemma 6.20. Proposition 6.5 allows us to obtain a recursive representation of the support function $\phi_\beta(\boldsymbol{\alpha})$. To that end, for a generic function $\psi(\boldsymbol{\alpha})$, define operator T_β as:

$$\begin{aligned} [T_\beta \psi](\boldsymbol{\alpha}) = \sup_{\mathbf{u},\mathbf{p},\mathbf{w}} \; & \boldsymbol{\alpha}^\intercal \mathbf{u} \\ \text{s.t.} \; & \text{(IC)}, \text{(FA)}, \text{(PK}(\mathbf{u})) \\ & \tilde{\boldsymbol{\alpha}}^\intercal \mathbf{w}(\mathbf{v}) \leq \psi(\tilde{\boldsymbol{\alpha}}) \,, \quad \forall \mathbf{v} \in [0, \bar{v}]^n, \; \tilde{\boldsymbol{\alpha}} \in \mathbb{R}^n_+ \,. \end{aligned} \tag{6.29}$$

The definition of operator T_β involves an infinite-dimensional linear optimization problem. We can solve it approximately by considering constraints (6.29) only for $\mathbf{v}$ and $\tilde{\boldsymbol{\alpha}}$ on a grid, such that $\tilde{\boldsymbol{\alpha}} \geq 0$ and $\|\tilde{\boldsymbol{\alpha}}\|_1 = 1$. Furthermore, we can define the operator of recursively implementing operator T_β as $T^k_\beta \psi \triangleq T_\beta(T^{k-1}_\beta \psi)$. Our T_β operator resembles the techniques developed by Judd et al. (2003) to compute equilibrium payoffs of infinitely repeated games. However, the specific constraints in Judd et al. (2003) are different from ours as our set operator B_β is different from the set operator B in Abreu et al. (1990). Specifically, Judd et al. (2003) consider the equilibrium payoffs of infinitely repeated games with perfect monitoring whose action spaces are finite.

The following result states that ϕ_β can be obtained by repeatedly applying operator T_β. The proof, again, follows from Proposition 6.5, and the fact that convergence of convex and compact sets is equivalent to convergence of support functions.

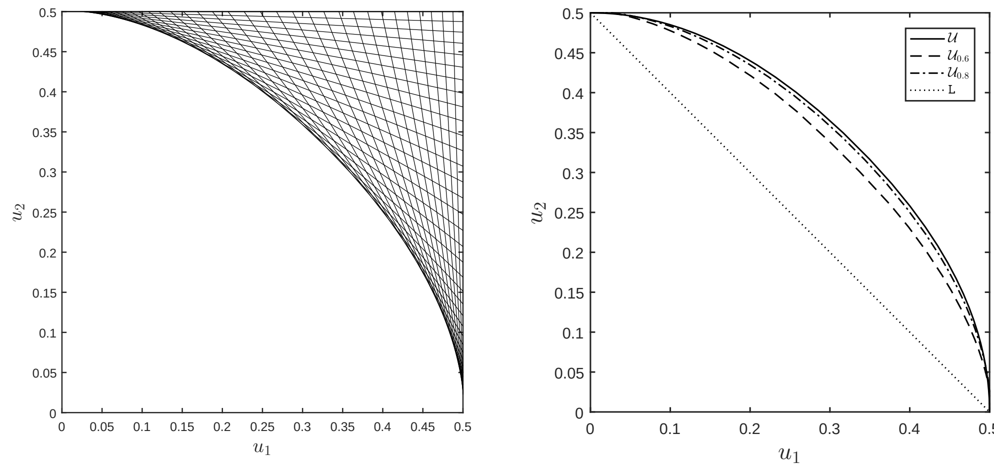

(a) **(b)**

Figure 6.1: Illustration of achievable sets in a two-agent case.

Corollary 6.2. *The support function ϕ_β is a fixed point of T_β, i.e., $\phi_\beta = T_\beta \phi_\beta$. Furthermore, we have $\phi_\beta(\boldsymbol{\alpha}) = \lim_{k\to\infty} [T_\beta^k \phi^0](\boldsymbol{\alpha})$ for all $\boldsymbol{\alpha}$, starting from $\phi^0(\boldsymbol{\alpha}) = \mathbb{E}[v]$.*

Figure 6.1a demonstrates the set of achievable utilities, $\mathcal{U}_\beta$ with two agents. The lines in the graph are the supporting hyperplanes, which are computed from the iterative procedure described in Corollary 6.2. The blank area outlined by the supporting hyperplanes is the set $\mathcal{U}_\beta$. Numerically, one can start the iterative algorithm of Corollary 6.2 from ϕ, instead of ϕ^0, which leads to faster convergence. In fact, Figure 6.1a is generated by starting from ϕ. Figure 6.1b compares the boundaries of sets $\mathcal{U}$ and $\mathcal{U}_\beta$ for different β values. The set $\mathcal{U}_\beta$ is monotonically increasing as β increases, and, therefore, the planner is able to achieve higher total utility. Furthermore, it appears that as the time discount factor approaches one, the set $\mathcal{U}_\beta$ approaches the full information achievable set $\mathcal{U}$.

Although the procedure described in Corollary 6.2 generates the achievable set $\mathcal{U}_\beta$, it does not directly provide PBIC mechanisms achieving each state in the set.

6.8.2 Proof of Proposition 6.5

We provide some preliminary results before proving Proposition 6.5. First, we endow the space of stage mechanism with a topology and prove that the set of mechanism satisfying our constraints is compact. Second, we show that the operator B_β is monotone, and preserves convexity and compactness.

We assume that the functions of the stage mechanism $\mathbf{p}$ and $\mathbf{w}$ lie in the space

$$L^\infty \triangleq L^\infty([0,\bar{v}]^n, (\mathbb{R}^n, \|\cdot\|_1))$$

of essentially bounded vector functions from $[0,\bar{v}]^n$ to $(\mathbb{R}^n, \|\cdot\|_1)$, i.e, a measurable vector function $\boldsymbol{\rho} : [0,\bar{v}]^n \to \mathbb{R}^n$ lies in L^∞ if $\operatorname{ess\,sup}_{\mathbf{v}} \|\boldsymbol{\rho}(\mathbf{v})\|_1 < \infty$ where $\|\boldsymbol{\rho}(\mathbf{v})\|_1 = \sum_{i=1}^n |\rho_i(\mathbf{v})|$. We denote by $L^1 \triangleq L^1([0,\bar{v}]^n, (\mathbb{R}^n, \|\cdot\|_\infty))$ the pre-dual of L^∞, i.e., a measurable vector function $\boldsymbol{\psi} : [0,\bar{v}]^n \to \mathbb{R}^n$ lies in L^1 if $\int_{[0,\bar{v}]^n} \|\boldsymbol{\psi}(\mathbf{v})\|_\infty \mathrm{d}\mathbf{v} < \infty$ where $\|\boldsymbol{\rho}(\mathbf{v})\|_\infty = \max_{i=1}^n |\rho_i(\mathbf{v})|$. Since the spaces L^∞ and L^1 are equivalence classes of functions which agree almost everywhere, all point-wise inequalities are understood to hold almost everywhere (a.e.). To simplify the exposition we drop the "almost everywhere" qualifier.

We endow L^∞ with the weak-* topology, the coarsest topology under which every element $\boldsymbol{\psi} \in L^1$ corresponds to a continuous (linear) map. A sequence $\boldsymbol{\rho}^m \in L^\infty$ converges to $\boldsymbol{\rho}$ in the weak-* topology if

$$\int \boldsymbol{\rho}^m(\mathbf{v}) \cdot \boldsymbol{\psi}(\mathbf{v}) \mathrm{d}\mathbf{v} \to \int \boldsymbol{\rho}(\mathbf{v}) \cdot \boldsymbol{\psi}(\mathbf{v}) \mathrm{d}\mathbf{v} \quad \forall \boldsymbol{\psi} \in L^1 ,$$

where $\boldsymbol{\rho}(\mathbf{v}) \cdot \boldsymbol{\psi}(\mathbf{v}) = \sum_{i=1}^n \rho_i(\mathbf{v})\psi_i(\mathbf{v})$ denotes the standard inner product in $\mathbb{R}^n$. In this case, one writes $\boldsymbol{\rho}^m \rightharpoonup \boldsymbol{\rho}$ as $m \to \infty$.

Lemma 6.21. *The stage mechanisms satisfy the following properties:*

1. *Let $\mathcal{A} \subset \mathbb{R}^n$ be a closed and convex set. The set $\{\mathbf{w} \in L^\infty : \mathbf{w}(\mathbf{v}) \in \mathcal{A}$ for almost all $\mathbf{v} \in [0,\bar{v}]^n\}$ is weak-* closed.*

2. *Let $\mathcal{A} \subset \mathbb{R}^n$ be a compact and convex set. The set $\mathcal{M}_{\mathcal{A}} = \{(\mathbf{p},\mathbf{w}) \in L^\infty \times L^\infty : \sum_{i=1}^n p_i(\mathbf{v}) \leq 1,\ \mathbf{p}(\mathbf{v}) \geq \mathbf{0},$ and $\mathbf{w}(\mathbf{v}) \in \mathcal{A}$ for almost all $\mathbf{v} \in [0,\bar{v}]^n\}$ is weak-* compact.*

3. *The set $\{(\mathbf{p},\mathbf{w}) \in L^\infty \times L^\infty : (\mathbf{p},\mathbf{w})$ satisfy (IC)$\}$ is weak-* closed.*

4. *The function $E : L^\infty \times L^\infty \to \mathbb{R}^n$ given by $E(\mathbf{p},\mathbf{w}) = (\mathbb{E}_{\mathbf{v}}[(1-\beta)v_i p_i(\mathbf{v}) + \beta w_i(\mathbf{v})])_{i=1}^n$ is weak-* continuous.*

We next provide some properties on the operator B_β defined in (3.3).

Lemma 6.22. *The operator $B_\beta : 2^{\mathbb{R}^n} \to 2^{\mathbb{R}^n}$ satisfies the following properties:*

1. *(B_β is monotone) If $\mathcal{X} \subseteq \mathcal{Y} \subset \mathbb{R}^n$ then $B_\beta(\mathcal{X}) \subseteq B_\beta(\mathcal{Y})$.*

2. *(B_β preserves convexity) If $\mathcal{A} \subset \mathbb{R}^n$ is convex, then $B_\beta(A)$ is convex.*

3. *(B_β preserves compactness) If $\mathcal{A} \subset \mathbb{R}^n$ is compact and convex, then $B_\beta(A)$ is compact.*

We defer the proofs of these results to the end of the subsection. We are now in position to prove the main result.

Proof of Proposition 6.5. Let $\mathcal{U}^\infty = \lim_{m\to\infty} \mathcal{U}^m$ with $\mathcal{U}^0 = [0, \mathbb{E}[v]]^n$ and $\mathcal{U}^m = B_\beta(\mathcal{U}^{m-1})$ for all $m \geq 1$. We prove this result in two steps. We first prove that the limit $\mathcal{U}^\infty$ exists and that $\mathcal{U}_\beta \subseteq \mathcal{U}^\infty$. We then show that $\mathcal{U}^\infty \subseteq \mathcal{U}_\beta$.

Step 1. We first claim that (i) $B_\beta(\mathcal{U}^0) \subseteq \mathcal{U}^0$ and (ii) $\mathcal{U}_\beta \subseteq \mathcal{U}^0$. Claim (i), together with Lemma 6.22, item 1, implies that $\mathcal{U}^m \subseteq \mathcal{U}^{m-1}$, and thus $\mathcal{U}^\infty = \lim_{m\to\infty} \mathcal{U}^m$ exists. Moreover, because $\mathcal{U}_\beta$ is self-generating we obtain by (ii) that $\mathcal{U}_\beta \subseteq \mathcal{U}^m$ for all m. This implies that $\mathcal{U}_\beta \subseteq \mathcal{U}^\infty$. We next prove the claims.

We show that $B_\beta(\mathcal{U}^0) \subseteq \mathcal{U}^0$ by proving that for all $\mathbf{u} \in B_\beta(\mathcal{U}^0)$ we have $\mathbf{u} \in \mathcal{U}^0$. Because $\mathbf{u} \in B_\beta(\mathcal{U}^0)$, there exist functions $(\mathbf{p}, \mathbf{w})$ satisfying (IC), (PK($\mathbf{u}$)) and (FA) such that $\mathbf{w}(\mathbf{v}) \in \mathcal{U}^0$ for all $\mathbf{v}$. Since $\mathbf{p}$ satisfies (FA), it follows that $\mathbf{u}^1 := (\mathbb{E}[v_i p_i(\mathbf{v})])_{i=1}^n$ satisfies $\mathbf{u}^1 \in \mathcal{U}^0$. Note that $\mathcal{U}^0$ is a convex set. Thus, the fact that $\mathbf{w}(\mathbf{v}) \in \mathcal{U}^0$ implies that $\mathbf{u}^2 := (\mathbb{E}[w_i(\mathbf{v})])_{i=1}^n$ satisfies $\mathbf{u}^2 \in \mathcal{U}^0$. Finally, (PK($\mathbf{u}$)) implies that $\mathbf{u} = (1-\beta)\mathbf{u}^1 + \beta\mathbf{u}^2$ and hence $\mathbf{u} \in \mathcal{U}^0$ because $\mathbf{u}$ is a convex combination of two points in $\mathcal{U}^0$.

We argue that $\mathcal{U}_\beta \subseteq \mathcal{U}^0$. Fix a point $\mathbf{u} \in \mathcal{U}_\beta$. Recognizing the definition of $\mathcal{U}_\beta$ in (3.2) it follows that $u_i = V_i(\boldsymbol{\pi}, \mathbf{I})$ for some PBIC mechanism $\boldsymbol{\pi}$. Because mechanism $\boldsymbol{\pi}$ satisfies $\boldsymbol{\pi}_t \in \mathcal{P}$, we conclude that $u_i \in [0, \mathbb{E}[v]]$ for all i. The claim follows.

Step 2. In this step, we prove the other direction, i.e., $\mathcal{U}^\infty \subseteq \mathcal{U}_\beta$. By Proposition 3.1, it is sufficient to prove that $\mathcal{U}^\infty \subseteq B_\beta(\mathcal{U}^\infty)$. Fix a point $\mathbf{u} \in \mathcal{U}^\infty$. In order to prove that $\mathbf{u} \in B_\beta(\mathcal{U}^\infty)$, we need to show that there exists a pair of functions $(\mathbf{p}, \mathbf{w})$ satisfying (IC), (FA), (PK($\mathbf{u}$)) and $\mathbf{w}(\mathbf{v}) \in \mathcal{U}^\infty$ for all $\mathbf{v}$.

As shown in the first step, $\mathcal{U}^\infty = \bigcap_{m\geq 1} \mathcal{U}^m$. Therefore, $\mathbf{u} \in \mathcal{U}^m$ for all m, and there exists a sequence of pairs of functions $(\mathbf{p}^m, \mathbf{w}^m)$ which satisfy (IC), (FA), (PK($\mathbf{u}$)) and $\mathbf{w}^m(\mathbf{v}) \in \mathcal{U}^{m-1} \subseteq \mathcal{U}_0$.

Because $\mathcal{U}_0$ is convex and compact, we obtain by Lemma 6.21, item 2 that $(\mathbf{p}^m, \mathbf{w}^m)$ lie in the weak-* compact set $\mathcal{M}_{\mathcal{U}_0}$ with the set $\mathcal{M}_{\mathcal{U}_0}$ defined in Lemma 6.21, item 2. By passing to a subsequence if necessary we obtain that $\mathbf{p}^m$ and $\mathbf{w}^m$ weak-* converge to some $(\mathbf{p}, \mathbf{w}) \in \mathcal{M}_{\mathcal{U}_0}$.

By Lemma 6.21, item 3 the set of incentive compatible stage mechanisms is weak-* closed, and thus $(\mathbf{p}, \mathbf{w})$ is incentive compatible. For a fixed m, we have $\mathbf{w}^j(\mathbf{v}) \in \mathcal{U}^m$ for all $j > m$ because the sequence $\mathcal{U}^m$ is non-increasing. Because $\mathcal{U}^0$ is convex and compact, we obtain that $\mathcal{U}^m$ is convex and compact because the operator B_β preserves convexity and compactness from Lemma 6.22, items 2 and 3. Because the set $\mathcal{U}^m$ is closed and convex, we obtain by Lemma 6.21, item 1 that weak limit verifies $\mathbf{w}(\mathbf{v}) \in \mathcal{U}^m$. Therefore, $\mathbf{w}(\mathbf{v}) \in \bigcap_{m \geq 1} \mathcal{U}^m = \mathcal{U}^\infty$.

We conclude by showing that $(\mathbf{p}, \mathbf{w})$ satisfies the promise keeping constraint with $\mathbf{u}$. For each m we have that the promised keeping constraint is equivalently given by $\mathbf{u} = E(\mathbf{p}^m, \mathbf{w}^m)$ with the function E defined in Lemma 6.21, item 4. Because the function E is weak-* continuous, we obtain by taking limits that $\mathbf{u} = E(\mathbf{p}, \mathbf{w})$. Thus, $\mathbf{u} \in B_\beta(\mathcal{U}^\infty)$ since the functions $(\mathbf{p}, \mathbf{w})$ satisfy (IC), (FA), (PK($\mathbf{u}$)) and $\mathbf{w}(\mathbf{v}) \in \mathcal{U}^\infty$ for all $\mathbf{v}$. □

Proof of Lemma 6.21

Proof. We prove each item at a time.

Item 1. Let $\mathcal{A} \subset \mathbb{R}^n$ be a closed and convex set. We need to show that the set $\mathcal{C} = \{\mathbf{w} \in L^\infty : \mathbf{w}(\mathbf{v}) \in \mathcal{A}$ for almost all $\mathbf{v} \in [0, \bar{v}]^n\}$ is weak-* closed. Convexity of $\mathcal{A}$ implies that $\mathcal{C}$ is convex. Because $\mathcal{A}$ is closed, we obtain that the set $\mathcal{C}$ is (strongly) closed. Because $\mathcal{C}$ is convex and (strongly) closed, we obtain that $\mathcal{C}$ is weak-* closed by the Hahn-Banach separation theorem (see Aliprantis and Border, 2006, Theorem 5.98 in p. 214).

Item 2. Let $\mathcal{A} \subset \mathbb{R}^n$ be a compact and convex set. We need to show that $\mathcal{M}_\mathcal{A} = \{(\mathbf{p}, \mathbf{w}) \in L^\infty \times L^\infty : \sum_{i=1}^n p_i(\mathbf{v}) \leq 1,\ \mathbf{p}(\mathbf{v}) \geq \mathbf{0},$ and $\mathbf{w}(\mathbf{v}) \in \mathcal{A}$ for almost all $\mathbf{v} \in [0, \bar{v}]^n\}$ is weak-* compact. We write $\mathcal{M}_\mathcal{A} = \mathcal{M}_\mathbf{p} \times \mathcal{M}_\mathbf{w}$, where $\mathcal{M}_\mathbf{p} = \{\mathbf{p} \in L^\infty : \sum_{i=1}^n p_i(\mathbf{v}) \leq 1,\ \mathbf{p}(\mathbf{v}) \geq \mathbf{0}$ for almost all $\mathbf{v} \in [0, \bar{v}]^n\}$ and $\mathcal{M}_\mathbf{w} = \{\mathbf{w} \in L^\infty : \mathbf{w}(\mathbf{v}) \in \mathcal{A}$ for almost all $\mathbf{v} \in [0, \bar{v}]^n\}$.

Because the pre-dual L^1 is normed, following Alaoglu's compactness theorem (see Aliprantis and Border, 2006, Theorem 5.105 in p. 218), we obtain that the closed unit ball in L^∞ given by $\{\boldsymbol{\rho} \in L^\infty : \operatorname{ess\,sup}_\mathbf{v} \|\boldsymbol{\rho}(\mathbf{v})\|_1 \leq 1\}$ is weak-* compact. Because $\mathcal{A} \subset \mathbb{R}^n$ is a compact and convex set, item 1 implies that $\mathcal{M}_{\mathbf{w}}$ is weak-* compact because the intersection of a weak-* compact set and a weak-* closed set is weak-* compact. A similar argument shows that $\mathcal{M}_{\mathbf{p}}$ is weak-* compact. The result follows by the Tychonoff product theorem (see Aliprantis and Border, 2006, Theorem 2.61 in p. 52) because the product space $\mathcal{M}_{\mathcal{A}} = \mathcal{M}_{\mathbf{p}} \times \mathcal{M}_{\mathbf{w}}$ is also weak-* compact.

Item 3. Let us assume that the interim functions of the stage mechanism for a single agent $W_i(v_i)$ and $P_i(v_i)$ lie in the space $L^{\infty,1} \triangleq L^{\infty,1}([0,\bar{v}], (\mathbb{R}, |\cdot|))$ of essentially bounded vector functions from $[0,\bar{v}]$ to $(\mathbb{R}, |\cdot|)$, i.e., a measurable function $\rho : [0,\bar{v}] \to \mathbb{R}$ lies in $L^{\infty,1}$ if $\operatorname*{ess\,sup}_{v\in[0,\bar{v}]} |\rho(v)| < \infty$. Let $T : L^{\infty} \to (L^{\infty,1})^n$ be the interim operator that projects an element $\boldsymbol{\rho} \in L^{\infty}$ to $(T\boldsymbol{\rho})_i(v_i)$ for each $i = 1, \ldots, n$ by taking expectations over the other agents' values. This operator is given by

$$(T\boldsymbol{\rho})_i(v_i) = \int_{[0,\bar{v}]^{n-1}} \rho_i(v_i, \mathbf{v}_{-i}) \prod_{j\neq i} f_j(v_j) \mathrm{d}\mathbf{v}_{-i} \,. \tag{6.30}$$

We claim that T is weak-* continuous. Define the space of interim incentive compatible mechanisms for a single agent as follows:

$$\mathcal{C}^1 = \{(\rho,\omega) \in L^{\infty,1} \times L^{\infty,1} : (1-\beta)v\rho(v) + \beta\omega(v) \geq (1-\beta)v\rho(v') + \beta\omega(v') \text{ for almost all } v, v' \in [0,\bar{v}]\} \,.$$

Note that $\mathcal{C}^1$ is convex and (strongly) closed. This follows because the constraints in $\mathcal{C}^1$ correspond to closed halfspaces, so the set of points satisfying these constraints is closed and convex. This implies that $\mathcal{C}^1$ is weak-* closed (see the proof of Lemma 6.21, Item 1).

Let $\mathcal{C} = \{(\mathbf{p},\mathbf{w}) \in L^{\infty} \times L^{\infty} : (\mathbf{p},\mathbf{w}) \text{ satisfy (IC)}\}$. Then, we have that $\mathcal{C} = (T \times T)^{-1}((\mathcal{C}^1)^n)$. Because T is weak-* continuous, then the cartesian product of functions $T \times T$ is weak-* continuous in the product topology. Because $T \times T$ is weak-* continuous and the cartesian product $(\mathcal{C}^1)^n$ is weak-* closed, $\mathcal{C}$ is weak-* closed (see Aliprantis and Border, 2006, Theorem 2.27 in p. 36). The result follows.

We now prove the claim that the interim operator $T : L^{\infty} \to (L^{\infty,1})^n$ is weak-* continuous. Following Proposition 4 in Anderson and Nash (1987, p. 37), it is sufficient to show that the adjoint T^* of the operator T maps the pre-dual of $(L^{\infty,1})^n$ into the pre-dual of L^{∞}. Let $L^{1,1} \triangleq L^{1,1}([0,\bar{v}], (\mathbb{R}, |\cdot|))$ denote the pre-dual of $L^{\infty,1}$, i.e., a measurable function $\psi : [0,\bar{v}] \to \mathbb{R}$ lies in $L^{1,1}$ if $\int_{[0,\bar{v}]} |\psi(v)| \mathrm{d}v < \infty$. The pre-dual of $(L^{\infty,1})^n$ is given by the space $(L^{1,1})^n$ with the norm of a typical element $(\psi_i(v_i))_{i=1}^n \in (L^{1,1})^n$ given by $\max_{i=1,\ldots,n} \int_{[0,\bar{v}]} |\psi_i(v_i)| \mathrm{d}v_i$. Specifically, we need to show that $T^*((\psi_i)_{i=1}^n) \in L^1$ where

$(\psi_i)_{i=1}^n \in (L^{1,1})^n$.

We start with determining the adjoint $T^* : (L^{1,1})^n \to L^1$. By definition, the adjoint satisfies the following property for all $\boldsymbol{\rho} \in L^\infty$ and $(\psi_i)_{i=1}^n \in (L^{1,1})^n$:

$$\sum_{i=1}^n \int_{[0,\bar{v}]} (T\boldsymbol{\rho})_i(v_i)\psi_i(v_i)\mathrm{d}v_i = \int_{[0,\bar{v}]^n} T^*((\psi_i)_{i=1}^n)(\mathbf{v}) \cdot \boldsymbol{\rho}(\mathbf{v})\mathrm{d}\mathbf{v}\,. \tag{6.31}$$

We need to find T^* satisfying the above condition. We claim that the adjoint is given by

$$T^*((\psi_i)_{i=1}^n)(\mathbf{v}) = \left(\psi_i(v_i)\prod_{j\neq i} f_j(v_j)\right)_{i=1}^n\,. \tag{6.32}$$

This follows from equations (6.31) and (6.30) because

$$\begin{aligned}\sum_{i=1}^n \int_{[0,\bar{v}]} (T\boldsymbol{\rho})_i(v_i)\psi_i(v_i)\mathrm{d}v_i &= \sum_{i=1}^n \int_{[0,\bar{v}]} \int_{[0,\bar{v}]^{n-1}} \rho_i(v_i,\mathbf{v}_{-i}) \prod_{j\neq i} f_j(v_j)\mathrm{d}\mathbf{v}_{-i}\psi_i(v_i)\mathrm{d}v_i \\ &= \int_{[0,\bar{v}]^n} \sum_{i=1}^n \left(\psi_i(v_i)\prod_{j\neq i} f_j(v_j)\right)\rho_i(\mathbf{v})\mathrm{d}\mathbf{v}\,,\end{aligned}$$

where we used Fubini's theorem to change the order of integrals, because these functions are integrable; and exchanged summation with integration, because the sum is finite.

We next show that $T^*((\psi_i)_{i=1}^n) \in L^1$ for $(\psi_i)_{i=1}^n \in (L^{1,1})^n$.

Note that $(\psi_i)_{i=1}^n \in (L^{1,1})^n$ if $\max_{i=1,\dots,n} \int_{[0,\bar{v}]} |\psi_i(v_i)|\mathrm{d}v_i < \infty$. We have

$$\begin{aligned}\int_{[0,\bar{v}]^n} \|T^*((\psi_i)_{i=1}^n)(\mathbf{v})\|_\infty \mathrm{d}\mathbf{v} &= \int_{[0,\bar{v}]^n} \max_{i=1,\dots,n} \left|\psi_i(v_i)\prod_{j\neq i} f_j(v_j)\right| \mathrm{d}\mathbf{v} \\ &\le \int_{[0,\bar{v}]^n} \sum_{i=1}^n |\psi_i(v_i)| \prod_{j\neq i} f_j(v_j)\mathrm{d}\mathbf{v} = \sum_{i=1}^n \int_{[0,\bar{v}]} |\psi_i(v_i)|\,\mathrm{d}v_i \\ &\le n \max_{i=1,\dots,n} \int_{[0,\bar{v}]} |\psi_i(v_i)|\mathrm{d}v_i < \infty\,,\end{aligned}$$

where the first equation follows from (6.32); the first inequality because $\|\mathbf{x}\|_\infty \le \|\mathbf{x}\|_1$ for $\mathbf{x} \in \mathbb{R}^n$; the second equation from exchanging summation with integration because the sum is finite, using Tonelli's theorem because the functions are nonnegative and using that $f(\cdot)$ is a density function; and the third inequality because $\|\mathbf{x}\|_1 \le n\|\mathbf{x}\|_\infty$ for $\mathbf{x} \in \mathbb{R}^n$. Finally, since $(\psi_i)_{i=1}^n \in (L^{1,1})^n$, the statement follows.

Item 4. We need to show that the function $E(\mathbf{p}, \mathbf{w}) = (\mathbb{E}_{\mathbf{v}}[(1-\beta)v_i p_i(\mathbf{v}) + \beta w_i(\mathbf{v})])_{i=1}^n$ is weak-* continuous. Consider $\mathbf{p}^m \rightharpoonup \mathbf{p}$ and $\mathbf{w}^m \rightharpoonup \mathbf{w}$ as $m \to \infty$. For a fixed i, consider the function $\boldsymbol{\psi}^{\mathbf{p}} \in L^1$ given by $\psi_i^{\mathbf{p}}(\mathbf{v}) = v_i \prod_{k=1}^n f_k(v_k)$ and $\psi_j^{\mathbf{p}}(\mathbf{v}) = 0$ for all $j \neq i$, and the function $\boldsymbol{\psi}^{\mathbf{w}} \in L^1$ given by $\psi_i^{\mathbf{w}}(\mathbf{v}) = \prod_{k=1}^n f_k(v_k)$ and $\psi_j^{\mathbf{w}}(\mathbf{v}) = 0$ for all $j \neq i$. By definition, we know that

$$\mathbb{E}_{\mathbf{v}}[v_i p_i^m(\mathbf{v})] \to \mathbb{E}_{\mathbf{v}}[v_i p_i(\mathbf{v})] \text{ and } \mathbb{E}[w_i^m(\mathbf{v})] \to \mathbb{E}[w_i(\mathbf{v})]\,.$$

This implies that $E(\mathbf{p}^m, \mathbf{w}^m) \to E(\mathbf{p}, \mathbf{w})$ as $m \to \infty$. □

Proof of Lemma 6.22

Proof. We prove each item at a time.

Item 1 (B_β is monotone). Fix a point $\mathbf{u} \in B_\beta(\mathcal{X})$ and let $(\mathbf{p}, \mathbf{w})$ be the functions corresponding to $\mathbf{u}$ satisfying (IC), (PK($\mathbf{u}$)), (FA), and $\mathbf{w}(\mathbf{v}) \in \mathcal{X}$ for all $\mathbf{v}$. Since $\mathcal{X} \subseteq \mathcal{Y}$, it follows that $\mathbf{w}(\mathbf{v}) \in \mathcal{Y}$ for all $\mathbf{v}$. Thus, the existence of the functions $(\mathbf{p}, \mathbf{w})$ satisfying the previous requirements implies that $\mathbf{u} \in B_\beta(\mathcal{Y})$.

Item 2 (B_β preserves convexity). Fix points $\mathbf{u}^1, \mathbf{u}^2 \in B_\beta(\mathcal{A})$. We need to show that $\mathbf{u} = \lambda^1\mathbf{u}^1 + \lambda^2\mathbf{u}^2 \in B_\beta(\mathcal{A})$ for $\lambda^1, \lambda^2 \geq 0$ with $\lambda^1 + \lambda^2 = 1$. For each $\mathbf{u}^j$ with $j = \{1, 2\}$ there exists some functions $(\mathbf{p}^j, \mathbf{w}^j)$ satisfying (IC), (PK($\mathbf{u}^j$)), (FA), and $\mathbf{w}^j(\mathbf{v}) \in \mathcal{A}$ for all $\mathbf{v}$. Because constraints are linear, we have that functions $(\mathbf{p}, \mathbf{w})$ given by $\mathbf{p}(\mathbf{v}) = \lambda^1\mathbf{p}^1(\mathbf{v}) + \lambda^2\mathbf{p}^2(\mathbf{v})$ and $\mathbf{w}(\mathbf{v}) = \lambda^1\mathbf{w}^1(\mathbf{v}) + \lambda^2\mathbf{w}^2(\mathbf{v})$ satisfy (IC), (PK($\mathbf{u}$)) and (FA). We conclude that $\mathbf{u} \in B_\beta(\mathcal{A})$ because $\mathbf{w}(\mathbf{v}) \in \mathcal{A}$ since $\mathcal{A}$ is convex.

Item 3 (B_β preserves compactness). Recall that, by the Heine-Borel theorem, the set $\mathcal{A} \subset \mathbb{R}^n$ is compact if and only if $\mathcal{A}$ is closed and bounded. Suppose that $\mathcal{A}$ is closed and bounded. We need show that $B_\beta(\mathcal{A})$ is closed and bounded. Boundedness follows trivially from the promise keeping constraint because the allocation is bounded and values have finite means. We next prove closedness.

Consider a sequence $\mathbf{u}^m \in B_\beta(\mathcal{A})$ converging to some $\mathbf{u} \in \mathbb{R}^n$. We need to show that $\mathbf{u} \in B_\beta(\mathcal{A})$. For each m there exists functions $(\mathbf{p}^m, \mathbf{w}^m)$ which satisfy (IC), (FA), (PK($\mathbf{u}^m$)) and $\mathbf{w}^m(\mathbf{v}) \in \mathcal{A}$. Because $\mathcal{A}$ is convex and compact, we obtain by Lemma 6.21, item 2 that $(\mathbf{p}^m, \mathbf{w}^m)$ lie in the weak-* compact set $\mathcal{M}_{\mathcal{A}}$. By passing to a subsequence if necessary we obtain that $\mathbf{p}^m$ and $\mathbf{w}^m$ weak-* converge to some $(\mathbf{p}, \mathbf{w}) \in \mathcal{M}_{\mathcal{A}}$. By Lemma 6.21, item 3 the set of incentive compatible stage mechanisms is weak-* closed, and thus $(\mathbf{p}, \mathbf{w})$ is incentive compatible. Because $\mathcal{A}$ is closed and convex, we obtain by Lemma 6.21, item 1 that $\mathbf{w}(\mathbf{v}) \in \mathcal{A}$ for almost all $\mathbf{v} \in [0, \bar{v}]^n$. We conclude by showing that $(\mathbf{p}, \mathbf{w})$ satisfies the promise keeping constraint with $\mathbf{u}$. For each m we have that the promised keeping constraint is equivalently given by $\mathbf{u}^m = E(\mathbf{p}^m, \mathbf{w}^m)$ with the function E defined in Lemma 6.21, item 4. Because the function E is weak-* continuous, we obtain by taking limits that $\mathbf{u} = E(\mathbf{p}, \mathbf{w})$. Thus, $\mathbf{u} \in B_\beta(\mathcal{A})$ since the functions $(\mathbf{p}, \mathbf{w})$ satisfy (IC), (FA), (PK($\mathbf{u}$)) and $\mathbf{w}(\mathbf{v}) \in \mathcal{A}$. □

6.8.3 Proof of Lemma 6.20

Proof. We first prove that $\mathcal{U}_\beta$ is convex and compact, and then show that $\mathcal{U}_\beta = \mathbb{R}^n_+ \cap \operatorname{hyp}(\mathcal{U}_\beta)$.

Step 1 ($\mathcal{U}_\beta$ is convex and compact). From Proposition 6.5 we have $\mathcal{U}^\infty = \bigcap_{m \geq 1} \mathcal{U}^m$ with $\mathcal{U}^0 = [0, \mathbb{E}[v]]^n$ and $\mathcal{U}^m = B_\beta(\mathcal{U}^{m-1})$ for all $m \geq 1$. Because $\mathcal{U}^0$ is convex and compact, we obtain that $\mathcal{U}^m$ is convex and compact because the operator B_β preserves convexity and compactness from Lemma 6.22, items 2 and 3. Therefore, $\mathcal{U}^\infty$ is convex and compact, since the intersection of arbitrary convex and compact sets in $\mathbb{R}^n$ is convex and compact. Convexity and compactness of $\mathcal{U}_\beta$ follows because $\mathcal{U}^\infty = \mathcal{U}_\beta$ from Proposition 6.5.

Step 2 ($\mathcal{U}_\beta = \mathbb{R}^n_+ \cap \operatorname{\mathbf{hyp}}(\mathcal{U}_\beta)$). Note that $\mathcal{U}_\beta$ is a subset of $\mathbb{R}^n_+$ by definition. Moreover, any set is in its hypo-graph. Thus, we get $\mathcal{U}_\beta \subseteq \mathbb{R}^n_+ \cap \operatorname{hyp}(\mathcal{U}_\beta)$. To prove the converse, fix a point $\mathbf{u} \in \mathbb{R}^n_+ \cap \operatorname{hyp}(\mathcal{U}_\beta)$. We need to show that $\mathbf{u} \in \mathcal{U}_\beta$. Because $\mathbf{u} \in \mathbb{R}^n_+ \cap \operatorname{hyp}(\mathcal{U}_\beta)$, this implies that $\mathbf{u} \geq \mathbf{0}$ and there exists $\bar{\mathbf{u}} \in \mathcal{U}_\beta$ such that $\bar{\mathbf{u}} \geq \mathbf{u}$. Let $\bar{\boldsymbol{\pi}}$ be the mechanism corresponding to $\bar{\mathbf{u}}$.

Consider the 2^n states that can be obtained by replacing each subset of components of $\bar{\mathbf{u}}$ with zero. For example, when $n = 2$, we get $\bar{\mathbf{u}}^0 = (0, 0)$, $\bar{\mathbf{u}}^1 = (0, \bar{u}_2)$, $\bar{\mathbf{u}}^2 = (\bar{u}_1, 0)$, $\bar{\mathbf{u}}^4 = \bar{\mathbf{u}}$. All these 2^n states are in $\mathcal{U}_\beta$. In particular, for a state $\bar{\mathbf{u}}^m$, we can construct a PBIC mechanism $\bar{\boldsymbol{\pi}}^m$ satisfying $\bar{\mathbf{u}}^m_i = V_i(\bar{\boldsymbol{\pi}}^m, \mathbf{I})$ for all i by setting $\bar{\pi}^m_i = 0$ if $\bar{u}^m_i = 0$ and $\bar{\pi}^m_i = \bar{\pi}_i$ otherwise. Therefore, $\bar{\mathbf{u}}^m \in \mathcal{U}_\beta$ for all m. The polytope whose extreme points are the states $\bar{\mathbf{u}}^m$ given by $\{\tilde{\mathbf{u}} \in \mathbb{R}^n_+ : 0 \leq \tilde{u}_i \leq \bar{u}_i \ \ \forall i\}$ is a subset of $\mathcal{U}_\beta$ because all extreme points are in $\mathcal{U}_\beta$ and $\mathcal{U}_\beta$ is convex. The result follows because $\mathbf{u}$ lies in this polytope. □

6.8.4 Proof of Corollary 6.2

Proof. Because the set $\mathcal{U}_\beta$ is convex following Lemma 6.20, it can be characterized in terms of its support function ϕ_β. Proposition 3.2 implies that the support function ϕ_β is a fixed point of T_β. In step 2 of the proof of Proposition 6.5 we showed that the sequence of sets $\mathcal{U}^m = B_\beta^m(\mathcal{U}^0)$ is convex and compact, and converges to the set $\mathcal{U}_\beta$, which is convex and compact by Lemma 6.20. Therefore, we obtain from Salinetti and Wets (1979, Corollary P4.A in p.31) that $\phi_\beta(\boldsymbol{\alpha}) = \lim_{m\to\infty}[T_\beta^m \phi^0](\boldsymbol{\alpha})$ for all $\boldsymbol{\alpha}$. □

6.9 Generalizations of Proposition 3.4 and Proposition 3.5

In this section, we provide generalizations of Proposition 3.4 and Proposition 3.5 to scaled settings. Specifically, we consider these propositions for the scaled main phase mechanism $(\mathbf{p}^{(k)}, \mathbf{w}^{(k)})$ which is introduced in Chapter 4. Note that, Proposition 3.4 and Proposition 3.5 are special cases with $k = n$.

Proposition 6.6. *Suppose we are given a constant* $\underline{u}_\beta^{(n)} \in [0, \mathbb{E}[v]]$ *and a sequence* $\left\{\mathsf{F}_\beta^{(k,n)} \in (0,1)\right\}_{k=1}^n$. *Let* Ω_β *is determined by those constants. For any given state* $\mathbf{u} \in \Omega_\beta$ *with* k *active agents, the future promise vector* $\mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u})$ *lies within a plane in* $\mathbb{R}^k$. *Specifically, the plane is described by the following equation.*

$$\left(\mathbf{u}^{(k)} - \mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u})\right)^\intercal \boldsymbol{a} = \frac{1-\beta}{\beta}\left(s(\mathbf{u})\nabla\phi^{(k)}(\boldsymbol{a}) - \mathbf{u}^{(k)}\right)^\intercal \boldsymbol{a} \geq 0$$

where $\boldsymbol{a} = \boldsymbol{\alpha}^{(k)}(\mathbf{u}^{(k)}/(s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}))$.

Proof. We first show that b such that $\boldsymbol{a}^\intercal \mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u}) + b = 0$. Let $b = \sum_{i=1}^k a_i[(1-\beta)s(\mathbf{u})(\nabla\phi^{(k)}(\boldsymbol{a}))_i - u_i^{(k)}]/\beta$. Using the definition of $\mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u})$, we obtain the following equation.

$$\boldsymbol{a}^\intercal \mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u}) = \sum_{i=1}^k a_i \mathbb{E}_{\hat{v}_i}[W_i^{(k)}(\hat{v}_i|\mathbf{u})]\,.$$

Moreover, (PK($\mathbf{u}$)) implies that

$$\beta\mathbb{E}_{\hat{v}_i}[W_i^{(k)}(\hat{v}_i|\mathbf{u})] = \left[u_i^{(k)} - (1-\beta)s(\mathbf{u})\mathbb{E}[v_i \mathbf{1}\{a_i v_i \geq \max_{j\neq i} a_j v_j\}]\right].$$

Recall that $(\nabla\phi^{(k)}(\boldsymbol{a}))_i = \mathbb{E}[v_i\mathbf{1}\{a_i v_i \geq \max_{j\neq i} a_j v_j\}]$. Combining these three observations, it follows that $\boldsymbol{a}^\intercal \mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u}) + b = 0$.

The second part of the proposition follows from the following properties of the support function $\phi^{(k)}(\cdot)$. Note that $\boldsymbol{x}^\intercal \nabla\phi^{(k)}(\boldsymbol{x}) = \phi^{(k)}(\boldsymbol{x})$ for all $\boldsymbol{x}$ by the definition of $\phi^{(k)}(\cdot)$. Furthermore, for a vector $\boldsymbol{x}$ the support function satisfies that $\phi^{(k)}(\boldsymbol{x}) \geq \boldsymbol{x}^\intercal \boldsymbol{z}$ for all $\boldsymbol{z} \in \mathcal{U}^{(k)}$ (see Schneider, 2013).

Because $\mathbf{u}^{(k)} \in s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathcal{U}^{(k)}$, we obtain the following inequality, which concludes the proof.

$$\left(s(\mathbf{u})\nabla\phi^{(k)}(\boldsymbol{a}) - \mathbf{u}^{(k)}\right)^\intercal \boldsymbol{a} = s(\mathbf{u})\nabla\phi^{(k)}(\boldsymbol{a})^\intercal \boldsymbol{a} - \left(\mathbf{u}^{(k)}\right)^\intercal \boldsymbol{a} = s(\mathbf{u})\phi^{(k)}(\boldsymbol{a}) - \left(\mathbf{u}^{(k)}\right)^\intercal \boldsymbol{a} \geq 0\,. \qquad \square$$

Proposition 6.7. *Suppose we are given a constant* $\underline{u}_\beta^{(n)} \in [0, \mathbb{E}[v]]$ *and a sequence* $\left\{\mathsf{F}_\beta^{(k,n)} \in (0,1)\right\}_{k=1}^n$. *Let* Ω_β *is determined by those constants. Fix a state* $\mathbf{u} \in \Omega_\beta$ *with* k *active agents. Then, the ex-post future promise* $\mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u})$ *is an optimal solution of the following optimization problem:*

$$\begin{aligned} &\max_{\tilde{\mathbf{w}}(\cdot)} \min_{\mathbf{v}} \boldsymbol{\alpha}^{(k)}(\mathbf{u}^{(k)}/(s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}))^\intercal(\mathbf{u}^{(k)} - \tilde{\mathbf{w}}(\mathbf{v})) \\ &s.t.\ \mathbb{E}_{\mathbf{v}_{-i}}[\tilde{w}_i(v_i, \mathbf{v}_{-i})] = W_i^{(k)}(v_i|\mathbf{u}) \qquad \forall v_i, i = 1, \ldots, k \end{aligned} \tag{6.33}$$

Proof. For simplicity, we use $\boldsymbol{a}$ to represent $\boldsymbol{a}^{(k)}(\mathbf{u}^{(k)}/(s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}))$ in this proof. First, we reformulate problem (6.33) as the following linear programming problem:

$$\max_{\tilde{\mathbf{w}}(\cdot)} \quad z$$

$$\text{st.} \quad \mathbb{E}[\tilde{w}_i(v_i, \mathbf{v}_{-i})] = W_i^{(k)}(v_i|\mathbf{u}) \quad \forall v_i, i = 1, \ldots, k \tag{6.34}$$

$$z \leq \boldsymbol{a}^\intercal(\mathbf{u}^{(k)} - \tilde{\mathbf{w}}(\mathbf{v})) \qquad \forall \mathbf{v} \tag{6.35}$$

Next, we find an upper bound for this problem by relaxing the constraints over z.

The support function of set $s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathcal{U}^{(k)}$ is $s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\phi^{(k)}(\cdot)$, where $\phi^{(k)}(\boldsymbol{x}) = \mathbb{E}[\max_{i\in k} x_i v_i]$. Because $\mathbf{u}^{(k)} \in \mathcal{E}(s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathcal{U}^{(k)})$, we have $\mathbf{u}^{(k)}/s(\mathbf{u})\mathsf{F}_\beta^{(k,n)} \in \mathcal{E}(\mathcal{U}^{(k)})$, which implies $\mathbf{u}^{(k)}/s(\mathbf{u})\mathsf{F}_\beta^{(k,n)} = \nabla\phi^{(k)}(\boldsymbol{a})$, i.e., $u_i^{(k)} = s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}\mathbb{E}[v_i\mathbf{1}\{a_i v_i \geq \max_{j\neq i} a_j v_j\}]$. Moreover, $u_i^{(k)} = (1-\beta)s(\mathbf{u})\mathbb{E}[v_i\mathbf{1}\{a_i v_i \geq \max_{j\neq i} a_j v_j\}] + \beta\mathbb{E}[W_i^{(k)}(v_i|\mathbf{u})]$, following (PK($\mathbf{u}^{(k)}$)). These two observations imply that $u_i^{(k)} = \Upsilon\mathbb{E}[W_i^{(k)}(v_i|\mathbf{u})]$, where $\Upsilon = \dfrac{\beta s(\mathbf{u})\mathsf{F}_\beta^{(k,n)}}{s(\mathbf{u})\mathsf{F}_\beta^{(k,n)} - (1-\beta)}$.

Therefore, the constraint (6.35) becomes

$$z \leq \sum_{i=1}^k a_i(\Upsilon\mathbb{E}[W_i^{(k)}(v_i|\mathbf{u})] - \tilde{w}_i(\mathbf{v})) \text{ for all } \mathbf{v}\ .$$

Because z has to satisfy this constraint for all $\mathbf{v}$, we can relax it by replacing these constraints with a single constraint where the right hand side is the expectation over $\mathbf{v}$. This operation corresponds to taking expectation of $\tilde{w}_i(\mathbf{v})$. Following constraint (6.34), we obtain

$$z \leq \sum_{i=1}^{k} a_i(\Upsilon - 1)\mathbb{E}[W_i^{(k)}(v_i|\mathbf{u})] \tag{6.36}$$

Therefore, $\sum_{i=1}^{k} a_i(\Upsilon - 1)\mathbb{E}[W_i^{(k)}(v_i|\mathbf{u})]$ is an upper bound for the optimization problem (6.33). To show $\mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u})$ is an optimal solution, we first show that $\mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u})$ is feasible, and the objective function evaluated at $\mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u})$ is equal to this upper bound.

By its definition, the expectation of $w_i^{(k)}(\mathbf{v}|\mathbf{u})$ with respect to $\mathbf{v}_{-i}$ is $W_i^{(k)}(v_i|\mathbf{u})$, which implies feasibility. By Proposition 6.6, we also know that $\boldsymbol{a}^\intercal(\mathbf{u}^{(k)} - \mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u})) = \sum_{i=1}^{k} a_i(\Upsilon - 1)\mathbb{E}[W_i^{(k)}(v_i|\mathbf{u})]$ which, in turn, implies that $\mathbf{w}^{(k)}(\mathbf{v}|\mathbf{u})$ is an optimal solution for (6.33). $\square$

Bibliography

Abreu, Dilip, Benjamin Brooks, Yuliy Sannikov. 2017. A "pencil-sharpening" algorithm for two player stochastic games with perfect monitoring. Working Paper.

Abreu, Dilip, David Pearce, Ennio Stacchetti. 1990. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica: Journal of the Econometric Society* 1041–1063.

Aliprantis, Charalambos D, Kim Border. 2006. *Infinite dimensional analysis: a hitchhiker's guide*. Springer Science & Business Media.

Anderson, Edward J, Peter Nash. 1987. *Linear programming in infinite-dimensional spaces: theory and applications*. John Wiley & Sons.

Arrow, Kenneth. 1979. *The property rights doctrine and demand revelation under incomplete information*. Economics and human welfare. New York Academic Press.

Balseiro, Santiago, Huseyin Gurkan, Peng Sun. 2017. Multi-agent mechanism design without money. Forthcoming in Operations Research.

Bergemann, Dirk, Maher Said. 2011. *Dynamic auctions: A survey*. John Wiley & Sons, Inc.

Bergemann, Dirk, Juuso Välimäki. 2010. The dynamic pivot mechanism. *Econometrica* **78**(2) 771–789.

Boyd, Stephen, Lieven Vandenberghe. 2009. *Convex Optimization*. Cambridge University Press.

d'Aspremont, Claude, Louis-André Gérard-Varet. 1979. Incentives and incomplete information. *Journal of Public economics* **11**(1) 25–45.

De Farias, Daniela Pucci, Benjamin Van Roy. 2003. The linear programming approach to approximate dynamic programming. *Operations research* **51**(6) 850–865.

Esö, Péter, Gabor Futo. 1999. Auction design with a risk averse seller. *Economics Letters* **65**(1) 71–74.

Fernandes, Ana, Christopher Phelan. 2000. A recursive formulation for repeated agency with history dependence. *Journal of Economic Theory* **2**(91) 223–247.

Friedman, Eric J, Joseph Y Halpern, Ian Kash. 2006. Efficiency and nash equilibria in a scrip system for p2p networks. *Proceedings of the 7th ACM conference on Electronic commerce*. ACM, 140–149.

Fudenberg, Drew, David Levine, Eric Maskin. 1994. The folk theorem with imperfect public information. *Econometrica: Journal of the Econometric Society* 997–1039.

Gershkov, Alex, Benny Moldovanu. 2010. Efficient sequential assignment with incomplete information. *Games and Economic Behavior* **68**(1) 144–154.

Gorokh, Artur, Siddhartha Banerjee, Krishnamurthy Iyer. 2016. Near-efficient allocation using artificial currency in repeated settings. Working Paper.

Guo, Mingyu, Vincent Conitzer, Daniel M Reeves. 2009. Competitive repeated allocation without payments. *International Workshop on Internet and Network Economics*. Springer, 244–255.

Guo, Yingni, Johannes Hörner. 2015. Dynamic mechanisms without money. Tech. rep., Cowles Foundation for Research in Economics, Yale University.

Jackson, Matthew O, Hugo F Sonnenschein. 2007. Overcoming incentive constraints by linking decisions. *Econometrica* **75**(1) 241–257.

Johnson, Kris, David Simchi-Levi, Peng Sun. 2014. Analyzing scrip systems. *Operations Research* **62**(3) 524–534.

Johnson, Terence R. 2014. Dynamic mechanism design without transfers: Promises and confidentiality. Working paper.

Judd, Kenneth L, Sevin Yeltekin, James Conklin. 2003. Computing supergame equilibria. *Econometrica* **71**(4) 1239–1254.

Kash, Ian A, Eric J Friedman, Joseph Y Halpern. 2007. Optimizing scrip systems: Efficiency, crashes, hoarders, and altruists. *Proceedings of the 8th ACM conference on Electronic commerce*. ACM, 305–315.

Kash, Ian A, Eric J Friedman, Joseph Y Halpern. 2012. Optimizing scrip systems: crashes, altruists, hoarders, sybils and collusion. *Distributed Computing* **25**(5) 335–357.

Kash, Ian A, Eric J Friedman, Joseph Y Halpern. 2015. An equilibrium analysis of scrip systems. *ACM Transactions on Economics and Computation* **3**(3) 13.

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

Luenberger, David G. 1969. *Optimization by Vector Space Methods*. 1st ed. John Wiley & Sons, Inc., New York, NY, USA.

Marcinkiewicz, J, A Zygmund. 1937. Quelques théoremes sur les fonctions indépendantes. *Fund. Math* **29** 60–90.

Milgrom, Paul, Ilya Segal. 2002. Envelope theorems for arbitrary choice sets. *Econometrica* **70**(2) 583–601.

Myerson, R. 1981. Optimal auction design. *Mathematics of Operations Research* **6**(1) 58–73.

Nisan, Noam, Tim Roughgarden, Eva Tardos, Vijay V Vazirani. 2007. *Algorithmic game theory*, vol. 1. Cambridge University Press Cambridge.

Parkes, David C, Satinder P Singh. 2004. An mdp-based approach to online mechanism design. *Advances in neural information processing systems*. 791–798.

Pringle, RM, AA Rayner. 1970. Expressions for generalized inverses of a bordered matrix with application to the theory of constrained linear models. *Siam Review* **12**(1) 107–115.

Ross, Sheldon M. 1996. *Stochastic Processes*. John Wiley and Sons.

Salinetti, Gabriella, Roger J.-B. Wets. 1979. On the convergence of sequences of convex sets in finite dimensions. *SIAM Review* **21**(1) 18–33.

Schneider, Rolf. 2013. *Convex bodies: the Brunn–Minkowski theory*. 151, Cambridge University Press.

Shapiro, Alexander, Darinka Dentcheva, et al. 2014. *Lectures on stochastic programming: modeling and theory*, vol. 16. SIAM.

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

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

Vohra, Rakesh V. 2011. *Mechanism Design: A Linear Programming Approach*. Cambridge University Press.

Yeltekin, Sevin, Yongyang Cai, Kenneth L. Judd. 2017. Computing equilibria of dynamic games. *Operations Research* **65**(2) 337–356.

Zhang, H. 2012a. Analysis of a dynamic adverse selection model with asymptotic efficiency. *Mathematics of Operations Research* **37**(3) 450–474.

Zhang, H. 2012b. Solving a dynamic adverse selection model through finite policy graphs. *Operations Research* **60**(4) 850–864.