UC Santa Barbara
UC Santa Barbara Electronic Theses and Dissertations

Title
Strategic Mechanisms in Multi-Agent Coordination

Permalink
https://escholarship.org/uc/item/55b2j5fq

Author
Chandan, Rahul

Publication Date
2022

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

University of California
Santa Barbara

Strategic Mechanisms in Multi-Agent Coordination

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Electrical and Computer Engineering

by

Rahul Chandan

Committee in charge:

Professor Jason R. Marden, Chair
Professor Mahnoosh Alizadeh
Professor João P. Hespanha
Professor Daniel Kovenock, Chapman University
Professor Dario Paccagnan, Imperial College London

September 2022

The Dissertation of Rahul Chandan is approved.

Professor Mahnoosh Alizadeh

Professor João P. Hespanha

Professor Daniel Kovenock, Chapman University

Professor Dario Paccagnan, Imperial College London

Professor Jason R. Marden, Committee Chair

July 2022

Strategic Mechanisms in Multi-Agent Coordination

Copyright © 2022

by

Rahul Chandan

To my parents, Anil and Vandana Chandan.

Acknowledgements

Foremost, I would like to express my deep gratitude to my advisor, Prof. Jason R. Marden, who supported me throughout my Ph.D. studies, for his valuable insights regarding technical writing and presentations, for providing me with every opportunity to excel, and for sharing his research vision. I could not ask for a better advisor, mentor and friend with whom to spend five years working on the various projects that make up this dissertation.

Besides my advisor, I would like to thank my unofficial co-advisors: the two “Prof. P’s,” Prof. Dario Paccagnan at Imperial College London and Prof. Keith Paarporn at University of Colorado, Colorado Springs, for their patience and enthusiasm when I began working with them as a fresh graduate student, and for being outstanding role models since then. I would also like to thank the rest of my thesis committee: Professors Mahnoosh Alizadeh, João P. Hespanha and Daniel Kovenock for sharing an immense wealth of knowledge, asking insightful questions and for suggesting exciting avenues for future research.

My sincere thanks goes also to my labmates, and other colleagues at the Center for Control, Dynamical Systems, and Computation (CCDC) at the University of California, Santa Barbara. Though they are too numerous to list, each played a part in moulding me into the researcher and individual I am today.

This thesis would not have been possible without the support of the any friends that I made over the past five years, as well as numerous friends that I have fortunately carried over from childhood, high school, and my undergraduate studies. My heartfelt gratitude to them for every call, get-together, or message, especially during the recent years of pandemic.

Finally, I would like to thank my family: my parents Anil and Vandana Chandan, and my brothers Nakul and Nikhil Chandan, for fostering my passion to explore the unknown, and for their continued support through all my life’s adventures.

Curriculum Vitæ

Rahul Chandan

Education

2022	Ph.D. in Electrical and Computer Engineering, University of California, Santa Barbara.
2019	M.S. in Electrical and Computer Engineering, University of California, Santa Barbara.
2017	B.A.Sc. in Engineering Science – Electrical and Computer Engineering Option, University of Toronto.

Publications

Journal Publications

1.

Conference Publications

1.

Awards/Honours

1. **2022 Dissertation Fellowship**, UCSB
2. **2022 Outstanding TA Award**, **ECE Department**, UCSB
3. **2021 Best Community-voted Presentation**, *EAI GameNets 2021*, Online
4. **2021 Outstanding TA Award**, **ECE Department**, UCSB
5. **2020 Outstanding TA Award**, **ECE Department**, UCSB
6. **2017 UC Regents' Fellowship**, University of California
7. **2017 CCDC Outstanding Scholar Fellowship**, UCSB
8. **2017 Gordon Cressy Student Leadership Award**, University of Toronto

Abstract

Strategic Mechanisms in Multi-Agent Coordination

by

Rahul Chandan

Strategic interactions in multi-agent systems can be conveniently modeled, manipulated and characterized within the analytical framework provided by game theory and mechanism design, and used to extract engineering insights regarding systems of interest. Accordingly, in this dissertation, we adopt such a framework and pursue research directions aimed at understanding how information and externalities impact the strategic outcomes that can emerge in systems with multiple, non-cooperative decision makers. We consider two types of mechanisms: *decision-based mechanisms* and *preference-based mechanisms* which respectively manipulate either the players' action sets or the players' utilities – the two major building blocks of a game. We conduct our analysis of decision-based mechanisms on Colonel Blotto games which are popular models for competitive resource allocation in adversarial environments. Within this framework, we first consider settings where the information to a competitor is obfuscated, and quantify the value of information relating to competitive objectives and the opponent's strength. We then consider the role of pre-emption under this framework, and show that – perhaps surprisingly – revealing information to competitors can also offer strategic benefits in competitive interactions. Our results on preference-based mechanisms focus on the design of taxes in congestion games to optimize the system performance. In our study, we consider three performance measures corresponding with the worst-case equilibrium efficiency (Price of Anarchy), the best-case equilibrium efficiency (Price of Stability), and the transient system performance (Price of Urgency). Within this context, our first set of results focus on optimizing the Price of Anarchy; we derive tractable methodologies for computing the optimal taxes within this setting. We then investigate the consequences of optimizing for the worst case on the other performance measures: we show that the taxes that optimize the Price of Anarchy necessarily have Price of Stability equal to the Price of Anarchy, and that optimal Price of Anarchy guarantees can correspond with arbitrarily poor Price of Urgency. We supplement this last set of results by proposing techniques for characterizing the respective trade-off curves. We conclude with a discussion on future directions for both decision- and preference-based mechanisms.

Contents

Chapter 1

Introduction and Overview

Over the past century, the classical discipline of control theory has evolved from its humble beginnings in regulating oscillations in engines and the establishment of control stability criteria for linear systems, to PID control and nonlinear control, and, more recently, to adaptive control, hybrid control and optimal control. The classical approach is *extremely* successful in its application to systems operated under the authority of a single decision maker (DM), including, for example, power converters, refridgerators, automobiles, and HVAC systems. Though implementing such a control approach is often the most straightforward and efficient, it is impractical when deploying large-scale systems that must satisfy imposing communication and security constraints, and generally inadmissible when considering systems involving human DMs. It is unsurprising, then, that as the engineered systems around us continue to grow larger, more complex and more human-centric, we are seeing research in control theory trend away from the classical discipline toward relatively new and unexplored fields. Consider, for example, the increasing popularity of non-traditional disciplines such as machine learning, human-robot interaction and game theory at control theory conferences such as the IEEE Conference on Decision and Control (CDC) and American Control Conference (ACC).

In this dissertation, we pursue the study of control theory using models and analytical tools borrowed from game theory. While the field of game theory originates from mathematics and economics, it has received significant interest in engineering and beyond, for example, in biology and medicine [1, 2], computer science [3, 4], and philosophy and political science [5, 6]. Game theory is well-suited to the study of large-scale systems as well as systems with human DMs since it allows us to represent a system, its DMs and the DMs' interactions under the abstraction of a set of players (or agents) in a game, each with a set of actions describing her decision space and a utility function describing her preferences. This abstraction allows us to study the set of system outcomes that emerge from the DM's strategic behaviour (in other words, we assume that the DMs act rationally).

Under the abstraction provided by game theory, influencing the set of emergent system outcomes amounts to altering the agents' action sets and/or utility functions, which is the central focus in the discipline of *mechanism design* [7]. In our application of mechanism design to engineering, we treat established game models for systems of interest as the baseline/nominal settings, and study how the emergent system outcomes change in response to designed perturbations on the agents' action sets and/or preferences. Accordingly, we refer to such perturbations as *mechanisms*. By selecting the appropriate mechanisms and identifying the differences between the nominal and perturbed systems' outcomes, we hope to develop engineering insights that inform how one should – or, perhaps more importantly, *should not* – design such systems in practice.

In this work, we propose the study of mechanisms under two classifications:

– *Decision-based mechanisms:* Mechanisms that modify the agents' action sets.

– *Preference-based mechanisms:* Mechanisms that modify the agents' utility functions.

Observe that applying either class of mechanism to a game model may lead to dramatically different strategic outcomes. We study decision-based and preference-based mechanisms under well-studied game theoretic models respectively known as Colonel Blotto games and congestion games. Interestingly, the original versions of both these game models pre-date game theory itself: though the foundation of game theory as a field is widely attributed to John von Neumann's *On the Theory of Games of Strategy* in 1928, while the original Colonel Blotto game was proposed by Émile Borel no later than 1921 and Arthur C. Pigou studied a form of congestion game as far back as 1920.

The goal of this dissertation is to further the applications of game theory in engineering (especially in control theory) by providing novel perspectives and methodologies for analyzing multi-agent interactions and coordination. Although some of the forthcoming insights may appear to be prescriptive – in other words, they may seemingly apply *directly* to practical problem settings – it is important to note that our insights only hold under the multitude of modeling assumptions that make up our analytical framework. If there is any truth to our results, it is that the emergent behaviour in multi-agent systems – particularly in those involving human DMs – is even more complex than it appears in our analysis, as even the simple models studied within this dissertation are rich with performance trade-offs, mis-directions and counter-measures, among other fascinating phenomena.

1.1 Why game theory?

On page 2 of their 1998 book *Dynamic noncooperative game theory*, Başar and Olsder [8] attribute to game theory "the development of suitable concepts to describe and understand conflict situations." Indeed, game theory may be used as a catch-all term for any study of systems with multiple DMs. Consider, for example, that any control theory problem can be posed as a game where the controllers are the players, the output spaces are their action sets, and the control objectives encode their utilities.

Nevertheless, a variety of well-established analytical techniques have already been proposed and used within control theory for the study of systems with multiple DMs. This prompts the question: "Why should we study game theory when other techniques already seek to address multi-agent coordination?" The simple answer is that game theory offers an alternative perspective on multi-agent coordination that these other perspectives cannot provide. A more complete answer is that game theory uniquely studies the set of stable outcomes (i.e., equilibria) that emerge under any form of decision making, and offers convenient models and analytical methodologies for studying the design of multi-agent systems around this set of outcomes.

To the authors' best knowledge, the most popular alternatives to game theory for the control of multi-agent systems are stochastic and robust control, dynamical systems, and distributed optimization. In the following, we summarize these approaches, and discuss how game theory offers complementary perspectives on multi-agent interactions and coordination:

Stochastic and robust control. In the disciplines of stochastic and robust control, we seek to design controllers that guarantee that a given system operates within a desired performance regime under noisy (stochastic) or worst-case (robust) disturbances. Often, these disturbances can be said to model interactions between the system and an added DM (either nature or an adversary) in the system. For example, a robust control approach assumes that the added DM is purely adversarial toward the system. Though games can and sometimes do possess purely adversarial DMs, with game theory we can also model systems with nuanced DM preferences that are neither purely adversarial nor pure uncorrelated with the preferences of the other DMs.

Dynamical systems. Dynamical systems is the study of any system governed by a set of dynamics [9]. As the name suggests, the study of systems from this perspective attributes significant importance to the underlying system dynamics. In the case of multi-agent systems, dynamical systems have been extremely successful in modelling phenomena such as the synchronization of oscillators [10] and their applications to the electric power grid [11], social networks and biological systems [12, 13]. Although dynamics, especially dynamics that converge to equilibrium, are an important topic of study within game theory (see, for example, [14, 15]), the focus in game theory is also on the outcomes of strategic decision making and, thus, the system dynamics are not a necessary component of the system analysis. This is an especially desirable property when the system dynamics are too complex to be modelled (for example, human decision making).

Distributed optimization. Another approach is distributed optimization which captures the objective and constraints of a given multi-agent system as an optimization problem and uses structural properties and/or relaxations of this optimization problem to formulate distributed control protocols [16]. Such approaches are often associated with (near-)optimal approximation guarantees and are especially useful for designing distributed algorithms for systems with programmable DMs [17, 18]. The study of game theory differs with this approach in spirit in that we are often interested in studying and contrasting the strategic outcomes in different systems without ever posing the optimization question.

A final, additional benefit to studying game theory is that it is inherently a multi-disciplinary field. Thus, when posed as a game theory problem, a seemingly new research direction within control theory may already be extensively studied in computer science, economics, operations research, etc.

1.2 Summary of contributions

This dissertation is divided into two parts corresponding to the two classifications of mechanisms proposed in the previous sections. Part I presents our work relating to agent-level mechanisms and their applications in Colonel Blotto games and variations thereof. Part II presents our work relating to system-level mechanisms and their applications in congestion games. In this section, we outline the research agenda and summarize our contributions for each of these two parts.

1.2.1 Part I: Decision-based mechanisms

In the first part of this dissertation, we focus on *decision-based mechanisms* and their application to competitive resource allocation settings. Specifically, we adopt the framework of *Colonel Blotto games* and variations thereof (especially *General Lotto games*) that model the competitive interactions of budget-constrained players over valuable battlefields, and investigate the role that information plays in the emergent strategic behaviour of such games. First, we aim to identify the value associated with knowing the opponents' budgets and knowing the battlefields' values. Second, we wish to understand whether a competitor can strategically reveal information to her opponents to improve her competitive position.

Outline. In Chapter 3, we introduce the Colonel Blotto and General Lotto game models considered, and provide a discussion on possible applications and related works. In Chapter 4, we introduce Bayesian formulations of the General Lotto game with asymmetric information on either the values of the battlefields or players' budgets, and identify the value of information on these pieces of information to a competitor. In Chapter 5, we identify the strategic benefits of pre-emption both in the form of battlefield and value concessions in three-player General Lotto games (which we review in the chapter) as well as pre-allocated resources in General Lotto games.

Contributions. The main contributions in Part I are discussed in Chapters 4 and 5, and are summarized as follows:

1. In Section 4.1, we propose the General Lotto game with incomplete value information, which permits us to study the value of battlefield information in the General Lotto game. We provide a complete equilibrium characterization in the setting with completely asymmetric information (i.e., one player can see the realized system state while the other only has access to the prior distribution governing the system state), the various system states are all the permutations of a given vector of battlefield values, and the prior distribution is uniform (Theorem 4.1.1). We also characterize how the players' equilibrium payoffs change as we transition from the completely asymmetric information to complete information settings (Theorem 4.1.2).

2. In Section 4.2, we propose the General Lotto game with asymmetric budget information, which permits us to study the value of budget information in the General Lotto game. First, we provide a complete equilibrium characterization of the class of games where the

informed player's budget is governed by a Bernoulli prior distribution (Theorem 4.2.1). We then compare the players' equilibrium payoffs in the setting under Bernoulli prior distribution against the complete information setting, and the setting where the informed player selects the optimal prior distribution (Corollaries 4.2.1 and 4.2.2).

3. In Section 5.1, we propose the three-player General Lotto game with budget concessions and the three-player General Lotto game with value concessions. First, we show that budget concessions never offer strategic opportunities to any player (Theorem 5.1.1). Then, we establish that value concessions can strictly improve a player's equilibrium payoff and fully characterize the set of pure strategy Nash equilibria for any given three-player General Lotto game with value concessions (Theorem 5.1.2). Finally, we show that neither concession format offers strategic opportunities in the (two-player) General Lotto game.

4. In Section 5.2, we propose the General Lotto game with pre-allocations. First, we fully characterize the players' equilibrium payoffs to both players within such games (Theorem 5.2.1) which allows us to identify the level sets within the parameter space (Theorem 5.2.2). Based on these results, we establish that regular resources are at least twice as effective as pre-allocated resources (Corollary 5.2.1), and solve the optimal investment problem when these two types of resources have linear costs (Corollary 5.10).

1.2.2 Part II: Preference-based mechanisms

In the second part of this dissertation, we focus on *preference-based mechanisms* and their application to non-cooperative multi-agent systems such as selfish routing in congestible networks and distributed resource allocation. Specifically, we adopt the framework of *congestion games* that model the noncooperative interactions of self-interested decision makers whose local utilities are distinct from an overarching system objective, and investigate how taxes can be used to improve the system performance associated with the emergent strategic outcomes. First, we consider a robust design approach in which we propose tractable methodologies for computing the taxes that optimize the worst-case equilibrium efficiency. We then consider the consequences of the worst-case design approach on the best-case equilibrium efficiency and transient system behaviour, identifying and characterizing the trade-offs between the worst-case equilibrium efficiency and both these performance measures.

Outline. In Chapter 6, we introduce the congestion game model, taxation mechanisms and relevant performance measures, and provide a discussion on possible applications and related works. In Chapter 7, we review and draw connections between existing analytical techniques for characterizing bounds on the equilibrium efficiency in games including smoothness and primal-dual techniques, and propose game parameterizations that balance the tractability and tightness of these bounds. In Chapter 8, we consider the design of taxes to optimize the worst-case equilibrium efficiency, and provide tractable methodologies to accomplish this. In Chapter 9, we study the consequences of optimizing the worst-case equilibrium efficiency on the best-case equilibrium efficiency and transient system performance, identifying and characterizing performance trade-offs corresponding with both these directions.

Contributions. The main contributions in Part II are contained in Chapters 7, 8 and 9, and are summarized as follows:

1. In Section 7.2, we introduce a novel smoothness notion – termed generalized smoothness - that represents a slight modification of Roughgarden's smoothness notion. Nevertheless, we show that generalized smoothness applies to a broader class of problems, and offers improved bounds on the equilibrium efficiency when compared to Roughgarden's smoothness (Proposition 7.2.1).

2. In Section 8.1, we consider the problem of designing taxes that optimize the worst-case equilibrium efficiency in congestion games. We derive a methodology based on tractable linear programs for computing the optimal taxes and the corresponding equilibrium efficiency guarantees (Theorem 8.1.1). We then show that, for nondecreasing, convex resource

cost functions, these linear programs can be further simplified, and that the equilibrium efficiency guarantees can be written in closed-form (Theorem 8.1.2).

3. In Section 8.2, we seek to identify universal performance guarantees on the equilibrium efficiency without relying on linear programming. For the class of resource allocation games with nonnegative, concave nondecreasing welfare functions, we show that there exist utility mechanisms that guarantee that the equilibrium efficiency is no lower than $1 - c/e$ where c is the curvature of the welfare function and e is Euler's constant (Theorem 8.2.1). Perhaps surprisingly, this efficiency guarantee matches the best-achievable approximation ratio among all polynomial-time centralized algorithm for this class of problems.

4. In Section 9.1, we seek to characterize the tension between the worst- and best-case equilibrium efficiency in congestion games. In this context, we identify the existence of two separate trade-offs between these two performance measures. First, we show that the taxes that optimize the worst-case equilibrium efficiency have corresponding best-case equilibrium efficiency guarantees equal to the worst-case equilibrium efficiency (Theorem 9.1.1). We then propose techniques for characterizing upper and lower bounds on the trade-off curve in the joint optimization of the worst- and best-case equilibrium efficiency (Theorems 9.1.2 and 9.1.3). Next, we consider the existence of an inner trade-off between the worst- and best-case equilibrium efficiency pairs that are jointly achieved within the same game. Though such an inner trade-off does not exist under the taxes that optimize either the worst- or best-case equilibrium efficiency (Theorem 9.1.1), we show that it does in general (i.e., in the setting without taxes) (Theorem 9.1.2).

5. In Section 9.2, we study the transient system performance in resource allocation games with nondecreasing, concave resource welfare functions. We consider the setting where agents perform $\kappa \geq 1$ round-robin best response sequences before arriving at the solution of interest. We first derive a linear program for computing the utility mechanism that optimizes the system performance guarantees after a one-round walk ($\kappa = 1$) (Theorem

9.2.1). Using this result, we show that the best-achievable system performance after a one-round walk is $1 - c/2$ where c is the curvature of the resource welfare function, and that the best-achievable system perofrmance after $\kappa > 1$ is no better (Theorem 9.2.3). Finally, we show that a trade-off exists between the system performance guarantees after a one-round walk and at equilibrium (Theorems 9.2.4 and 9.2.5). Notably, the performance guarantees after a one-round walk can be arbitrarily poor under the utility design that optimizes the worst-case equilibrium efficiency.

Chapter 2

Mathematical preliminaries

2.1 Game theory and strategic decision making

In this section, we review several important concepts from game theory. However, this section is not intended to provide a complete introduction to game theory. For more complete introductions to this topic, we refer the interested reader to [3, 4].

2.1.1 Formal game definition

Formally, a game consists of a set of n players, which we denote as $N = \{1, \ldots, n\}$. Each player $i \in N$ has a *set of permissible actions*, which we denote as $\mathcal{A}_i$. During play, each player selects an action $a_i \in \mathcal{A}_i$, resulting in an *action profile* $a = (a_1, \ldots, a_n)$ at the end of play. We use $\mathcal{A} = \Pi_{i=1}^{n} \mathcal{A}_i$ to denote the set of all possible action profiles of the game.

Each player in a game associates a degree of satisfaction with each possible action profile of the game, imposing a (possibly weak) *preference ordering* over the set $\mathcal{A}$. In the definition of a game, it is *critical* that each player i's preference ordering be over the action profiles $a \in \mathcal{A}$ and not simply over actions $a_i \in \mathcal{A}_i$ to ensure that i's satisfaction is dependent not solely on her own action, but on the other players' actions as well. In general, the preference ordering over $\mathcal{A}$ is different for each player. In this dissertation, we will specify the preference ordering for each player by assigning a value to each possible action profile. Depending on what is appropriate for the application we consider, this value will represent either the payoff or cost experienced by the player under the corresponding action profile. We will use $\pi_i : \mathcal{A} \to \mathbb{R}$ and $c_i : \mathcal{A} \to \mathbb{R}$ to denote player i's payoff and cost functions, respectively. Note that payoffs and costs can be used interchangeably, since $\pi_i(a) = -c_i(a)$ preserves the preference ordering. Accordingly, we can denote any game as a tuple $G = (N, \{\mathcal{A}_i\}_{i \in N}, \{\pi_i\}_{i=1}^n)$.

Before our discussion on strategic outcomes of games, it is important to explore the various models of interaction in game theory. As we will see, the model of interaction will inform the solution concepts that we consider. In this dissertation, we focus on simultaneous, sequential, and incomplete information games, which we define below:

Definition 2.1.1 (Simultaneous game)**.** *A simultaneous game – also known as* static *or* one-shot *game – are games where each player selects her action without access to information on the other players' actions. Therefore, these games are useful in modelling systems in which the decision making takes place* simultaneously*. The classic example of a simultaneous game is rock-paper-scissors, while other practical examples include the interaction between offense and defense before each down in American football.*

Definition 2.1.2 (Sequential game)**.** *A sequential game – also known as,* Stackleberg *or* turn-based *game – is any game that is not simultaneous. In other words, at least one player has access to information on at least one other player's action. Naturally, such games model systems in which the decision making has a strict temporal ordering. The classic example of a sequential game is tic-tac-toe. Other practical examples include board games like chess and backgammon, and card games like bridge and hearts.*

Definition 2.1.3 (Incomplete information game)**.** *An incomplete information game is a game in which at least one player does not possess all the information about the game parameters. Thus, a player without complete information must select her action, for example, by reasoning about her utility or cost in expectation if a prior distribution over the possible game parameters is available. The classic example of an incomplete information game is an auction, as players must select their bids without knowing the other players' payoff functions.*

2.1.2 Nash equilibrium and other solution concepts

In any analytical framework, we would like to define solution concepts that we use to represent the set of system outcomes within our model. In the case of game theory, the standard assumption is that the players are self-interested in that each player $i \in N$ will select from the actions that she most prefers if we fix the actions of the other players, i.e., given any action profile $a \in \mathcal{A}$ and payoff functions $\pi_i \mathcal{A} \to \mathbb{R}$ for all $i \in N$, player i will select an action

$$a_i^* \in \arg\max_{a_i' \in \mathcal{A}_i} \pi_i(a_i', a_{-i}), \tag{2.1}$$

where a_i', a_{-i} denotes the action profile wherein player i selects the action $a_i' \in \mathcal{A}_i$, and every other player selects her action in a. We will refer to any action $a_i^* \in \mathcal{A}_i$ that satisfies (2.1) as player i's *best-response strategy* to action profile a. Since the preference ordering specified by the payoff functions can be weak, a given player's best-response strategy to a given action profile need not be unique in general. If a given action $a_i^* \in \mathcal{A}_i$ is player i's best-response strategy for any action profile a, it is additionally termed a *dominant strategy*. The definition of best-response strategies immediately prompts our first equilibrium notion: A *pure Nash equilibrium* of a given simultaneous game G is any action profile $a^{\mathrm{ne}} \in \mathcal{A}$ under which each player i's action a_i^{ne} is a best-response strategy to a^{ne}.

In general, a simultaneous game can have one or many pure Nash equilibria, but can also have no pure Nash equilibria as well. Consider, as a simple example, the rock-paper-scissors game. If player 1 selects the action "Rock," then player 2's best-response strategy is the action "Paper," which then makes player 1's best-response strategy "Scissors," making player 2's best-response strategy "Rock," and so on. Of course, anyone who has played rock-paper-scissors knows that it is best to mix between "Rock," "Paper" and "Scissors," or to *randomize* one's strategy. When players have *mixed strategies*, i.e., each player $i \in N$ can choose a probability distribution over her actions $\sigma_i \in \Delta(\mathcal{A}_i)$, then we assume that they act to maximize their *expected payoffs* under the joint probability distribution over action profiles $\sigma = \Pi_{i=1}^n \sigma_i$

where $\sigma \in \Delta(\mathcal{A})$. Here we have implicitly assumed that the players' mixed strategies are independently distributed. If we augment a game G by endowing each player with the choice of any mixed strategy over her action set, a *mixed Nash equilibrium* is a joint probability distribution $sigma^{\mathrm{ne}} = \Pi_{i=1}^{n} \sigma_i^{\mathrm{ne}}$ such that each player i's choice σ_i is a best-response mixed strategy to σ. One of the most important results – if not *the* most important result – in game theory was John Nash's proof that every game with a finite number of players, each with a finite set of actions, has a mixed Nash equilibrium.

In certain scenarios, we may wish to drop the assumption that players' mixed strategies are independently distributed, for example, if the players share communication channels or there is a game coordinator that chooses the players' strategies on their behalf. Under such correlated joint probability distributions $\sigma \in \Delta(\mathcal{A})$, we may consider the notion of *correlated equilibrium* as the solution concept, which is any distribution $\sigma \in \Delta(\mathcal{A})$ such that the marginal distribution σ_i is a best-response mixed strategy to σ for each player $i \in N$. For any given game, observe that the set of pure Nash equilibria is a subset of the set of mixed Nash equilibria, which is itself a subset of the set of correlated equilibria.

Observe that sequential and incomplete information games naturally give rise to their own set of solution concepts akin to those defined above for simultaneous games. First, an action profile in a sequential game is termed a *subgame perfect Nash equilibrium* if each player's action at each stage of the game is her best-response strategy given the actions played at earlier stages as well as the current stage of the game. Finally, we will assume that players with incomplete information use Bayes' rule based on given priors distribution over the possible game parameters to compute their expected payoffs. This leads to *Bayes Nash equilibrium* as the relevant solution concept.

2.1.3 Mechanisms

Classically, a mechanism is defined as a mapping from the set of all possible action profiles $\mathcal{A}$ to the set of players' utility or cost values [19]. In this sense, a mechanism defines a game upon a given set of players $N = \{1, \ldots, n\}$ with action sets $\{\mathcal{A}_i\}_{i=1}^n$. For this reason, the study of mechanisms and mechanism design – which refers to the design of a mechanism to induce a desired set of solutions – is also referred to as *game design* and *reverse game theory*.

In this dissertation, we propose a modified definition of mechanism which we believe retains the original spirit of the study of mechanisms while broadening its applications within engineering. In our definition, a mechanism is a mapping $\mathcal{M} : \mathcal{G} \times \Theta \to \mathcal{G}'$ from a nominal family of games $\mathcal{G}$ and some parameter space Θ to a family of modified games $\mathcal{G}'$. The idea is to compare the players' equilibrium payoffs in a given nominal game $G \in \mathcal{G}$ against their equilibrium payoffs in the modified game $G' = \mathcal{M}(G, \theta)$ corresponding to parameter vector $\theta \in \Theta$ to characterize the impact of this modification. Thus, under this definition, a mechanism can be used to modify a complete information game to study the effect of withholding some information from a player, to modify a simultaneous game to study the strategic outcomes in a sequential counterpart, among other interesting engineering perspectives.

Part I

Decision-Based Mechanisms

Chapter 3

Introduction

3.1 Model: Colonel Blotto games and variations

A Colonel Blotto game consists of two players A and B. Each possesses an amount $X_i > 0$, $i \in \{A, B\}$, of resources to distribute among a set of $B \geq 2$ battlefields. Each battlefield $b \in \{1, \ldots, B\}$ has an associated value $v_b \geq 0$. Each player $i \in \{A, B\}$ selects a vector $\mathbf{a}_i = (a_{i,1}, \ldots, a_{i,B})$ from the set of vectors $\mathcal{A}_i = \Delta(X_i) := \{\mathbf{a} \in \mathbb{R}^n \text{ s.t. } \mathbf{a} \geq 0, 1_B^T \mathbf{a} \leq X_i\}$. Given action profile $(\mathbf{a}_A, \mathbf{a}_B)$, player i's utility is given by

$$u_i(\mathbf{a}_i, \mathbf{a}_{-i}) := \sum_{b=1}^{B} v_b \cdot U(a_{i,b}, a_{-i,b}) \tag{3.1}$$

where

$$U(x, y) := \begin{cases} 1, & \text{if } x > y \\ 1/2, & \text{if } x = y \cdot \\ 0, & \text{if } x < y \end{cases} \tag{3.2}$$

Observe that a mixed strategy for player i is any B-variate distribution F_i on $\Delta(X_i)$. Hence, any pure strategy $\mathbf{a}$ drawn from the distribution F_i satisfies the budget constraint with probability one. Each player i's expected utility under distributions F_i and F_{-i} is evaluated as

$$\mathbb{E}_{\mathbf{a}_i \sim F_i} \mathbb{E}_{\mathbf{a}_{-i} \sim F_{-i}} \; u_i(\mathbf{a}_i, \mathbf{a}_{-i}). \tag{3.3}$$

We will refer to a particular instance of the Colonel Blotto game as $\mathrm{CB}(X_A, X_B, \mathbf{w})$.

3.1.1 General Lotto game

A *General Lotto* game is a popular variation of the Colonel Blotto game. The two models are identical in their definitions except for the players' action sets. In the General Lotto game, each player $i \in \{A, B\}$ selects a B-variate distribution F_i from the set of distributions $\mathcal{A}_i := \mathcal{L}(X_i)$, where we let $\mathcal{L}(X_i)$ denote the set of all distributions F on $\mathbb{R}_+^B$ that satisfy the following budget constraint:

$$\sum_{b=1}^{B} \mathbb{E}_{a_b \sim F}[a_b] \leq X_i \tag{3.4}$$

Note that this relaxes the budget constraint from the Colonel Blotto game that is satisfied with probability one to a budget constraint that is satisfied in expectation. Players' expected utilities under distributions F_i and F_{-i} are evaluated as in (3.3). It is important to note that the players' utility structures in the Colonel Blotto and General Lotto games define constant-sum games, since the players' utilities always sum to $\sum_{b=1}^{B} \; v_b$. It is well known that the players' equilibrium payoffs are unique in General Lotto games, as summarized in the following theorem:

Theorem 3.1.1 (General Lotto payoffs[20]). *Consider the General Lotto game with player budgets $X_A, X_B \geq 0$ and battlefield values $\mathbf{v} \geq 0$ with cumulative worth $\Phi = 1^T \mathbf{v}$. Player $i \in \{A, B\}$ has unique equilibrium payoff*[1]

$$u_i^*(X_i, X_{-i}, \Phi) = \begin{cases} \Phi \cdot \frac{X_i}{2X_{-i}} & \text{if } X_i \leq X_{-i}, \\ \Phi \cdot \left(1 - \frac{X_{-i}}{2X_i}\right) & \text{if } X_i > X_{-i}. \end{cases} \tag{3.5}$$

[1]Candidate equilibrium strategies are characterized in [21, 20], and consist of non-trivial probability distributions over $\mathbb{R}^n_{\geq 0}$. However, unlike the players' equilibrium payoffs, the equilibrium strategies in General Lotto games are not generally unique.

Observe that the players' equilibrium payoffs in the General Lotto game are independent of the specific values in the vector of battlefield values $\mathbf{v}$, and only depend on the cumulative value $\Phi = 1^T\mathbf{v}$, and the players' relative budgets. Thus, we will denote an instance of the General Lotto game as $\mathrm{GL}(X_A, X_B, \Phi)$. Figure 5.1a depicts the two-player General Lotto game.

3.2 Applications

Resource allocation decisions in adversarial environments are central to the design and operation of networked multi-agent systems. Adversarial models are especially prevalent in studies of cyber-physical systems control (e.g., perimeter defense [22, 23], attack identification [24, 25], data protection [26, 27]), and robust network control under disturbances (e.g., economic dispatch [28], leader-follower control [29]). The common objective in these settings is to ensure the best performance under strategic interference by the adversary. In adversarial contexts, zero- and constant-sum game models (like Colonel Blotto and General Lotto games) are particularly popular, as the gains of a given player necessarily come at a cost to the others. Such models have been applied to pursuit-evasion [30, 31], threat detection [32, 33] and secure control [34, 35], among other problems in cyber-physical systems.

3.3 Related work

The primary line of research in Colonel Blotto games focuses on characterizing the equilibria for a given competitive environment. Since Borel's initial study, many works have advanced this thread (see, e.g., [36, 37, 20, 38, 39, 40, 41]). However, analytical solutions to the most general settings remain as open problems. As such, there are several variations of Colonel Blotto games that have been studied extensively. Of these variations, General Lotto games [42, 21, 20, 43] are the most popular. Notably, the players' equilibrium payoffs in General Lotto games have been fully characterized [21, 20]. Due to its tractability, the General Lotto game is often adopted in studies of more complex adversarial environments, including engineering domains such as network security [44, 45, 46] and the security of cyber-physical systems [47, 48, 49].

Our work on pre-emption relates to recent threads in the literature on two-stage General Lotto game models, where players have the option to publicly announce their strategic intentions ahead of play. Relevant to our study of concessions in Section 5.1 are [48, 49, 50], in which the authors propose game models in which players facing a common adversary have the opportunity to negotiate an alliance that takes the form of a pre-emptive, unilateral budget transfer between the players. However, mechanisms requiring mutual coordination – including alliances – are often not practical, as the necessary channels for coordination between the players may not be available, or players' budgets may not be directly transferable. This further motivates our study of concessions, as they model decisions that do not require mutual coordination. Our study of pre-allocation in Section 5.2 draws significant inspiration from the *General Lotto game with favoritism* (GL-F) proposed in [51]. Favoritism refers to the fact that pre-allocated resources provide an inherent advantage to one player's competitive chances over the other's. Their work establishes existence of equilibria and develops computational methods to calculate them to arbitrary precision. However, this prior work considers pre-allocated resources as exogenous parameters of the game. In contrast, we model the deployment of pre-allocated resources as a strategic element of the competitive interaction. Furthermore, we provide the first analytical characterization of equilibria and the corresponding payoffs in GL-F games. Finally, our study of strategic concessions and pre-allocations falls under a larger research thread on the potential benefits associated with revealing information to an adversary (see, e.g., [52, 53, 54]).

Chapter 4

Obfuscating mechanisms: The value of concealing information

Beyond the results presented in this chapter, few contributions in the Colonel Blotto literature characterize equilibrium solutions in incomplete information settings [55, 56, 57, 58]. Moreover, the study of incomplete, asymmetric information in the literature is largely unavailable.

In this chapter, we investigate General Lotto games with incomplete and asymmetric information about battlefield values and about players' budget endowments. There are many motivations for considering problem settings with informational asymmetries. From an analytical perspective, characterizing equilibria and the corresponding equilibrium payoffs in such scenarios directly identifies the "value of information" pertaining to these vital pieces of information. That is, how does knowledge (or lack thereof) about the strategic objectives or about an opponent's strength impact the viable strategies and resulting performance? On the other hand, this characterization can also be leveraged to inform how one should obfuscate information in competitive scenarios. For instance, what is the return on investment for concealing information related to the value of the various strategic objectives and one's overall strength?

4.1 The value of battlefield information

4.1.1 Model

We define the following Bayesian framework for *General Lotto games with incomplete value information*: Before play, the vector of batttlefield values $\mathbf{v} \in \mathbb{R}_+^n$ is drawn from a distribution $\mathbf{p}$ that is common knowledge. The distribution is over a finite collection of vectors $\mathcal{V}$, where $\mathbf{v} \in \mathcal{V}$ is drawn with probability $p_\mathbf{v} \geq 0$, and $\sum_{\mathbf{v}\in\mathcal{V}} p_\mathbf{v} = 1$. We will call an element of $\mathcal{V}$ a *state*, and $\mathcal{V}$ the *state space*.

The information available to each player $i \in \{A, B\}$ about the realized state $\mathbf{v}$ can be represented by a mapping from states to types, $\tau_i : \mathcal{V} \to \mathcal{T}_i$. Here, $\mathcal{T}_i$ is the set of types for player i, where $1 \leq |\mathcal{T}_i| \leq |\mathcal{V}|$. The mapping τ_i describes how well i can distinguish between different realized states. For example, if τ_i is a bijection (i.e., player i has a unique type for each state), then player i is *informed* about the state. Otherwise, if $\tau_i(\mathbf{v}) = t$ (i.e., player i has one type for all the states), then player i is *uninformed*, i.e., cannot refine its posterior from the common prior $\mathbf{p}$. As is standard in Bayesian games, the mappings are common knowledge.

Each player is tasked with allocating her resources over the n battlefields. An allocation is a vector $\mathbf{x}_i = (x_{i,b})_{b=1}^B \in \mathbb{R}_+^n$. Given type spaces $\mathcal{T}_i$, an admissible strategy for player i is a tuple of n-variate distributions $F_i = \{F_i^t\}_{t\in\mathcal{T}_i} \in \mathcal{L}(X_i)^{|\mathcal{T}_i|}$, where $F_i^t \in \mathcal{L}(X_i)$ means that

$$\mathbb{E}_{\mathbf{x}_i \sim F_i^t}\left[\sum_{b\in\mathcal{B}} x_{i,b}\right] \leq X_i. \tag{4.1}$$

In words, player i in type $t \in \mathcal{T}_i$ can randomize over any allocation in $\mathbb{R}_+^n$ as long as its budget is not violated *in expectation*. Given state $\mathbf{v} \in \mathcal{V}$ is realized, the type of player i is $t_i = \tau_i(\mathbf{v})$, and the resulting payoff to both players is given by

$$U_i(F_i^{t_i}, F_{-i}^{t_{-i}}; \mathbf{v}) := \sum_{b\in\mathcal{B}} v_b \int_0^\infty F_{-i,b}^{t_{-i}}(x_{i,b}) dF_{i,b}^{t_i} \tag{4.2}$$

where $F_{i,b}^{t_i}$ is the univariate marginal distribution on player i's allocation to battlefield b. The

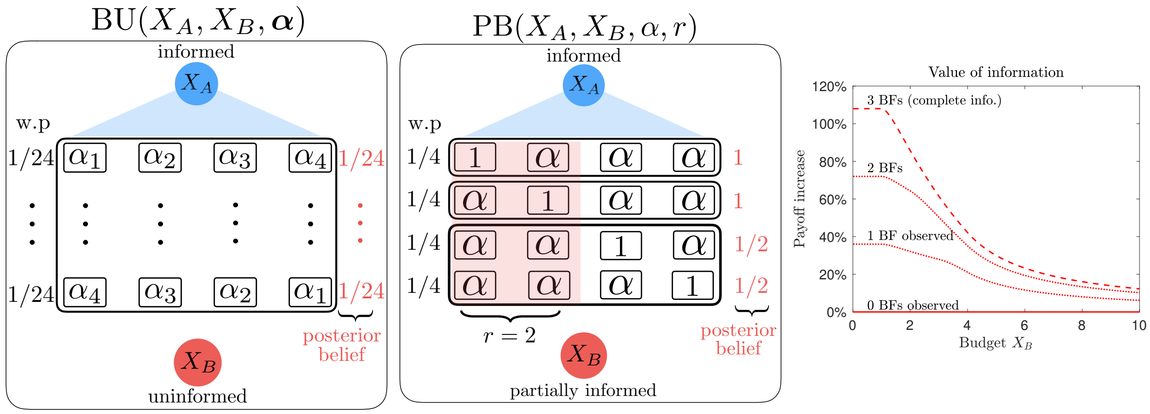

Figure 4.1: **Left:** The game BU(X_A, X_B, α). Each row is a state $\mathbf{v} \in \mathcal{V}$ that represents one equiprobable set of battlefield values. Each state is a permutation of n underlying values $(\alpha_i)_{i=1}^n$. Player A observes the realization, whereas player B does not. In this example, there are four underlying values, and there are $4! = 24$ states. **Middle:** The game PB(X_A, X_B, α, r). Given there are only two distinct underlying values, 1 and α, and hence there are only four distinct permutations (states). In this depiction, player B can observe the first two battlefields (of four) from the realized state. It thus refines its posterior belief on the state realization. **Right:** Percent increase in payoff, $100 \times (\frac{\pi_B(X_A,X_B,\alpha,r)}{\pi_B(X_A,X_B,\alpha,0)} - 1)$, for player B from observing $r = \{0, 1, 2, 3\}$ battlefields when $n = 4$. Here we fix $X_A = 1$ and $\alpha = 0.1$. The percent increase is explicitly given in Corollary 4.1.2 for smaller budgets $X_B < X_A$.

integral term is the probability that player i allocates more resources to b. The ex-ante expected payoff to player i is thus given by

$$\Pi_i(F_i, F_{-i}) := \sum_{\mathbf{v} \in \mathcal{V}} p_{\mathbf{v}} \cdot U_i(F_i^{\tau_i(\mathbf{v})}, F_{-i}^{\tau_{-i}(\mathbf{v})}; \mathbf{v}). \tag{4.3}$$

An *equilibrium* is a strategy profile (F_i^*, F_{-i}^*) such that

$$\Pi_i(F_i^*, F_{-i}^*) \geq \Pi_i(F_i, F_{-i}^*) \tag{4.4}$$

for each $i \in \{A, B\}$ and any $F_i \in \mathcal{L}(X_i)^{|\mathcal{T}_i|}$.

4.1.2 Completely asymmetric information

While an analytical characterization of equilibrium is desirable for arbitrary state spaces $\mathcal{V}$, prior $\mathbf{p}$, and information structures τ_A, τ_B, such a task proves difficult to accomplish. Nonetheless, we provide complete analytical equilibrium characterizations given that the following assumptions hold.

Assumption 4.1.1. *Player A is informed and player B is uninformed.*

Thus, we consider the completely asymmetric information setting. We will also consider state spaces with the following structure.

Assumption 4.1.2. *Given n underlying values $\alpha = (\alpha_1, \ldots, \alpha_n)$ with $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_n$, the state space $\mathcal{V}$ is comprised of all possible permutations of the elements in α. Moreover, the prior $\mathbf{p}$ is uniform.*

Assumption 4.1.2 asserts that the players always compete over the same set of n values, but their orientation is different in each state of the world.[1] Such scenarios can arise when a defender (player A) attempts to obfuscate the location of its valuable assets, which are targeted by an attacker (player B). An illustration of our setup is depicted in Figure 4.1 (left).

We refer to this class of games as *General Lotto games with asymmetric value information* (GL-V), and denote an instance with GL-V(X_A, X_B, α). We will denote the equilibrium payoff to player i as $\pi_i(X_A, X_B, \alpha)$, where it holds that $\pi_B = \sum_{i=1}^n \alpha_i - \pi_A$. A complete characterization of the players' equilibrium payoffs in the GL-V game is given in the main result below.

Theorem 4.1.1. *The unique equilibrium payoff $\pi_B(X_A, X_B, \alpha)$ to player B (uninformed) in the game GL-V(X_A, X_B, α) is given as follows. If $X_B \in [0, X_A)$, it is*

$$\frac{X_B}{2X_A} \sum_{i=1}^{n} \frac{\alpha_i}{n} (2i - 1) \,. \tag{4.5}$$

[1] Assumption 4.1.2 is also equivalent to considering the state space with n states, where the elements of α are simply shifted one position to the right (modulo n) in each subsequent state. This equivalence holds because each battlefield b is still equally likely to hold the value α_k. Setups where the type mappings are arbitrary, the values do not have the permutation structure of Assumption 4.1.2, or the distribution on states is not uniform, prove difficult to analyze with the methods developed in this paper. Such generalizations are left for future work.

If $X_B \in \left[\frac{nX_A}{k}, \frac{nX_A}{k-1}\right)$ *($k \in [n]$), then it is*

$$\begin{aligned} &\sum_{i=k}^{n} \alpha_i + \frac{X_B}{2X_A} \sum_{i=1}^{k-1} \frac{\alpha_i}{n} \left(1 + 2 \sum_{j=i+1}^{k-1} \frac{\alpha_j}{\alpha_i} \right) \\ &- \frac{n\alpha_k X_A}{2X_B} \left(1 - \frac{k-1}{n} \frac{X_B}{X_A} \right)^2 . \end{aligned} \tag{4.6}$$

The unique equilibrium payoff to player A is $\pi_A = \sum_{i\in[n]} \alpha_i - \pi_B$.

Comparison with complete information. We compare the players' equilibrium payoffs $\pi_A(G)$, $\pi_B(G)$ in the GL-V game with their payoffs in the nominal, complete information game, where players A and B play the General Lotto game $\text{GL}(X_A, X_B, \alpha)$. Player A's equilibrium payoff in the nominal game is provided in Fact 3.1.1, where $\Phi = \sum_k \alpha_k$. We show in the next result that the uninformed player can receive n times more payoff in the complete information game.

Corollary 4.1.1. *Player B can experience an n-fold performance improvement by acquiring complete information about the battlefield values. Specifically, for* $X_B < X_A$*,*

$$\max_{\alpha\in\mathbb{R}^n_{++}} \frac{\pi_B^{CI}(X_A, X_B)}{\pi_B(X_A, X_B, \alpha)} = n.$$

Proof. Player B's complete information payoff is $\pi_B^{\text{CI}} = \frac{X_B}{2X_A} \sum_{i=1}^{n} \alpha_i$ (see Fact 3.1.1). When it has incomplete information, the payoff is given by (4.5). The ratio $\frac{\pi_B^{\text{CI}}}{\pi_B}$ takes its highest value of n when all value is concentrated at $\alpha_1 > 0$, i.e. $\alpha_2 = \cdots = \alpha_n = 0$. □

4.1.3 Partial battlefield uncertainty

We next consider scenarios where player B is partially informed about the state – it can observe $r \in \{0, 1, \ldots, n-1\}$ of the realized battlefield valuations, while player A remains fully informed. Note that Assumption 4.1.1 does not hold if $r \neq 0$. To formulate such a scenario, we consider a simplified setting with two underlying values: one priority type[2] of value $\alpha_1 = 1$, and $n-1$ non-priority types each with value $\alpha \equiv \alpha_2 = \cdots = \alpha_n \in [0,1]$. Considering all permutations of the values, there are thus n distinct and equiprobable states of the world. Each state of the world, $i \in [n]$, takes the form $\mathbf{v}^i = (\alpha, \ldots, 1, \ldots, \alpha)$ where the 1 is in the i-th component (see Figure 4.1)b.

Formally, player B has $r+1$ types $\mathcal{T}_B = \{t_1, \ldots, t_r, t_{r+1}\}$. The type mapping τ_B is given by

$$\tau_B(\mathbf{v}^i) = \begin{cases} t_i, & \text{if } i \in [r] \\ t_{r+1}, & \text{if } i \in \{r+1, \ldots, n\} \end{cases}. \tag{4.7}$$

In words, if the priority valuation belongs to one of the first r battlefields (w.p. $\frac{r}{n}$), player B is fully informed about the state (it is one of $\{\mathbf{v}^1, \ldots, \mathbf{v}^r\}$), and the General Lotto game becomes one of complete information. In the event that the r valuations player B observes are non-priority (w.p. $1 - \frac{r}{n}$), player B holds a uniform posterior belief on the $n-r$ states $\{\mathbf{v}^{r+1}, \ldots, \mathbf{v}^n\}$. A strategy for player A belongs to $\mathcal{L}(X_A)^{n-r}$, and a strategy for player B belongs to $\mathcal{L}(X_B)$. Their payoffs are thus evaluated according to (4.3) with $\{\mathbf{v}^{r+1}, \ldots, \mathbf{v}^n\}$ as the state space with a uniform prior.

We refer to this class of games as *General Lotto games with partial value information*, and denote an instance as GL-PV(X_A, X_B, α, r). A diagram of this setup is given in Figure 4.1 (middle). Note the class of BU games from the previous section covers the case $r = 0$ for a wider range of possible battlefield values. We completely characterize the equilibrium payoffs $\pi_i(X_A, X_B, \alpha, r)$ in the main result below.

Theorem 4.1.2. *The unique equilibrium payoff $\pi_B(X_A, X_B, \alpha, r)$ to player B in the game GL-PV(X_A, X_B, α, r) is given as follows. If $X_B \in [0, X_A]$, then it is*

$$\frac{r}{n}\pi_B^{CI} + \frac{n-r}{2\gamma n}\left(\frac{\phi - r\alpha}{n-r} + (n-1)\alpha\right). \tag{4.8}$$

[2]The value of 1 is without loss of generality. If there are only two distinct values, e.g. $\beta \geq \alpha$, then one can always normalize by β.

where $\gamma := X_A/X_B$. *If* $X_B \in (X_A, (n-r)X_A]$, *then it is*

$$\frac{r}{n}\pi_B^{CI} + \frac{n-r}{n}\left[\frac{\gamma^{-1}}{2(n-r)} + \alpha\left(n-1-\frac{\gamma}{2}(n-r)(1-\frac{\gamma^{-1}}{n-1})^2\right)\right]. \tag{4.9}$$

If $X_B > (n-r)X_A$, *then it is*

$$\frac{r}{n}\pi_B^{CI} + \frac{n-r}{n}\left(\phi - (n-(1-\alpha)r)\frac{\gamma}{2}\right). \tag{4.10}$$

Here, we denote $\phi := 1 + (n-1)\alpha$. *The unique equilibrium payoff to player* A *is* $\pi_A = \phi - \pi_B$.

For low budgets, we can analytically characterize player B's payoff improvement compared to when it observes no battlefields (also see Figure 4.1 (right) for a numerical plot).

Corollary 4.1.2. *For* $X_B < X_A$, *the performance improvement factor for player* B *as a result of acquiring information about* $r \le n-1$ *battlefield valuations is*

$$\frac{\pi_B(X_A, X_B, \alpha, r)}{\pi_B(X_A, X_B, \alpha, 0)} = 1 + r \cdot \frac{1-\alpha}{(1-\alpha) + \alpha n^2}. \tag{4.11}$$

For each additional battlefield observed, the improvement increases by a constant linear factor. Figure 4.1 (right) also suggests that the improvement factor decreases to 1 as X_B grows large. In other words, it appears that battlefield information is less valuable to stronger players.

4.2 The value of budgetary information

4.2.1 Model

In our formulation of the *General Lotto game with asymmetric budget information* (GL-B), player A's true budget is drawn from a common (Bernoulli) prior probability distribution, where player B has knowledge of the prior distribution but does not observe the realization of player A's budget. Specifically, player A's budget is of the high type ($X_A = X_H$) with probability p, and of the low type ($X_A = X_L$) with probability $1-p$, where $X_H \ge X_L \ge 0$. Player B's budget X_B is fixed and is common knowledge. Additionally, the games we consider have a single battlefield of value 1.[3] A formal definition is given below.

Definition 4.2.1. *An instance of the GL-B game is specified by four parameters: $X_H \geq 0, X_L \geq 0, X_B \geq 0$, and $p \in [0,1]$, such that $X_H \geq X_L$. Here, t_H indicates the "high" budget type, occurring with probability p, and t_L is the "low" budget type, occurring with probability $1-p$. A strategy for player A is a pair of distributions $F_A := (F_A(t_H) \in \mathcal{L}(X_H), F_A(t_L) \in \mathcal{L}(X_L))$. A strategy for player B is a single distribution $F_B \in \mathcal{L}(X_B)$. The expected payoff to each player is given by*

$$\begin{aligned}\Pi_A(F_A, F_B) &:= p \cdot U_A(F_A(t_H), F_B) + (1-p) \cdot U_A(F_A(t_L), F_B) \\ \Pi_B(F_B, F_A) &:= p \cdot U_B(F_B, F_A(t_H)) + (1-p) \cdot U_B(F_B, F_A(t_L))\end{aligned} \tag{4.12}$$

where it holds that $\Pi_B(F_B, F_A) = 1 - \Pi_A(F_A, F_B)$. We will often refer to player A as the informed player*, and to player B as the* uninformed player*. Let the tuple $G = (X_H, X_L, X_B, p)$ represent one instance of this game, which we will sometimes refer to as GL-B(X_H, X_L, X_B, p) (or as GL-B when the context is clear).*

4.2.2 Complete characterization of equilibrium payoffs.

Our goal is to evaluate the equilibrium payoffs for any given GL-B game. Throughout the paper, we will denote the *budget ratios* as $\gamma_i := \frac{X_i}{X_B}$, $i \in \{H, L\}$, and define $\bar{\gamma} := p\gamma_H + (1-p)\gamma_L$, when convenient.

Our complete characterization of the players' equilibrium payoffs in the GL-B game is given as follows.

Theorem 4.2.1. *Consider a GL-B game $G = (X_H, X_L, X_B, p)$ with a single battlefield with*

[3]Recall that a General Lotto game over a single battlefield with value 1 is mathematically equivalent to any General Lotto game with n battlefields that have cumulative value of 1.

value 1. Player A's equilibrium payoff is

$$\pi_A(G) = \begin{cases} \frac{\bar{\gamma}}{2}, & \text{if } (\gamma_H, \gamma_L) \in \mathcal{R}_1 \\ 1 - \frac{1}{2\bar{\gamma}}, & \text{if } (\gamma_H, \gamma_L) \in \mathcal{R}_2 \\ p + (1-p)\left(1 - \frac{1}{2\gamma_L}\right), & \text{if } (\gamma_H, \gamma_L) \in \mathcal{R}_3 \\ p + (1-p)\frac{\gamma_L}{2}, & \text{if } (\gamma_H, \gamma_L) \in \mathcal{R}_4 \\ p + (1-p)\frac{\gamma_L}{\gamma_H} + \frac{\sqrt{\bar{\gamma}(\bar{\gamma} - p\gamma_H)}}{\gamma_H} - \frac{\left(\sqrt{(1-p)\gamma_L} + \sqrt{\bar{\gamma}}\right)^2}{2\gamma_H^2}, & \text{if } (\gamma_H, \gamma_L) \in \mathcal{R}_5 \end{cases} \tag{4.13}$$

where $\mathcal{R}_k$, $k = 1, \ldots, 5$, are disjoint subsets of $\mathcal{R} := \{(\gamma_H, \gamma_L) \in \mathbb{R}_+^2 : \gamma_H \geq \gamma_L\}$, given by

$$\begin{aligned} \mathcal{R}_1 &= \{(\gamma_H, \gamma_L) \in \mathcal{R} : \bar{\gamma} \leq 1\} \setminus \mathcal{R}_5 \\ \mathcal{R}_2 &= \left\{(\gamma_H, \gamma_L) \in \mathcal{R} : \bar{\gamma} \geq 1 \text{ and } \gamma_L \geq \frac{1-p}{2-p}\right\} \\ \mathcal{R}_3 &= \left\{(\gamma_H, \gamma_L) \in \mathcal{R} : \gamma_H \geq 2 + \frac{p}{1-p} \text{ and } 1 \leq \gamma_L \leq \frac{1-p}{2-p}\gamma_H\right\} \\ \mathcal{R}_4 &= \left\{(\gamma_H, \gamma_L) \in \mathcal{R} : \gamma_H \geq 2 + \frac{p}{1-p} \text{ and } \frac{p}{(1-p)(\gamma_H - 2)} \leq \gamma_L \leq 1\right\} \\ \mathcal{R}_5 &= \{(\gamma_H, \gamma_L) \in \mathcal{R} : \gamma_L \leq G(\gamma_H)\} \end{aligned} \tag{4.14}$$

with $G(\gamma_H)$ defined as

$$G(\gamma_H) := \begin{cases} 0, & \text{if } \gamma_H \in [0, 1) \\ \frac{p(\gamma_H - 1)^2}{(1-p)(2-\gamma_H)}, & \text{if } \gamma_H \in [1, 2-p] \\ \frac{1-p}{2-p}\gamma_H, & \text{if } \gamma_H \in \left(2-p, 2 + \frac{p}{1-p}\right] \\ \frac{p}{(1-p)(\gamma_H - 2)}, & \text{if } \gamma_H > 2 + \frac{p}{1-p} \end{cases} \tag{4.15}$$

In the proof of Theorem 4.2.1, we establish a connection to two-player all-pay auctions with asymmetric information, and leverage solutions to such auctions to derive a system of non-linear equations associated with the Lotto budget constraints (3.4). Note that such solutions can only be applied to the sub-region $\mathcal{R}_5$. The equilibrium characterization for the remaining four regions is then done *ad hoc*.

Comparison with complete information. We identify the effect of asymmetric budget information by comparing the players' equilibrium payoffs $\pi_A(G)$, $\pi_B(G)$ in the GL-B game with their payoffs in the nominal, complete information game, where players A and B play the (complete information) General Lotto game $\text{GL}(X_H, X_B, \mathbf{v})$ with probability p and the game $\text{GL}(X_L, X_B, \mathbf{v})$ with probability $1-p$. Player A's equilibrium payoff in the nominal game is provided in [21]:

$$\pi_A^{\text{CI}}(G) := \Phi \cdot [pL(X_H, X_B) + (1-p)L(X_L, X_B)] \tag{4.16}$$

where $\Phi = \sum_{b=1}^{B}[v_b] = 1$ and

$$L(X,Y) = \begin{cases} X/(2Y) \quad \text{if } X \leq Y, 1 - Y/(2X) \quad \text{if } X > Y. \end{cases} \tag{4.17}$$

Corollary 4.2.1. *Given any GL-B game G, we have that $\pi_A(G) = \pi_A^{CI}(G)$ when $(\gamma_H, \gamma_L) \in \mathcal{R}_1 \cap \{(\gamma_H, \gamma_L) \in \mathbb{R}_+^2$ s.t. $\gamma_L \leq \gamma_H \leq 1\}$, or when $(\gamma_H, \gamma_L) \in \mathcal{R}_2 \cap \{(\gamma_H, \gamma_L) \in \mathbb{R}_+^2$ s.t. $\gamma_H \geq \gamma_L \geq 1\}$. Otherwise, we have that $\pi_A(G) > \pi_A^{CI}(G)$.*

Comparison with other prior distributions. Previously, we identified the value of budget information when player A's budget is governed by a Bernoulli prior distribution. We may also consider the setting where player A is endowed with some budget X_A, and may choose *any* prior distribution $f_A \in \Delta(\mathbb{R}_+)$ such that $\mathbb{E}_{x \sim f_A}[x] \leq X_A$. It is straightforward to show that an optimal prior distribution is simply player A's equilibrium strategy from the General Lotto game $\text{GL}(X_A, X_B, \mathbf{v})$. Therefore, player A's equilibrium payoff under her optimal prior distribution is $\pi_A^{\text{opt}}(G) = \Phi \cdot L(X_A, X_B)$, where $X_A = pX_H + (1-p)X_L$.

From the characterization of the equilibrium payoffs in Theorem 4.2.1, we observe that player A does not benefit from further obfuscating her budget (i.e., $\pi_A(G) = \pi_A^{\text{opt}}(G)$) if $(\gamma_H, \gamma_L) \in \mathcal{R}_1 \cup \mathcal{R}_2$. Otherwise, player A's equilibrium payoff under the optimal prior distribution is strictly greater than under the Bernoulli distribution. We formally state this observation in the following:

Corollary 4.2.2. *Given any GL-B game G, we have that $\pi_A(G) = \pi_A^{\text{opt}}(G)$ when $(\gamma_H, \gamma_L) \in \mathcal{R}_1 \cup \mathcal{R}_2$. Otherwise, we have that $\pi_A(G) < \pi_A^{\text{opt}}(G)$.*

4.3 Chapter proofs

This section contains the proofs of the main results in this chapter.

4.3.1 Proofs from Section 4.1

Proof of Theorem 4.1.1. Let $s \in [n!]$ be an indexing of the state space $\mathcal{V}$, which is all permutations of the vector α (Assumption 4.1.2). Also, denote v_b^s as the value of battlefield b in state s and $F_{A,b}^s$ as A's marginal distribution on allocations to battlefield b in state s. We will consider strategies for player A that satisfy $F_{A,b}^s = F_{A,b'}^{s'}$ whenever $v_b^s = v_{b'}^{s'} = \alpha_i$ (for some $i \in [n]$). That is, it applies the same marginal distribution for the battlefields identified with the underlying value α_i. We will henceforth refer to such strategies as *symmetric*, denoting $F_A \in \mathcal{S}_A(X_A)$. Such a strategy is thus represented as a tuple of n univariate marginals $\{F_A^i\}_{i\in[n]} \in \mathcal{L}(X_A)^n$, where player A employs F_A^i on the battlefield identified with α_i in any state s.

Similarly, we will consider symmetric strategies $F_B \in \mathcal{S}_B(X_B)$ for player B that satisfy $F_{B,b} = F_{B,b'}$ for every $b, b' \in [n]$, i.e. it uses the same distribution for every battlefield. It is represented with a *single* univariate distribution F_B (i.e. without explicitly referring to marginals $F_{B,b}$). The following result provides a method to find an equilibrium of BU games.

Lemma 4.3.1. *Consider the game $BU(X_A, X_B, \alpha)$. Suppose $\lambda = (\lambda_A, \lambda_B) \in \mathbb{R}_+^2$ solves the following system of equations*

$$X_A = \sum_{i=1}^{\bar{k}(\lambda)} \int_{T_{i+1}}^{T_i} \frac{n\lambda_B}{\alpha_i} x \, dx \text{ and } X_B = \sum_{i=1}^{\bar{k}(\lambda)} \int_{T_{i+1}}^{T_i} \frac{n\lambda_A}{\alpha_i} x \, dx \tag{4.18}$$

where

$$\bar{k}(\lambda) := \begin{cases} k, & \text{if } \frac{n}{k} \le \frac{\lambda_A}{\lambda_B} < \frac{n}{k-1},\ k \in [n] \\ n+1, & \text{if } \frac{\lambda_A}{\lambda_B} < 1 \end{cases} \tag{4.19}$$

and T_i are defined as follows. If $\bar{k} \in [n]$,

$$T_i := \begin{cases} \sum_{j=i}^{\bar{k}-1} \frac{\alpha_j}{n\lambda_B} + \frac{\alpha_{\bar{k}}}{\lambda_A}\left(1 - \frac{\lambda_A(\bar{k}-1)}{n\lambda_B}\right), & \text{if } i \in [\bar{k}] \\ 0, & \text{if } i > \bar{k} \end{cases} \tag{4.20}$$

If $\bar{k} = n+1$, $T_i := \sum_{j=i}^{n} \frac{\alpha_j}{n\lambda_B}$. Then an equilibrium $(F_A, F_B) \in \mathcal{S}_A(X_A) \times \mathcal{S}_B(X_B)$ of $BU(X_A, X_B, \alpha)$ is given as follows. For player A, F_A^i ($i \in [n]$) is given by

$$\begin{cases} \mathit{Unif}(T_{i+1}, T_i), & \text{if } i \in [\bar{k}-1] \\ \left(1 - \frac{n\lambda_B T_{\bar{k}}}{\alpha_{\bar{k}}}\right)\delta_0 + \frac{n\lambda_B T_{\bar{k}}}{\alpha_{\bar{k}}}\,\mathit{Unif}(0, T_{\bar{k}}), & \text{if } i = \bar{k} \\ \delta_0, & \text{if } i > \bar{k} \end{cases} \tag{4.21}$$

and F_B is given by

$$\begin{cases} \sum_{i=1}^{\bar{k}} \frac{\lambda_A(T_i - T_{i+1})}{\alpha_i}\,\mathit{Unif}(T_{i+1}, T_i), & \text{if } \bar{k} \in [n] \\ \left(1 - \frac{\lambda_A}{\lambda_B}\right)\delta_0 + \sum_{i=1}^{n} \frac{\lambda_A(T_i - T_{i+1})}{\alpha_i}\,\mathit{Unif}(T_{i+1}, T_i), & \text{if } \bar{k} = n+1 \end{cases} \tag{4.22}$$

The equilibrium payoff to player A is

$$\lambda_A X_A + \sum_{i \in [\bar{k}-1]} \left((\frac{\alpha_i}{\alpha_{\bar{k}}} - 1)\lambda_A T_{\bar{k}} + \sum_{k=i+1}^{\bar{k}-1} \frac{\sigma(\alpha_i - \alpha_k)}{n} \right) \tag{4.23}$$

where $\sigma := \frac{\lambda_A}{\lambda_B}$. The equilibrium payoff to player B is

$$\lambda_B X_B + \alpha_{\bar{k}} - n\lambda_B T_{\bar{k}} + \sum_{i=\bar{k}+1}^{n} \alpha_i. \tag{4.24}$$

Before proceeding with the proof, a few remarks are in order. The system of equations (4.18) corresponds to the expected budget constraints for both players (4.1), where each term in the sum is the expected allocation to the battlefield identified with α_i under the strategies F_A, F_B (4.21), (4.22). For a solution λ, player A competes on the most valuable $\bar{k}(\lambda)$ battlefields.

The marginals of F_A, F_B are precisely the equilibrium bidding strategies of a two-player all-pay auction with asymmetric information [59]. It can be shown that the necessary condition for equilibrium in the BU game coincides with that of n independent all-pay auctions – each corresponding to one of the n battlefields [60]. The connection between all-pay auctions and the equilibria of General Lotto games is often utilized in the literature [39, 20, 60, 51]. Lastly, the sum of the equilibrium payoffs (4.23), (4.24) is $\sum_{i=1}^n \alpha_i$ (constant-sum game).

Proof. We show that the marginals (4.21) and (4.22) constitute an equilibrium profile for $\text{BU}(X_A, X_B, \alpha)$. For space concerns, we provide the proof for the case that a solution λ exists when $\bar{k} \in [n]$ (the case $\bar{k} = n+1$ follows similar arguments). The approach is to show that F_A is a best-response to F_B, and then vice versa. Suppose $\bar{F}_A \in \mathcal{S}_A(X_A)$ is another symmetric[4] strategy for A. Player A's expected payoff against F_B then becomes identical regardless of the state. We can thus write $\Pi_A(\bar{F}_A, F_B)$ (4.3) as

$$\sum_{i\in[n]} C^i := \sum_{i\in[n]} \alpha_i \int_0^\infty F_B(x)\, d\bar{F}_A^i. \tag{4.25}$$

The approach is to derive a universal upper bound of the above expression for any $\bar{F}_A \in \mathcal{S}_A(X_A)$, and then show that F_A (4.21) meets the upper bound with equality. Using (4.22), and denoting $\sigma := \frac{\lambda_A}{\lambda_B}$ for compactness, we can write C_j as

$$\begin{aligned} C^i = &\frac{\lambda_A \alpha_i}{\alpha_{\bar{k}}} \int_0^{T_{\bar{k}}} x\, d\bar{F}_A^i + \alpha_i \int_{T_1}^\infty d\bar{F}_A^i \\ &+ \alpha_i \sum_{j=1}^{\bar{k}-1} \int_{T_{j+1}}^{T_j} (1 - \frac{j}{n}\sigma + \frac{\lambda_A}{\alpha_j}(x - T_{j+1}))\, d\bar{F}_A^i \end{aligned} \tag{4.26}$$

[4]Here, deviations can be restricted to $\mathcal{S}_A(X_A)$ in place of the more general space $\mathcal{L}(X_A)^n$, without loss of generality. Intuitively, this is due to the symmetry in the arrangement of the battlefields in our setup, and because F_B itself is a symmetric strategy for B.

We can write C^i as

$$\begin{aligned}
&\lambda_A \int_0^{T_1} x\, d\bar{F}_A^i + \alpha_i \int_{T_1}^{\infty} d\bar{F}_A^i \\
&+ \sum_{j=i+1}^{\bar{k}} (\frac{\alpha_i}{\alpha_j} - 1)\lambda_A \int_{T_{j+1}}^{T_j} x\, d\bar{F}_A^i - \sum_{j=1}^{\min(i,\bar{k})} (1 - \frac{\alpha_i}{\alpha_j})\lambda_A \int_{T_{j+1}}^{T_j} x\, d\bar{F}_A^i \\
&+ \sum_{j=1}^{\bar{k}-1} (\alpha_i (1 - \frac{j}{n}\sigma) - \frac{\alpha_i}{\alpha_j}\lambda_A T_{j+1}) \int_{T_{j+1}}^{T_j} d\bar{F}_A^i
\end{aligned} \tag{4.27}$$

It will be instructive to collect the terms above with respect to each interval $[T_{j+1}, T_j]$. For $j \in \{1, \ldots, i\}$, the term is $\alpha_i((1 - \frac{j}{n}\sigma) - \frac{\lambda_A T_{j+1}}{\alpha_j}) \int_{T_{j+1}}^{T_j} d\bar{F}_A^i - (1 - \frac{\alpha_i}{\alpha_j})\lambda_A \int_{T_{j+1}}^{T_j} x\, d\bar{F}_A^i$. It can be upper bounded by

$$\left((\frac{\alpha_i}{\alpha_{\bar{k}}} - 1)\lambda_A T_{\bar{k}} + \frac{\sigma}{n} \sum_{k=j+1}^{\bar{k}-1} (\alpha_i - \alpha_k) \right) \int_{T_{j+1}}^{T_j} d\bar{F}_A^i \tag{4.28}$$

The upper bound is due to an application of Markov's inequality, i.e. $-\int_a^b x\, dF \leq -a \int_a^b dF$ for any interval $[a, b]$ and distribution F. Let us denote the upper bound above as B_i. A similar bound can be derived for the $[T_{j+1}, T_j]$ terms when $j \in \{i + 1, \ldots, \bar{k}\}$:

$$\left((\frac{\alpha_i}{\alpha_{\bar{k}}} - 1)\lambda_A T_{\bar{k}} + \frac{\sigma}{n} \sum_{k=j}^{\bar{k}-1} (\alpha_i - \alpha_k) \right) \int_{T_{j+1}}^{T_j} d\bar{F}_A^i \tag{4.29}$$

Let us denote b_j^i, $j \in [\bar{k}]$, as the coefficient in the upper bounds (4.28),(4.29). Additionally, we can write the first two terms of (4.27), $\lambda_A \int_0^{T_1} d\bar{F}_A^i + \alpha_i \int_{T_1}^{\infty} d\bar{F}_A^i$, as

$$\begin{aligned}
&\lambda_A X_A^i + \alpha_i \int_{T_1}^{\infty} d\bar{F}_A^i - \lambda_A \int_{T_1}^{\infty} x\, d\bar{F}_A^i \\
&\leq \lambda_A X_A^i + (\alpha_i - \lambda_A T_1) \int_{T_1}^{\infty} d\bar{F}_A^i
\end{aligned} \tag{4.30}$$

where $X_A^i := \int_0^\infty x \, d\bar{F}_A^i$. The expression $d^i := \alpha_i - \lambda T_1$, $i \in [n]$, can be written as

$$\begin{cases} (\frac{\alpha_i}{\alpha_{\bar{k}}} - 1)\lambda_A T_{\bar{k}} + \frac{\sigma}{n}\sum_{k \neq i}^{\bar{k}-1}(\alpha_i - \alpha_k), & \text{if } j \in [\bar{k}] \\ \alpha_i - \alpha_{\bar{k}} + \frac{\sigma}{n}\sum_{k=1}^{\bar{k}-1}(\alpha_{\bar{k}} - \alpha_k), & \text{if } j > \bar{k} \end{cases} \tag{4.31}$$

Player A's expected payoff $\Pi_A(\bar{F}_A, F_B)$ can then be upper bounded by

$$\lambda_A X_A + \sum_{i \in [n]} \left(d^i \int_{T_1}^{\infty} d\bar{F}_A^i + \sum_{j=1}^{\bar{k}} b_j^i \int_{T_{j+1}}^{T_j} d\bar{F}_A^i \right) \tag{4.32}$$

where we used $\sum_{j=1}^{\bar{k}} \int_0^{T_1} x \, d\bar{F}_A^i = X_A - \sum_{j=1}^{\bar{k}} \int_{T_1}^\infty x \, d\bar{F}_A^i$. For any $i > \bar{k}$, we have $b_j^i < 0$ and $d^i < 0$, $\forall\, j \in [\bar{k}]$. For $i \in [\bar{k}]$, we observe that $b_i^i = \max_{j \in [\bar{k}]} b_j^i$, and that $d_i \leq b_i^i$ with equality if and only if $i = 1$. Therefore, the expression (4.32) is maximized to $\lambda_A X_A + \sum_{i \in [\bar{k}]} b_i^i$ if the mass of $\bar{F}_A^i$, $i \in [\bar{k}]$, is contained in the interval $[T_{i+1}, T_i]$, and each $\bar{F}_A^i$ represents a point mass at zero for $i > \bar{k}$. Note that F_A satisfies all of these properties. One can then show after algebraic steps, that (4.32) with $\bar{F}_A = F_A$ coincides with (4.23). Lastly, F_A satisfies the budget constraint because (λ_A, λ_B) satisfies the first equation of (4.18).

The proof is completed by showing that (4.24) serves as an upper bound for $\Pi_B(F_A, \bar{F}_B)$ for any $\bar{F}_B \in \mathcal{S}(X_B)$, and that $\bar{F}_B = F_B$ achieves the upper bound. The arguments are similar to the procedure when deriving the upper bound for $\Pi_A(\bar{F}_A, F_B)$, and hence we omit the details due to space concerns.

□

Proof of Theorem 4.1.1. From the system of equations (4.18), we can immediately deduce that $\sigma = \frac{\lambda_A}{\lambda_B} = \frac{X_B}{X_A}$. Also, note that $\bar{k}(\lambda)$ depends only on σ. Thus, $\bar{k} \in [n]$ is equivalent to the budget conditions $\frac{X_B}{X_A} \in [\frac{n}{k}, \frac{n}{k-1})$. When $\bar{k} = k \in [n]$, the unique solution of (4.18) is $\lambda_A = \frac{1}{X_A}\left[D_{\bar{k}} + \frac{n\alpha_{\bar{k}}\gamma}{2}\left(1 - \frac{\bar{k}-1}{n\gamma}\right)\left(1 + \frac{\bar{k}-1}{n\gamma}\right)\right]$ and $\lambda_B = \gamma\lambda_A$, where $\gamma := \frac{X_A}{X_B}$ and $D_{\bar{k}} := \sum_{i=1}^{\bar{k}-1} \frac{1}{n\gamma}\left(\frac{\alpha_i}{2} + \sum_{j=i+1}^{\bar{k}-1} \alpha_j\right)$. The expressions for equilibrium payoff in the statement of Theorem 4.1.1 can be recovered from these solutions and Lemma 4.3.1 ((4.23) and (4.24)). □

Proof of Theorem 4.1.2. In the game $\text{PB}(X_A, X_B, \alpha, r)$, player B is informed about the state of the world with probability $\frac{r}{n}$. In this event, its equilibrium payoff is π_B^{CI}. With probability $\frac{n-r}{n}$, player B holds a uniform posterior distribution on the last $n-r$ states, while player A remains completely informed. In this event, we say player B has *partial uncertainty.* The first r battlefields of the state are common knowledge (non-priority α), while player B does not know the location of the priority value 1 among the last $n-r$ battlefields. We thus observe that the first r battlefields are contested with "complete information", and the last $n-r$ battlefields are contested with asymmetric information precisely described by a BU game.

Our approach to deriving an equilibrium to the PB game is to generalize the system of equations (4.18) to the scenario when player B has partial uncertainty. Let us define

$$G_\ell(\lambda;\alpha) := \sum_{i=1}^{\bar{k}(\lambda)} \int_{T_{i+1}}^{T_i} \frac{n\lambda_{-\ell}}{\alpha_i} x\, dx, \quad \ell \in \{A, B\} \tag{4.33}$$

where n is the number of elements in α. Recall in BU games (Lemma 4.3.1), an equilibrium corresponds to the solution $\lambda \in \mathbb{R}_+^2$ of the system $X_A = G_A(\lambda;\alpha)$ and $X_B = G_B(\lambda;\alpha)$. For any λ, α, let us also denote $F_A^{(\lambda,\alpha)} := \{F_A^{i,(\lambda,\alpha)}\}_{i\in[n]}$ as the set of player A's marginals given in (4.21) and $F_B^{(\lambda,\alpha)}$ as player B's marginal given in (4.22).

Lemma 4.3.2. *Consider $PB(X_A, X_B, \alpha, r)$. Suppose $\lambda \in \mathbb{R}_+^2$ is a solution to the system of equations*

$$\begin{aligned} X_A &= G_A(\lambda;\alpha\mathbf{1}_r) + G_A(\lambda;\mathbf{v}_{n-r}) \\ X_B &= G_B(\lambda;\alpha\mathbf{1}_r) + G_B(\lambda;\mathbf{v}_{n-r}) \end{aligned} \tag{4.34}$$

where $\mathbf{1}_r$ is a vector of r ones, and $\mathbf{v}_{n-r} = (1, \alpha\mathbf{1}_{n-r-1})$. Then given player B has partial uncertainty, an equilibrium strategy for A is to deploy the marginals $\{F_A^{i,(\lambda,\alpha\mathbf{1}_r)}\}_{i\in[r]}$ on the first r battlefields, and marginals $\{F_A^{i,(\lambda,\mathbf{v}_{n-r})}\}_{i\in[n-r]}$ corresponding to the last $n-r$ battlefields. An equilibrium strategy for B is to deploy the marginal $F_B^{(\lambda,\alpha\mathbf{1}_r)}$ on each of the first r battlefields, and the marginal $F_B^{(\lambda,\mathbf{v}_{n-r})}$ on each of the last $n-r$ battlefields.

Proof. We note the expression $G_A(\lambda, \alpha\mathbf{1}_r)$ depends on whether $\sigma > 1$ or $\sigma \le 1$, and the expression $G_A(\lambda; \mathbf{v}_{n-r})$ depends on whether $\sigma \leq 1$, $1 < \sigma \leq n-r$, or $\sigma > n-r$. Hence, a solution falls into one of the three latter cases. Moreover, we can deduce from (4.34) that $\sigma = \frac{\lambda_A}{\lambda_B} = \frac{X_B}{X_A}$. We can also calculate player A and B's equilibrium payoffs from the equilibrium marginals generated in the statement of the lemma. We omit the details of the proof because the structure follows similar arguments and calculations from the proof of Lemma 4.3.1. □

4.3.2 Proofs from Section 4.2

This section is devoted to the proof of Theorem 4.2.1 – the characterization of equilibrium payoffs in the Bayesian Lotto game (Definition 4.2.1). To do so, we derive the equilibrium strategies for player A and B in any game instance $G \in \mathcal{G}$. First, we establish a connection between the necessary conditions for equilibrium in two-player all-pay auctions with asymmetric information and the BL game. Second, we leverage equilibrium solutions to such all-pay auctions, provided by [59], to formulate a system of non-linear equations in the Lagrange multipliers λ associated with the players' Lotto expected budget constraints (3.4). Finally, we solve this system of equations, and show that a solution can take one of three different forms, which correspond to disjoint regions in the multiplier space. Hence, solutions to this system are not only algebraic, but also case-dependent.

We completely characterize solutions to these equations in Proposition 4.3.1, finding they exist only for a subset of BL games we are interested in (region $\mathcal{R}_5$). Indeed, the algorithm in [59] constructs equilibrium strategies to the auctions when certain monotonicity conditions are met – players' types are required to be somewhat correlated to their valuations. While these informational requirements hold in the BL games we are interested in, we find the structure of the auction strategies cannot accommodate all possible combinations of budget parameters X_1, X_2, X_B, limiting the applicability of [59] to the sub-region $\mathcal{R}_5 \subset \mathcal{G}$. Nonetheless, we prove that solutions to the system of equations, when they exist, correspond to equilibria of the BL game (Proposition 4.3.1).

We then identify the remaining regions $\mathcal{R}_i$, $i = 1, \ldots, 4$ to have distinct equilibrium structures that cannot be calculated from the aforementioned system of equations. For these equilibria, we combine features of the equilibrium strategies from complete information Lotto games with those that were computed through the system of equations. These details and proofs are given in the Appendix, thus completing the proof of Theorem 4.2.1.

We first review all-pay auctions with asymmetric information, as studied by Siegel in [59].

All-pay auctions with asymmetric information and valuations. In an all-pay auction, two bidders (A and B) compete over a single item. Before bids are submitted, player A privately observes one of two possible types: t_1 (type 1) with a probability p and t_2 (type 2) with $1-p$. Here, it is assumed p is common knowledge and $p > 0$. Player B always observes the same type t_B. Hence, there are two possible type profiles[5], corresponding to whether A observes t_1 or t_2. In type $i \in \{1,2\}$, the players' valuations for the item are $v_{A,i}$ and $v_{B,i}$ for player A and B, respectively.

A (pure) strategy for player A is a pair $\mathbf{x}_A = (x_1, x_2) \in \mathbb{R}^2_+$, where x_i is its bid for the item contingent on receiving type i. A strategy for player B is a non-negative bid $x_B \geq 0$. The resulting payoffs are given by

$$\begin{aligned} &p\left(v_{A,1}\cdot W_A(x_1,x_B)-x_1\right)+(1-p)\left(v_{A,2}\cdot W_A(x_2,x_B)-x_2\right) \quad \text{(Player A)} \\ &p\left(v_{B,1}\cdot W_B(x_B,x_1)-x_B\right)+(1-p)\left(v_{B,2}\cdot W_B(x_B,x_2)-x_B\right) \quad \text{(Player B)} \end{aligned} \tag{4.35}$$

where W_ℓ, $\ell \in \{A,B\}$ is defined in (3.2). A mixed strategy for player A is a pair $\{F_A(t_i)\}_{i=1,2}$, where $F_A(t_i)$ is a univariate probability distribution on $\mathbb{R}_+$. A mixed strategy for B is a single univariate distribution F_B on $\mathbb{R}_+$. The payoffs are calculated as the expected payoffs with respect to the distribution p and the mixed strategies. We refer to this auction as $\text{APA}(v_A, v_B, p)$.

Note that we have adapted this model to the information structure of our BL game. In general, Siegel's model allows for arbitrary, finite type spaces $\mathcal{T}_A$ and $\mathcal{T}_B$ with a joint probability

[5]Siegel's model is general, allowing an arbitrary, finite number of types for each player. We review the model here with two types, since it pertains to our BL game, and for simpler exposition. Indeed, some of our forthcoming results can be generalized to situations where player A has any number of types, while player B still only has one. In particular, systems of equations ($\star$) can be derived with this generality. We can prove their solutions, when they exist, correspond to Bayes-Nash equilibria of BL.

distribution $\mathbf{p}$ on the type profiles. Siegel shows equilibria can be computed algorithmically, provided the following monotonicity condition is met for some ordering of the players' type spaces.

$$\text{For any } t_{-\ell} \in \mathcal{T}_{-\ell},\ v_\ell(t_\ell, t_{-\ell}) \cdot p(t_{-\ell}|t_\ell) \text{ is non-decreasing in } t_\ell \in \mathcal{T}_\ell, \quad \ell \in \{A, B\} \tag{WM}$$

where $p(t_{-\ell}|t_\ell)$ denotes conditional probabilities associated with the joint distribution $\mathbf{p}$, and $v_\ell(t_\ell, t_{-\ell})$ is player ℓ's item valuation in type profile $(t_\ell, t_{-\ell})$. When (WM) holds, a *monotonic*[6] equilibrium in mixed strategies to APA can explicitly be computed through an iterative procedure [59] (we refer to as "Siegel's algorithm"), where the number of iterations is at most $|\mathcal{T}_A| + |\mathcal{T}_B|$. When a strict version of (WM) holds, the constructed monotonic equilibrium is the unique equilibrium of APA. Observe that (WM) is always satisfied if player B only has a single type t_B, and the types of A are ordered according to the valuations $v_A(t_i, t_B)$. The informational structure of our BL game belongs to such "informed-uninformed" settings.

Equilibria of APA have been characterized when (WM) is not met. In fact, [61] provides an algorithm that generalizes Siegel's algorithm to calculate such non-monotonic equilibria. When (WM) is not met, this algorithm can become quite complex. In the "informed-uninformed" scenarios, this algorithm reduces to Siegel's algorithm.

Connection between APA and Bayesian General Lotto games. We now present some informal intuition that suggests a connection between the equilibria of APA and BL. These insights are analogous to those drawn between two-player all-pay auctions with complete information and the Colonel Blotto and General Lotto games [62, 39, 20].

Consider a game instance $G = (X_1, X_2, p, X_B) \in \mathcal{G}$. The Lotto budget constraint (3.4) must hold for the strategies associated in each type. Player A's ex-interim constrained optimization,

[6] An equilibrium is monotonic if any best-response bid for player ℓ in type t_ℓ against the equilibrium strategy of player $-\ell$ is not lower than any best-response bid in a lower type $t \preceq_\ell t_\ell$. When (WM) is not met, non-monotonic equilibria to APA have been characterized in [61]. However, for the central case of interest in this paper for which one player is informed and the other uninformed, (WM) is always met. We leave the characterization of Lotto equilibria when (WM) is not met as an open problem, e.g. players have partial and different information.

given type i is realized, can be written as

$$\max_{\{F_{A,j}(t_i)\}_{j\in[n]}} \sum_{j\in[n]} \int_0^\infty \left[v_j F_{B,j}(x_{A,j}) - \lambda_i x_{A,j}\right] \, dF_{A,j}(t_i) + \lambda_i X_i \tag{4.36}$$

where λ_i is the multiplier on player A's expected budget constraint for type i, and we denote $p_1 = p$ and $p_2 = 1 - p$. Player B's constrained optimization is written as

$$\max_{\{F_{B,j}\}_{j\in[n]}} \sum_{j\in[n]} \sum_{i=1,2} p_i \int_0^\infty \left[v_j F_{A,j}(t_i, x_{B,j}) - \lambda_B x_{B,j}\right] \, dF_{B,j} + \lambda_B X_B. \tag{4.37}$$

where λ_B is the multiplier on player B's budget. The necessary first-order conditions for equilibrium are

$$\begin{aligned} &\frac{d}{dx_{B,j}} \left[\sum_{i=1,2} p_i \left(v_j F_{A,j}(t_i, x_{B,j}) - \lambda_B x_{B,j}\right) \right] = 0 \\ &\frac{d}{dx_{A,j}} \left[v_j F_B(x_{A,j}) - \lambda_i x_{A,j}\right] = 0, \quad i = 1,2 \end{aligned} \tag{4.38}$$

Dividing by the associated (positive) multiplier in each condition, this coincides with the necessary first-order conditions for (Bayesian) equilibrium of n independent two-player all-pay auctions with incomplete information for which the item valuation in auction j for player A in type i is $v_{A,i} = \frac{v_j}{\lambda_i}$, and the valuation for player B when player A's type is i is $v_{B,i} = \frac{v_j}{\lambda_B}$.

The equilibria to each of the n APA games can be computed using Siegel's algorithm, as long as (WM) is satisfied[7]. However, since we do not know the actual ranking of player A's types, i.e. whether $\lambda_1 \leq \lambda_2$ or vice versa, we must proceed by first imposing such a ranking. For the sake of demonstration, let us suppose $\lambda_1 \leq \lambda_2$, so that type 1 is "higher" than type 2, for instance. This allows us to proceed with Siegel's algorithm, which generates bidding distributions $\{F_{A,j}^\lambda(t_i)\}_{i=1,2}$ and $F_{B,j}^\lambda$ for each $j \in [n]$. Here, the superscript indicates the expressions are in terms of the (still) unknown multipliers. These distributions must be consistent with the Lotto expected budget constraints (3.4), yielding a system of three equations

[7] As discussed earlier, the algorithm of [61] can handle APA when (WM) is not met. We do not detail this algorithm here, however, because of its complexity and because Siegel's algorithm suffices for the problems of interest in this paper.

in $\lambda = (\lambda_1, \lambda_2, \lambda_B)$.

$$\begin{aligned} &\sum_{j\in[n]} \mathbb{E}_{x_{A,j}\sim F_{A,j}^{\lambda}(t_i)}[x_{A,j}] = X_i, \quad i = 1,2 \\ &\sum_{j\in[n]} \mathbb{E}_{x_{B,j}\sim F_B^{\lambda}}[x_{B,j}] = X_B \\ &\text{such that } 0 < \lambda_1 \le \lambda_2 \end{aligned} \tag{4.39}$$

Hence, we seek to find a solution to (4.39). In the next sections, we detail the distributions $\{F_A^{\lambda}(t_i)\}_{i=1,2}$ and F_B^{λ} constructed from Siegel's algorithm and apply them to (4.39) to explicitly derive the system of equations.

Equilibrium strategies of APA. In the following, we summarize the resulting equilibria from applying Siegel's algorithm to the APA setup of the previous part.

Define

$$\bar{k} := \begin{cases} 1, & \text{if } p\frac{v_{B,1}}{v_{A,1}} \ge 1 \\ 2, & \text{if } p\frac{v_{B,1}}{v_{A,1}} + (1-p)\frac{v_{B,2}}{v_{A,2}} \ge 1 \text{ and } p\frac{v_{B,1}}{v_{A,1}} < 1 \\ 3, & \text{if } p\frac{v_{B,1}}{v_{A,1}} + (1-p)\frac{v_{B,2}}{v_{A,2}} < 1 \end{cases} \tag{4.40}$$

In brief, $\bar{k}$ is the iteration at which Siegel's algorithm terminates. Denoting $p_1 = p$ and $p_2 = 1-p$, define

$$L_1 = \begin{cases} v_{A,1}, & \text{if } \bar{k} = 1 \\ pv_{B,1}, & \text{if } \bar{k} \in \{2,3\} \end{cases}, \quad L_2 = \begin{cases} 0, & \text{if } \bar{k} = 1 \\ v_{A,2}\left(1 - p\frac{v_{B,1}}{v_{A,1}}\right), & \text{if } \bar{k} = 2 \\ (1-p)v_{B,2}, & \text{if } \bar{k} = 3 \end{cases} \tag{4.41}$$

The L_k are lengths of intervals for which the equilibrium marginals have support. Below, we provide expressions for the equilibrium strategies that result from applying Algorithm 1.

Lemma 4.3.3. *The equilibrium mixed strategies for APA are given as follows*[8]*:*

$$\begin{aligned}
\textit{If } \bar{k}=1: \quad & F_A(t_1) = \left(1-\frac{L_1}{pv_{B,1}}\right)\delta_0 + \frac{L_1}{pv_{B,1}}\,\textit{Unif}(0,L_1), \quad F_A(t_2)=\delta_0 \\
& F_B = \textit{Unif}(0,L_1) \\
\textit{If } \bar{k}=2: \quad & F_A(t_1) = \textit{Unif}(L_2, L_2+L_1) \\
& F_A(t_2) = \left(1-\frac{L_2}{(1-p)v_{B,2}}\right)\delta_0 + \frac{L_2}{(1-p)v_{B,2}}\,\textit{Unif}(0,L_2) \\
& F_B = \frac{L_2}{v_{A,2}}\,\textit{Unif}(0,L_2) + \frac{L_1}{v_{A,1}}\,\textit{Unif}(L_2,L_2+L_1) \\
\textit{If } \bar{k}=3: \quad & F_A(t_1) = \textit{Unif}(L_2,L_2+L_1), \quad F_A(t_2) = \textit{Unif}(0,L_2) \\
& F_B = \left(1-\sum_{i=1}^{2}\frac{L_i}{v_{A,i}}\right)\delta_0 + \frac{L_2}{v_{A,2}}\,\textit{Unif}(0,L_2) + \frac{L_1}{v_{A,1}}\,\textit{Unif}(L_2,L_2+L_1)
\end{aligned} \tag{4.42}$$

In summary, the marginals for player A are uniform distributions with shifted supports, and player B's marginal is a piece-wise uniform distribution.

Equilibria in the $\mathcal{R}_5$ region. We are now ready to apply the methods outlined above to explicitly state the system of equations (4.39). We then completely characterize the solutions to these equations in Proposition 4.3.1. In doing so, we identify the subset of game instances of $\mathcal{G}$ for which solutions to (4.39) exist. Recall this subset was identified as the $\mathcal{R}_5$ region in Theorem 4.2.1 (Figure 4.2). We also prove that such solutions and their associated strategies (constructed from Siegel's algorithm) constitute Bayes-Nash equilibria of the BL game. This serves as the proof of Theorem 4.2.1 in the $\mathcal{R}_5$ region.

The item valuations in one of the n APA games are given by $v_{A,i} = v_j/\lambda_i > 0$ for player A and $v_{B,i} = v_j/\lambda_B > 0$ for player B, in type $i \in \{1,2\}$. Since the high budget X_1 is associated with type t_1, we naturally impose the ranking $\lambda_1 \leq \lambda_2$. Indeed, the distribution $F_A(t_1)$ (4.42) will have a higher expected budget expenditure under this ranking of types.

The value of $\bar{k}$ is not known a priori, as it now depends on the multipliers. The values it

[8]To simplify exposition and notation where convenient, we sometimes explicitly write CDFs as a mixture of uniform and point mass distributions. Here, we denote $\text{Unif}(a,b) := \mathbf{1}(x \geq a)\min\{\frac{x}{b-a}, 1\}$ as the CDF of the uniform distribution on (a,b) and $\delta_0 := \mathbf{1}(x \geq 0)$ the CDF of a point mass centered at zero.

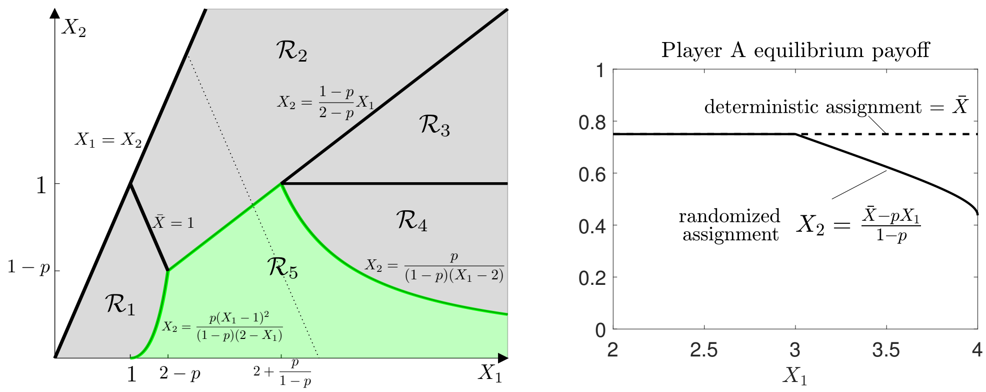

Figure 4.2: **Left:** Distinct parameter regions that encompass the entire class of Bayesian Lotto games (Definition 4.2.1), which are specified by the quadruple (X_1, X_2, p, X_B). Shown here is the space $X_1 \geq X_2$ for a fixed $p = 0.5$. We fix $X_B = 1$ here, though these regions are defined in general by the ratios X_1/X_B and X_2/X_B. The dashed black line segment indicates all instances of the game corresponding to randomized assignments of a fixed expected endowment $\bar{X} = pX_1 + (1-p)X_2 = 2$. **Right:** Player A's equilibrium payoff in the instances with fixed expected endowment $\bar{X} = 2$ (indicated by the dashed black line in the left panel). Here, the low endowment is parameterized by $X_2 = \frac{\bar{X}-pX_1}{1-p}$ with $X_1 \in [\bar{X}, \frac{\bar{X}}{p}]$. For instance, when $X_1 = 3.5$, the randomization (X_1, X_2, p) is given by $(3.5, 0.5, 0.5)$. Randomized assignments do not improve player A's payoff over deterministically assigning the expected endowment 2.

can take, $\bar{k} \in \{1, 2, 3\}$, correspond to transformed multipliers $\sigma = (\sigma_1, \sigma_2)$, with $\sigma_i := \frac{\lambda_i}{\lambda_B} > 0$ for $i \in \{1, 2\}$, lying in three disjoint regions of $\mathbb{R}^2_+$. These regions result directly from (4.40), and are given below.

$$\begin{aligned} \bar{k} = 1, \text{ if } \quad & p\sigma_1 \geq 1 \\ \bar{k} = 2, \text{ if } \quad & p\sigma_1 < 1 \text{ and } p\sigma_1 + (1-p)\sigma_2 \geq 1 \\ \bar{k} = 3, \text{ if } \quad & p\sigma_1 + (1-p)\sigma_2 < 1 \end{aligned} \tag{4.43}$$

Let us denote these regions as E_1, E_2, E_3, whose union is $\mathbb{R}^2_+$. The constructed strategies take three different forms described in Lemma 4.3.3, contingent on the value of $\bar{k}$. Thus, there are three cases the system of equations (4.39) can take. They are given below, where we seek to find multipliers $(\sigma_1, \sigma_2, \lambda_B)$ that satisfies one of the three cases for a given instance

$(X_1, X_2, p, X_B) \in \mathcal{G}$ – note that we can uniquely recover $(\lambda_1, \lambda_2, \lambda_B)$ from a tuple $(\sigma_1, \sigma_2, \lambda_B)$.

$$
\begin{array}{llll}
 & \underline{\text{Case 1}} & \underline{\text{Case 2}} & \underline{\text{Case 3}} \\
\text{(i)} & \dfrac{1}{2p\sigma_1^2} = \lambda_B X_1 & \dfrac{p}{2} + \dfrac{1-p\sigma_1}{\sigma_2} = \lambda_B X_1 & \dfrac{p}{2} + 1 - p = \lambda_B X_1 \\
\text{(ii)} & 0 = X_2 & \dfrac{(1-p\sigma_1)^2}{2(1-p)\sigma_2^2} = \lambda_B X_2 & \dfrac{1-p}{2} = \lambda_B X_2 \\
\text{(iii)} & p\sigma_1 X_1 = X_B & p\sigma_1 X_1 + (1-p)\sigma_2 X_2 = X_B & p\sigma_1 X_1 + (1-p)\sigma_2 X_2 = X_B \\
 & \text{such that} & \text{such that} & \text{such that} \\
\text{(iv)} & \sigma \in E_1 & \sigma \in E_2 & \sigma \in E_3 \\
\text{(v)} & \sigma_1 \leq \sigma_2 & \sigma_1 \leq \sigma_2 & \sigma_1 \leq \sigma_2
\end{array}
\qquad (\star)
$$

A solution $(\sigma_1, \sigma_2, \lambda_B)$ to $(\star)$ cannot satisfy two cases simultaneously, due to the E_i being disjoint. Observe that the individual battlefield values v_j (whose sum total is one) do not appear in these equations. In fact, one would arrive to the system $(\star)$ when considering Lotto games with any number $n \geq 1$ of battlefields whose total value is normalized to one – the individual battlefield values do not play a role in the analysis. For simplified exposition, we will henceforth consider players' strategies as allocations to a single battlefield of value one (F_A and F_B with no j dependence), noting this is mathematically equivalent to any arbitrary set of n battlefield valuations $\mathbf{v}$ that sum to one.

Also, note the multiplier σ_2 does not appear in the system of Case 1, but does appear in the condition (v). Here, σ_2 can be set to ∞ to satisfy (v), without affecting other variables. We detail the complete solutions to $(\star)$, and prove their associated strategies (from (4.42)) constitute Bayes-Nash equilibria in the result below.

Proposition 4.3.1. *The set of game instances in $\mathcal{G}$ for which a solution to $(\star)$ exists and their corresponding equilibrium strategies and payoffs are given as follows.*

Case 1: *Suppose $\frac{X_1}{X_B} \leq 1$ and $X_2 = 0$. The solution to $(\star)$ is given by $\lambda_1 = \frac{1}{2X_B}$, $\lambda_B = \frac{pX_1}{2X_B^2}$,* *and* $\lambda_2 \geq \frac{1}{2X_B}$. *The equilibrium strategies are*

$$F_A(t_1) = (1-\gamma_1)\delta_0 + \gamma_1 \, \mathit{Unif}(0, 2X_B), \quad F_A(t_2) = \delta_0, \quad F_B = \mathit{Unif}(0, 2X_B) \tag{4.44}$$

and the (ex-ante) equilibrium payoffs are

$$\pi_A = \frac{pX_1}{2X_B}, \qquad \pi_B = p(1 - \frac{X_1}{2X_B}) + (1-p) \tag{4.45}$$

Case 2: *Suppose* $\frac{X_2}{X_B} \leq H\left(\frac{X_1}{X_B}\right)$, *where* H *is defined in (4.15). The unique solution to* $(\star)$ *is* $\sigma_2 = (1 - \frac{X_B}{X_1})\sqrt{\frac{X_1/((1-p)X_2)}{p+(1-p)X_2/X_1}}$, $\sigma_1 = \frac{X_B-(1-p)\sigma_2 X_2}{pX_1}$, *and* $\lambda_B = \frac{\left(\sqrt{(1-p)X_2}+\sqrt{pX_1+(1-p)X_2}\right)^2}{2X_1^2}$. *The equilibrium strategies are*

$$\begin{aligned} F_A(t_1) &= \mathit{Unif}(L_2, L_2+L_1), \quad F_A(t_2) = \left(1 - \frac{1-p\sigma_1}{(1-p)\sigma_2}\right)\delta_0 + \frac{1-p\sigma_1}{(1-p)\sigma_2}\mathit{Unif}(0, L_2) \\ F_B &= (1-p\sigma_1)\,\mathit{Unif}(0, L_2) + p\sigma_1\,\mathit{Unif}(L_2, L_2+L_1) \end{aligned} \tag{4.46}$$

where $L_1 = \frac{p}{\lambda_B}$ *and* $L_2 = \frac{1-p\sigma_1}{\lambda_2}$. *The equilibrium payoffs are*

$$\pi_A = p(1-p\sigma_1)\left(1 - \frac{\sigma_1}{\sigma_2}\right) + \lambda_B X_B, \qquad \pi_B = \lambda_B X_B - \frac{1-p\sigma_1}{\sigma_2} + (1-p) \tag{4.47}$$

Case 3: *Suppose* $\frac{X_2}{X_1} = \frac{1-p}{2-p}$ *and* $2 - p < \frac{X_1}{X_B} < 2 + \frac{p}{1-p}$. *A solution to* $(\star)$ *is of the form* $\lambda_B = \frac{2-p}{2X_1}$, $\sigma_1 \in \left(\frac{\frac{X_B}{X_1}\left(2+\frac{p}{1-p}-\frac{X_1}{X_B}\right)}{p\left(1+\frac{p}{1-p}\right)}, \frac{\frac{X_B}{X_1}\left(2+\frac{p}{1-p}\right)}{p\left(2+\frac{p}{1-p}+\frac{1-p}{p}\right)}\right)$, *and* $\sigma_2 = \frac{X_B - p\sigma_1 X_1}{(1-p)X_2}$. *The equilibrium strategies are*

$$\begin{aligned} F_A(t_1) &= \mathit{Unif}(L_2, L_2+L_1), \quad F_A(t_2) = \mathit{Unif}(0, L_2) \\ F_B(x) &= (1 - p\sigma_1 + (1-p)\sigma_2)\delta_0 + (1-p)\sigma_2\,\mathit{Unif}(0, L_2) + p\sigma_1\,\mathit{Unif}(L_2, L_2+L_1) \end{aligned} \tag{4.48}$$

where $L_1 = \frac{p}{\lambda_B}$ *and* $L_2 = \frac{1-p}{\lambda_B} = 2X_2$. *The equilibrium payoffs are given by*

$$\pi_A = 1 - \lambda_B X_B, \qquad \pi_B = \lambda_B X_B. \tag{4.49}$$

Proof. We divide this proof into two parts. In the first part, we detail the steps used in each Case to calculate the algebraic solution to ($\star$) and the set of game instances for which it is valid. In the second part, we provide a proof that the corresponding strategies recovered from (4.42) do in fact constitute an equilibrium to the BL game. We will rely on shorthand notations $\gamma_i = X_i/X_B$ when convenient.

<u>Case 1</u>: The solution to ($\star$) can directly be found to be $\lambda_1 = \frac{1}{2X_B}$, $\lambda_B = \frac{pX_1}{2X_B^2}$, and any $\lambda_2 \geq \frac{1}{2X_B}$ (to satisfy (v)). Such a solution must also satisfy (iv), $p\sigma_1 = 1/X_1 \geq 1$. Combined with (ii), the set of valid game parameters is $\frac{X_1}{X_B} \leq 1$ and $X_2 = 0$: player A's budget in type 1 is smaller than player B's budget, and has a budget of zero in type 2. Since λ_2 does not appear in the algebraic equations of ($\star$) (only in the constraints), this is essentially unique. Plugging these values into (4.42), we obtain the resulting strategies.

<u>Case 2</u>: To solve for λ_B, we have $1 - p_1\sigma_1 = \sqrt{2(1-p)\lambda_2\sigma_2 X_2}$ from (ii). Substituting into (i), we obtain a quadratic equation in $\sqrt{\lambda_B} > 0$. Its (positive) solution yields the expression for λ_B.

Multiplying (ii) by $\frac{X_1}{X_2}$, the RHS of equations (i) and (ii) become equivalent. From (iv) of ($\star$), we use the substitution $1 - p\sigma_1 = 1 - \gamma_1^{-1} + (1-p)\sigma_2\frac{\gamma_2}{\gamma_1}$ to obtain $\sigma_2 = |1-\gamma_1^{-1}|\sqrt{\frac{\gamma_1/((1-p)\gamma_2)}{p+(1-p)\gamma_2/\gamma_1}}$. The condition (iv) requires $p\sigma_1 < 1$. Using the substitution $\sigma_1 = \frac{1-(1-p)\sigma_2\gamma_2}{p\gamma_1}$ from (iii), we deduce that $\gamma_1 > 1$:

$$\begin{aligned} p\sigma_1 &= \gamma_1^{-1}\left(1 - (1-p)\gamma_2|1-\gamma_1^{-1}|\sqrt{\frac{\gamma_1/((1-p)\gamma_2)}{p+(1-p)\gamma_2/\gamma_1}}\right) \\ &= \gamma_1^{-1} - |1-\gamma_1^{-1}|\sqrt{\frac{(1-p)\gamma_2/\gamma_1}{p+(1-p)\gamma_2/\gamma_1}} < 1 \\ &\Rightarrow 1-\gamma_1^{-1} > -|1-\gamma_1^{-1}| \Rightarrow \gamma_1 > 1 \end{aligned} \tag{4.50}$$

We can also deduce from (iii) and $X_1 \geq X_2$ that $\frac{X_2}{X_B} \leq 1$. The condition (iv) also requires $p\sigma_1 + (1-p)\sigma_2 \geq 1$. From this, we obtain

$$X_2 \leq \frac{1-p}{2-p}X_1. \tag{4.51}$$

Furthermore, the positivity of σ_2 is trivially satisfied. However, positivity of σ_1 requires that $1-(1-p)\sigma_2\gamma_2 > 0$. Plugging in the expression for σ_2, we obtain $\gamma_2(1-p)\left(2-\gamma_1\right) > -p$. Hence, the positivity constraint $\sigma_1 > 0$ is equivalent to

$$\gamma_1 \leq 2, \text{ or } \gamma_1 > 2 \text{ and } \gamma_2 < \frac{p}{(1-p)\left(\gamma_1-2\right)}. \tag{4.52}$$

Lastly, the constraint (v) requires $\sigma_1 \leq \sigma_2$. Plugging in the expression for σ_2, we deduce that $\gamma_2\left(1-(\gamma_1-1)^2\right) \leq (\gamma_1-1)^2\gamma_1\frac{p}{1-p}$. The term in parentheses on the LHS is positive when $\gamma_1 < 2$, and negative otherwise. Hence, we obtain

$$\gamma_2 \begin{cases} \leq \frac{\frac{p}{1-p}(\gamma_1-1)^2}{2-\gamma_1}, & \text{if } 1 < \gamma_1 \leq 2 \\ \geq 0, & \text{if } \gamma_1 > 2 \end{cases} \tag{4.53}$$

The intersection of conditions (4.53),(4.51), and (4.52) on the budget parameters X_1 and X_2, derived directly from (iv) and (v), yields $\gamma_2 \leq H\left(\gamma_1\right)$, where H was defined in (4.15). This establishes the set of games for which the system $(\star)$ has a solution in Case 2.

<u>Case 3</u>: We can directly obtain $\lambda_B = \frac{2-p}{2X_1}$. Note that $\lambda_B = \frac{1-p}{2X_2}$ as well, from which we obtain $X_2 = \frac{1-p}{2-p}X_1$. From (iii), we have $\sigma_2 = \frac{X_B-pX_1\gamma_1}{(1-p)X_2}$. Substituting this in the condition (iv), $p\sigma_1+(1-p)\sigma_2 < 1$, we obtain $\sigma_1 > \frac{\gamma_1^{-1}\left(2+\frac{p}{1-p}-\gamma_1\right)}{p\left(1+\frac{p}{1-p}\right)}$. Similarly, constraint (v), $\sigma_1 \leq \sigma_2$, yields $\sigma_1 \leq \frac{\gamma_1^{-1}\left(2+\frac{p}{1-p}\right)}{p\left(2+\frac{p}{1-p}+\frac{1-p}{p}\right)}$. A feasible σ_1 exists within these constraints if and only if $\gamma_1 > 1+\frac{1+\frac{1-p}{p}}{2+\frac{p}{1-p}+\frac{1-p}{p}} = 2-p$ (upper bound must be larger than lower bound), and $\gamma_1 < 2+\frac{p}{1-p}$ (lower bound must be positive). Subsequently, (4.42) recovers the strategies (4.48). The union of characterized parameter sets in all three cases constitutes the $\mathcal{R}_5$ region in Theorem 4.2.1.

<u>Part 2</u>: We prove the strategy profile (F_A, F_B) recovered from (4.42) is an equilibrium[9]. We can immediately deduce the strategies in Case 1 are equilibria to the BL game by observing

[9]This proof can be extended to scenarios where player A has an arbitrary number m endowment types. That is, if one can derive a solution to the system of $m+1$ equations (instead of just 3 in (4.39)), the profile (F_A, F_B) one constructs using Siegel's algorithm (in analogous manner to that used above) is an equilibrium to the BL game. Characterizing solutions to a system of $m+1$ equations, however, is a non-trivial extension that we leave for future study.

that player A has zero budget in type 2, and $(F_A(t_1), F_B)$ forms the unique equilibrium to the complete information General Lotto game [21] with a single battlefield of value p. We will focus on the strategies produced from Case 2, as the proof for Case 3 follows analogous arguments.

We first calculate the (ex-interim) payoffs from the strategies (4.46).

$$\begin{aligned}
U_A(F_A(t_1), F_B) &= \int_0^\infty F_B(x)\, dF_A(t_1) = \int_{L_2}^{L_2+L_1} [L_2\lambda_2 + \lambda_1(x - L_2)]\, \frac{\lambda_B}{p} dx \\
&= (1 - p\sigma_1)\left(1 - \frac{\sigma_1}{\sigma_2}\right) + \lambda_1 X_1 \\
U_A(F_A(t_2), F_B) &= \int_0^\infty F_B(x)\, dF_A(t_2) = \int_0^{L_2} \lambda_2 x \cdot \frac{\lambda_B}{1-p} dx = \lambda_2 X_2
\end{aligned} \tag{4.54}$$

The expected payoff (4.12) to player A is then $\pi_A = p(1 - p\sigma_1)\left(1 - \frac{\sigma_1}{\sigma_2}\right) + \lambda_B X_B$ (using (iii)). The payoff to B is $\pi_B = 1 - \pi_A$. We need to show F_A is a best-response to F_B, and vice versa.

For any $F_A'(t_1) \in \mathcal{L}(X_1)$, the payoff in type 1 is

$$\begin{aligned}
U_A(F_A'(t_1), F_B) &= \int_0^{L_2} \lambda_2 x\, dF_A'(t_1) + \int_{L_2}^{L_2+L_1} [L_2\lambda_2 + \lambda_1(x - L_2)]\, dF_A'(t_1) + \int_{L_2+L_1}^\infty dF_A'(t_1) \\
&= (\lambda_2 - \lambda_1)\left(\int_0^{L_2} x\, dF_A'(t_1) - L_2 \int_0^{L_2} dF_A'(t_1)\right) \\
&\quad + \lambda_1\left((L_2 + L_1)\int_{L_2}^{L_2+L_1} dF_A'(t_1) - \int_{L_2}^{L_2+L_1} x\, dF_A'(t_1)\right) \\
&\quad + \lambda_1 X_1 + L_2(\lambda_2 - \lambda_1) \\
&\leq \lambda_1 X_1 + L_2(\lambda_2 - \lambda_1) = \lambda_1 X_1 + (1 - \frac{\lambda_1}{\lambda_2})(1 - p\sigma_1).
\end{aligned} \tag{4.55}$$

In the second equality, we used the identities

$$\int_{L_2}^{L_2+L_1} x\, dF_A'(t_1) = X_1 - \int_0^{L_2} x\, dF_A'(t_1) - \int_{L_2+L_1}^\infty x\, dF_A'(t_1),$$

and

$$\int_{L_2}^{L_2+L_1} dF_A'(t_1) = 1 - \int_0^{L_2} dF_A'(t_1) - \int_{L_2+L_1}^\infty dF_A'(t_1).$$

The inequality follows from two applications of Markov's inequality:

$$\int_0^{L_2} x\, dF'_A(t_1) \leq L_2 \int_0^{L_2} dF'_A(t_1),$$

and

$$-\int_{L_2}^{L_2+L_1} x\, dF'_A(t_1) \leq -(L_2+L_1) \int_{L_2}^{L_2+L_1} dF'_A(t_1).$$

Hence, the payoff in type 1 is upper-bounded by $\lambda_1 X_1 + (1-\frac{\lambda_1}{\lambda_2})(1-p\sigma_1)$, which can be attained whenever $\text{supp}(F'_A(t_1)) \subseteq [L_2, L_2+L_1]$ (for which the Markov inequalities hold with equality). We provide analogous calculations for any $F'_A(t_2) \in \mathcal{L}(X_2)$:

$$\begin{aligned}
U_A(F'_A(t_2), F_B) &= \int_0^{L_2} \lambda_2 x\, dF'_A(t_2) + \int_{L_2}^{L_2+L_1} [L_2\lambda_2 + \lambda_1(x - L_2)]\, dF'_A(t_2) + \int_{L_2+L_1}^{\infty} dF'_A(t_2) \\
&= (\lambda_2 - \lambda_1)\left(L_2 \int_{L_2}^{L_2+L_1} dF'_A(t_2) - \int_{L_2}^{L_2+L_1} x\, dF'_A(t_2)\right) \\
&\quad + \left(\int_{L_2+L_1}^{\infty} dF'_A(t_2) - \lambda_2 \int_{L_2+L_1}^{\infty} dF'_A(t_2)\right) + \lambda_2 X_2 \\
&\leq -p(\sigma_2 - \sigma_1) \int_{L_2+L_1}^{\infty} dF'_A(t_2) + \lambda_2 X_2 \\
&\leq \lambda_2 X_2
\end{aligned} \tag{4.56}$$

The first inequality is similarly obtained from two applications of Markov's inequality. The second inequality follows from the condition (v), $\sigma_2 \geq \sigma_1$. Hence, the payoff in type 2 is upper-bounded by $\lambda_2 X_2$, which can be attained whenever $\text{supp}(F'_A(t_2)) \subseteq [0, L_2]$. The strategy F_A satisfies these properties, and hence is a best-response to F_B.

For any $F'_B \in \mathcal{L}(X_B)$, player B's expected payoff (4.12) is

$$\begin{aligned}
\Pi_B(F'_B, F_A) &= p\left[\int_{L_2}^{L_2+L_1} \frac{\lambda_B}{p}(x - L_2)\, dF'_B + \int_{L_2+L_1}^{\infty} dF'_B\right] \\
&\quad + (1-p)\left[\int_0^{L_2} \left(1 - \frac{\lambda_B L_2}{1-p} + \frac{\lambda_B}{1-p}x\right) dF'_B + \int_{L_2}^{\infty} dF'_B\right] \\
&= \lambda_B \left(\int_0^{L_2+L_1} x\, dF'_B - L_2 \int_0^{L_2+L_1} dF'_B\right) + p\int_{L_2+L_1}^{\infty} dF'_B + (1-p) \qquad (4.57) \\
&= \lambda_B X_B - \lambda_B L_2 + (1-p) + \lambda_B\left((L_1+L_2)\int_{L_2+L_1}^{\infty} dF'_B - \int_{L_2+L_1}^{\infty} x\, dF'_B\right) \\
&\leq \lambda_B X_B - \lambda_B L_2 + (1-p) = \pi_B
\end{aligned}$$

Player B's expected payoff is upper-bounded by π_B, which can be attained for any strategy with $\text{supp}(F'_B) \subseteq [0, L_2 + L_1]$. Because F_B is one such strategy, it is a best-response to F_A. $\square$

Equilibria in regions $\mathcal{R}_1$ - $\mathcal{R}_4$. Here, we provide all derivations of Bayes-Nash equilibria corresponding to the payoffs in regions $\mathcal{R}_i$, $i = 1, \ldots, 4$ (Theorem 4.2.1). As shown in Proposition 4.3.1, such equilibria cannot be generated through the standard method using Siegel's algorithm.

To first give some informal intuition, we provide some descriptions of the equilibria in these regions. The equilibrium strategies in $\mathcal{R}_1$ and $\mathcal{R}_2$ are convex combinations between an equilibrium strategy on the border[10] of $\mathcal{R}_5$ and equilibrium in its corresponding benchmark complete information game $\text{GL}(\bar{X}, X_B, \mathbf{v})$. As a result, the equilibrium payoff for any G in $\mathcal{R}_1$ or $\mathcal{R}_2$ coincides with the equilibrium payoff of its corresponding benchmark game. In regions $\mathcal{R}_3$ and $\mathcal{R}_4$, the "high" budget X_1 is disproportionately higher than the "low" budget X_2. We find an equilibrium strategy for the uninformed player is to not compete against the high budget at all, thus giving up a payoff p to the obfuscating player. In the forthcoming proofs, we make extensive use of Markov's inequality: $\int_a^b x\, dF \leq b\int_a^b dF$ for any distribution F.

Region $\mathcal{R}_3$: Suppose $G \in \mathcal{R}_3 = \left\{(\gamma_1, \gamma_2) \in \mathcal{R} : \gamma_1 \geq 2 + \frac{p}{1-p} \text{ and } 1 \leq \gamma_2 \leq \frac{1-p}{2-p}\gamma_1\right\}$. The fol-

[10] Since equilibria on the border are not necessarily unique, i.e. Case 3 parameters of Proposition 4.3.1, the equilibria in the regions $\mathcal{R}_1$ and $\mathcal{R}_2$ are not unique. However, all equilibria in one game instance yield identical payoffs, since it is a constant sum game (in ex-ante payoffs).

lowing is an equilibrium.

$$\begin{aligned}
&F_A(t_1) = \text{Unif}\left(2X_2, 2(X_1 - X_2)\right), \quad F_A(t_2) = \text{Unif}\left(0, 2X_2\right), \\
&F_B = (1-\gamma_2^{-1})\delta_0 + \gamma_2^{-1}\text{Unif}\left(0, 2X_2\right)
\end{aligned} \tag{4.58}$$

The equilibrium payoff is given by $\pi_A(G) = p + (1-p)\left(1 - \frac{1}{2\gamma_2}\right)$.

Proof. First, we show F_A is a best-response to F_B. For any $\{F_A'(t_i) \in \mathcal{L}(X_i)\}_{i=1,2}$, player A's expected payoff is

$$\begin{aligned}
&p\left[\int_0^{2X_2}\left(1 - \gamma_2^{-1} + \frac{\gamma_2^{-1}}{2X_2}x\right)\,dF_A'(t_1) + \int_{2X_2}^{\infty} dF_A'(t_1)\right] \\
&\quad + (1-p)\left[\int_0^{2X_2}\left(1 - \gamma_2^{-1} + \frac{\gamma_2^{-1}}{2X_2}x\right)\,dF_A'(t_2) + \int_{2X_2}^{\infty} dF_A'(t_2)\right] \\
&\leq p + (1-p)\left[\int_0^{2X_2}\left(1 - \gamma_2^{-1} + \frac{\gamma_2^{-1}}{2X_2}x\right)\,dF_A'(t_2) + \int_{2X_2}^{\infty} dF_A'(t_2)\right]
\end{aligned} \tag{4.59}$$

The inequality follows by selecting any $F_A'(t_1)$ such that $\text{supp}(F_A'(t_1)) \subset [2X_2, \infty)$, which awards player A the payoff p from state 1 outright. This is possible because $\gamma_1 \geq 2 + \frac{p}{1-p} > 2$, from the assumption. The expression above can be re-written and upper-bounded as follows:

$$\begin{aligned}
&p + (1-p)\left[\left(1 - \gamma_2^{-1}\right)\int_0^{2X_2} dF_A'(t_2) + \frac{\gamma_2^{-1}}{2} + \int_{2X_2}^{\infty} dF_A'(t_2) - \frac{\gamma_2^{-1}}{2X_2}\int_{2X_2}^{\infty} x\,dF_A'(t_2)\right] \\
&\quad \leq p + (1-p)\left(1 - \frac{\gamma_2^{-1}}{2}\right)
\end{aligned} \tag{4.60}$$

The inequality holds with equality if and only if $\text{supp}(F_A'(t_2)) \subseteq [0, 2X_2]$. We have thus established an upper bound on A's payoff to F_B that is achieved by F_A.

Now we show F_B is a best-response to F_A. Let $K := (1-p) - \frac{p\gamma_2}{\gamma_1 - 2\gamma_2} \geq 0$, which is non-negative due to the assumption $\gamma_2 \leq \frac{1-p}{2-p}\gamma_1$. For any $F'_B \in \mathcal{L}(X_B)$, player B's payoff is

$$\begin{aligned}
&p\left[\int_{2X_2}^{2(X_1-X_2)} \frac{x-2X_2}{2(X_1-2X_2)}\,dF'_B + \int_{2(X_1-X_2)}^{\infty} dF'_B\right] + (1-p)\left[\int_0^{2X_2} \frac{x}{2X_2}\,dF'_B + \int_{2X_2}^{\infty} dF'_B\right] \\
&= \frac{1-p}{2X_2}\int_0^{2X_2} x\,dF'_B + \frac{p}{2(X_1-2X_2)}\int_{2X_2}^{2(X_1-X_2)} x\,dF'_B \\
&\quad - \left(\frac{pX_2}{X_1-2X_2} - (1-p)\right)\int_{2X_2}^{2(X_1-X_2)} dF'_B + \int_{2(X_1-X_2)}^{\infty} dF'_B
\end{aligned} \tag{4.61}$$

Applying the identity $\int_0^{2X_2} x\,dF'_B = X_B - \int_{2X_2}^{2(X_1-X_2)} x\,dF'_B - \int_{2(X_1-X_2)}^{\infty} x\,dF'_B$, we then obtain

$$\begin{aligned}
&= K\left(-\frac{1}{2X_2}\int_{2X_2}^{2(X_1-X_2)} x\,dF'_B + \int_{2X_2}^{2(X_1-X_2)} dF'_B\right) + (1-p)\frac{\gamma_2^{-1}}{2} \\
&\quad + \int_{2(X_1-X_2)}^{\infty} dF'_B - \frac{1-p}{2X_2}\int_{2(X_1-X_2)}^{\infty} x\,dF'_B \\
&\leq (1-p)\frac{\gamma_2^{-1}}{2} + \left(p + (1-p)\left(2 - \frac{\gamma_1}{\gamma_2}\right)\right)\int_{2(X_1-X_2)}^{\infty} dF'_B \\
&\leq (1-p)\frac{\gamma_2^{-1}}{2}
\end{aligned} \tag{4.62}$$

The first inequality results from applying Markov's inequality to the expressions $\int_{2X_2}^{2(X_1-X_2)} x\,dF'_B$ and $\int_{2(X_1-X_2)}^{\infty} x\,dF'_B$. The second inequality follows from non-positivity of term in parentheses (from assumption of the Lemma). This inequality holds with equality if and only if $\text{supp}(F'_B) \subseteq [0, 2X_2]$. We have thus established an upper bound on B's payoff to F_A that is achieved by F_B. $\square$

<u>**Region $\mathcal{R}_4$**</u>: Suppose $G \in \mathcal{R}_4 := \left\{(\gamma_1,\gamma_2) \in \mathcal{R} : \gamma_1 \geq 2 + \frac{p}{1-p} \text{ and } \frac{p}{(1-p)(\gamma_1-2)} \leq \gamma_2 \leq 1\right\}$. The following is an equilibrium: $F_A(t_1) = \text{Unif}\,(2X_B, 2(X_1 - X_B))$, $F_A(t_2) = (1-\gamma_2)\delta_0 + \gamma_2\text{Unif}\,(0, 2X_B)$, $F_B = \text{Unif}\,(0, 2X_B)$, and the equilibrium payoff is given by $\pi_A(G) = p + (1-p)\frac{\gamma_2}{2}$.

Proof. First, show F_A is a best-response to F_B. Player A's payoff for any $\{F'_A(t_i) \in \mathcal{L}(X_i)\}_{i=1,2}$

is

$$
\begin{aligned}
&p\left[\int_0^{2X_B} \frac{x}{2X_B}\, dF'_A(t_1) + \int_{2X_B}^{\infty} dF'_A(t_1)\right] + (1-p)\left[\int_0^{2X_2} \frac{x}{2X_B}\, dF'_A(t_2) + \int_{2X_B}^{\infty} dF'_A(t_2)\right] \\
&\le p + (1-p)\left[\int_0^{2X_2} \frac{x}{2X_B}\, dF'_A(t_2) + \int_{2X_B}^{\infty} dF'_A(t_2)\right] \\
&= p + (1-p)\left[\frac{\gamma_2}{2} + \int_{2X_B}^{\infty} dF'_A(t_2) - \frac{1}{2X_B}\int_{2X_B}^{\infty} x\, dF'_A(t_2)\right] \\
&\le p + (1-p)\frac{\gamma_2}{2}.
\end{aligned}
\tag{4.63}
$$

The first inequality follows by selecting any $F'_A(t_1)$ such that $\text{supp}(F'_A(t_1)) \subset [2X_B, \infty)$, which awards player A p outright. This is possible because $\gamma_1 \ge 2$, from the assumption. The second inequality follows by applying Markov's inequality to $\int_{2X_B}^{\infty} x\, dF'_A(t_2)$. This inequality holds if and only if $\text{supp}(F'_A(t_2)) \subseteq [0, 2X_B]$.

Now we show F_B is a best-response to F_A. For any $F_B \in \mathcal{L}(X_B)$, player B's payoff is

$$
\begin{aligned}
&\frac{(1-p)\gamma_2}{2X_B}\int_0^{2X_B} x\, dF'_B + (1-p)\,(1-\gamma_2)\int_0^{2X_B} dF'_B \\
&\quad + \frac{p}{2(X_1 - 2X_B)}\int_{2X_B}^{2(X_1-X_B)} x\, dF'_B + \left((1-p)\gamma_2 - \frac{p}{\gamma_1 - 2}\right)\int_{2X_B}^{2(X_1-X_B)} dF'_B + \int_{2(X_1-X_B)}^{\infty} dF'_B \\
&= K\left(\int_{2X_B}^{2(X_1-X_B)} dF'_B - \frac{1}{2X_B}\int_{2X_B}^{2(X_1-X_B)} x\, dF'_B\right) \\
&\quad + (p + (1-p)\gamma_2)\int_{2(X_1-X_B)}^{\infty} dF'_B - \frac{(1-p)\gamma_2}{2X_B}\int_{2(X_1-X_B)}^{\infty} x\, dF'_B + (1-p)\left(1 - \frac{\gamma_2}{2}\right) \\
&\le (1-p)\left(1 - \frac{\gamma_2}{2}\right)
\end{aligned}
\tag{4.64}
$$

where $K := (1-p)\gamma_2 - \frac{p}{\gamma_1 - 2} > 0$, from the assumption. In the equality, we substituted $\int_0^{2X_B} x\, dF'_B = X_B - \int_{2X_B}^{2(X_1-X_B)} x\, dF'_B - \int_{2(X_1-X_B)}^{\infty} x\, dF'_B$. The inequality is due to applying Markov's inequality and noting the expression $p + (1-p)\gamma_2 - (1-p)\gamma_2\,(\gamma_1 - 1) = p + (1-p)\gamma_2\,(2-\gamma_1) \le 0$ holds from the assumptions. This holds with equality if and only if $\text{supp}(F'_B) \subseteq [0, 2X_B]$. □

Regions $\mathcal{R}_1$ and $\mathcal{R}_2$:

Consider the set of $(\gamma_1, \gamma_2) \in \mathcal{R}_1 \cup \mathcal{R}_2$ that have a fixed average budget $\bar{\gamma}$. Define

$$\gamma^{\mathrm{bd}} := \begin{cases} (\bar{\gamma}/p, 0) \in \mathcal{R}, & \text{if } \bar{\gamma} \le p \\ (2 - p/\bar{\gamma}, H(2 - p/\bar{\gamma})) \in \mathcal{R}, & \text{if } p < \bar{\gamma} \le 1 \\ ((2-p)\bar{\gamma}, (1-p)\bar{\gamma}) \in \mathcal{R}, & \text{if } 1 < \bar{\gamma} \end{cases} \tag{4.65}$$

where H is defined in (4.15). The points γ^{bd} specified above for $\bar{\gamma} \le 1$ are on the border of $\mathcal{R}_5$, whose equilibria are given in Proposition 4.3.1. The points for $1 < \bar{\gamma}$ are on the upper border of $\mathcal{R}_3$, where a equilibrium is given in (4.58). Define

$$F_A^{\mathrm{bd}} := \begin{cases} \text{given by (4.44) at } \gamma^{\mathrm{bd}}, & \text{if } \bar{\gamma} \le p \\ \text{given by (4.46) at } \gamma^{\mathrm{bd}}, & \text{if } p < \bar{\gamma} \le 1 \\ \text{given by (4.48) at } \gamma^{\mathrm{bd}}, & \text{if } 1 < \bar{\gamma} \end{cases} \tag{4.66}$$

Here, F_A^{bd} is an equilibrium strategy for player A at the boundary point γ^{bd}. Let us also define $(\bar{F}_A, \bar{F}_B)$ as the Nash equilibrium at $(\bar{\gamma}, \bar{\gamma}) \in \mathcal{R}$, which is simply the equilibrium in the corresponding complete information game. That is,

$$\begin{aligned} \bar{F}_A &= \begin{cases} (1-\bar{\gamma})\delta_0 + \bar{\gamma}\mathrm{Unif}([0, 2X_B]), & \text{if } \bar{\gamma} \le 1 \\ \mathrm{Unif}([0, 2\bar{X}), & \text{if } \bar{\gamma} > 1 \end{cases}, \\ \bar{F}_B &= \begin{cases} \mathrm{Unif}([0, 2X_B]), & \text{if } \bar{\gamma} \le 1 \\ (1-\bar{\gamma}^{-1})\delta_0 + \bar{\gamma}^{-1}\mathrm{Unif}([0, 2\bar{X}), & \text{if } \bar{\gamma} > 1 \end{cases} \end{aligned} \tag{4.67}$$

Let $\alpha \in [0,1]$ be the appropriate scaling that gives $\alpha\gamma^{\mathrm{bd}} + (1-\alpha)\cdot(\bar{\gamma}, \bar{\gamma}) = (\gamma_1, \gamma_2)$. We claim the strategy profile $(\alpha F_A^{\mathrm{bd}} + (1-\alpha)\bar{F}_A, \bar{F}_B)$ is an equilibrium for $(\gamma_1, \gamma_2) \in \mathcal{R}_1 \cup \mathcal{R}_2$.

Proof. Since we know that $(\bar{F}_A, \bar{F}_B)$ is an equilibrium, it will suffice to show that $(F_A^{\mathrm{bd}}, \bar{F}_B)$ is also an equilibrium for all $\bar{\gamma}$. After lengthy calculations, we see that the payoffs from $(F_A^{\mathrm{bd}}, \bar{F}_B)$

coincide with the payoffs from $(\bar{F}_A, \bar{F}_B)$.

For $\bar{\gamma} \leq p$, the equilibrium at γ^{bd} is given by Case 1 (4.44), where player B's strategy is precisely $\bar{F}_B$. For $p < \bar{\gamma} \leq 1$, the equilibrium at γ^{bd} is given by Case 2 (4.46), where it is also true that B's strategy is $\bar{F}_B$ (on these border points, $\sigma_1 = \sigma_2$). For $1 < \bar{\gamma} \leq \frac{1}{1-p}$, the equilibrium at γ^{bd} is given by Case 3 (4.48). We note that although player B's equilibrium strategy here is not unique, $\bar{F}_B$ is one such strategy.

Lastly, for $\frac{1}{1-p} < \bar{\gamma}$, $(F_A^{\text{bd}}, \bar{F}_B)$ is an equilibrium at γ^{bd}, where F_A^{bd} is player A's equilibrium strategy at the border of $\mathcal{R}_3$ (4.58). This strategy is also identical to the monotonic equilibrium strategy from Case 3 (4.48). Hence, the proof that $(F_A^{\text{bd}}, \bar{F}_B)$ is an equilibrium follows from the analysis in Proposition 4.3.1. Note that player B's $\mathcal{R}_3$ equilibrium strategy written in (4.58) is not $\bar{F}_B$. Indeed, $\bar{F}_B$ in general is not an equilibrium strategy in the interior of $\mathcal{R}_3$. At the border however, we know of at least two equilibria (giving the same payoffs), one of them being $(F_A^{\text{bd}}, \bar{F}_B)$. □

Chapter 5

Pre-emptive mechanisms: The value of revealing information

As we explored in the previous chapter, the conventional wisdom suggests that an agent improves her competitive position by increasing her resource budget and investing in more accurate information in adversarial resource allocation scenarios. By extension, an agent would prefer to face an adversary with fewer resources and less information, as this should sway the strategic outcomes in her favour. Note that these insights have been applied across a spectrum of problem settings in cyber-physical systems security [63, 64, 65] and beyond in, e.g., airport security [66, 67], wildlife protection [68], market economics [69], and political campaigning [70].

In this chapter, we challenge the conventional wisdom by explicitly considering *pre-emption* as a viable, alternative component in an agent's strategic decision making: In Section 5.1, we study *General Lotto games with concessions.* Here, we consider concessions under two formats: *budget concessions*, in which the agent willingly reduces her own budget; and, *value concessions*, which involve voluntary non-participation on a specified subset of the battlefields. As such, concessions contradict the conventional wisdom, as the conceding agent weakens herself while providing the adversary with more information. Our goal in this work is to understand whether these concession formats represent valid strategic options in adversarial resource allocation. Then, in Section 5.2, we study *General Lotto games with pre-allocations.* Under this model, one of the players can strategically decide how to deploy resources before the actual engagement takes place. The placement of the pre-allocated resources thus has an effect on how the allocation decisions are made at the time of competition.

Our goal is to understand the effect of including such pre-emptive mechanisms into an agent's decision space on the emergent strategic outcomes of the game. In other words, we wish to understand if and when a strategic agent will leverage these additional opportunities, or whether these never offer strategic advantages as the intuition suggests.

5.1 General Lotto games with concessions

5.1.1 Model

In the Colonel Blotto literature, researchers have proposed generalized game models involving more than two players [50, 71, 48, 49]. One such class of models, termed *Coalitional Lotto games* [50], examines a scenario where two self-interested entities compete against a common opponent. More formally, consider a three-player, constant-sum Stackleberg game where players 0, A and B, have respective budgets $X_0 > 0$ and $X_A, X_B \geq 0$. Players A and B are simultaneously engaged in independent General Lotto games against player 0 over the disjoint sets of battlefields $\mathcal{B}_A$ and $\mathcal{B}_B$, respectively. For each $i \in \{A, B\}$, let $\Phi_i \geq 0$ denote the cumulative value of battlefields in the set $\mathcal{B}_i$. The game is played over two stages, defined as follows:

- *Stage 1*: Player 0 allocates its budget as $X_{0,A}, X_{0,B} \geq 0$ such that $X_{0,A} + X_{0,B} \leq X_0$. Then, player 0's selection becomes common knowledge and binding;
- *Stage 2*: The games $\text{GL}(X_{0,i}, X_i, \Phi_i)$, $i = A, B$, are played.

Players' final payoffs are their equilibrium payoffs from the General Lotto games in Stage 2. We assume that player 0's Stage 1 allocation is made to maximize her final payoff, i.e., given a

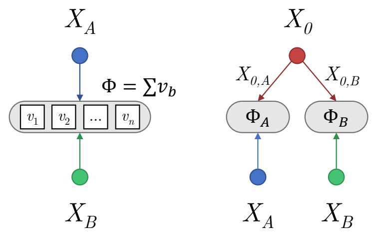

(a) Two-player. (b) Three-player.

Figure 5.1: *Illustration of the two- and three-player General Lotto game models. (a) In the two-player General Lotto game, players A and B with respective budgets $X_A, X_B \geq 0$ simultaneously compete over n battlefields with cumulative value $\Phi \geq 0$. (b) In Stage 1 of the three-player General Lotto game, player 0 allocates her budget X_0 between the two (two-player) General Lotto games against players A and B. Subsequently, in Stage 2, the two General Lotto games are played simultaneously.*

three-player General Lotto game, player 0's Stage 1 allocation $(X^*_{0,A}, X^*_{0,B}) \in \mathbb{R}^2_{\geq 0}$ satisfies

$$(X^*_{0,A}, X^*_{0,B}) \in \underset{\substack{X_{0,A}, X_{0,B} \geq 0, \\ X_{0,A}+X_{0,B} \leq X_0}}{\arg\max} \sum_{i \in \{A,B\}} u^*_0(X_{0,i}, X_i, \Phi_i).$$

Furthermore, for each player $i \in \{A, B\}$, we let

$$\Pi_i(X_0, X_A, X_B, \Phi_A, \Phi_B) = u^*_i(X_i, X^*_{0,i}, \Phi_i)$$

denote her final payoff under player 0's optimal Stage 1 allocation. For ease of notation, we use the shorthand $\Pi_i = \Pi_i(X_0, X_A, X_B, \Phi_A, \Phi_B)$, $i \in \{A, B\}$, when the dependence on the game parameters is clear. Figure 5.1b depicts the three-player General Lotto game.[1]

The authors of [50] investigate when players A and B can cooperate to change the outcome of the three-player General Lotto game in their favour. Specifically, they focus on unilateral budget transfers, where the following preliminary stage is played prior to Stages 1 and 2 of the nominal game:[2]

[1]One can easily verify that any two-player General Lotto game can be recast as a three-player General Lotto game.

[2]Though we model player 0's selection of her budget allocation, and of players A and B's alliances or concessions (defined in the forthcoming section) as sequential under our proposed framework, we note that the players' strategic interactions still occur simultaneously in Stages 0 and 2.

- *Stage 0:* Players A and B select a unilateral budget transfer $\delta \in [-X_B, X_A]$, to be made from A to B. Then, the transfer δ becomes common knowledge and binding, and the players' resulting budgets are $X'_A = X_A - \delta$ and $X'_B = X_B + \delta$, respectively.

Accordingly, under a unilateral budget transfer $\delta \in [-X_B, X_A]$ selected in Stage 0, each player $i \in \{A, B\}$ receives as her final payoff:

$$\Pi_i^{\mathrm{al}}(\delta) := \Pi_i(X_0, X_A - \delta, X_B + \delta, \Phi_A, \Phi_B).$$

We will refer to these unilateral budget transfers between players A and B as *alliances*. Observe that players' payoffs under no alliance (i.e., $\delta = 0$) are the same as their payoffs in the nominal game.

Interestingly, [50] shows that there exist scenarios under which alliances are mutually beneficial to A and B, i.e., there exists $\delta \in [-X_B, X_A]$ such that

$$\Pi_i^{\mathrm{al}}(\delta) > \Pi_i^{\mathrm{al}}(\delta), \text{ for each } i \in \{A, B\}.$$

This implies that a player $i \in \{A, B\}$ can improve her payoff by making herself weaker, and making her competitor, player $-i$, stronger. Figure 5.2a depicts the three-player General Lotto game with alliances.

5.1.2 Budget concessions

Recall that an alliance (as defined in [50]) consists of a unilateral budget transfer between players A and B, meaning one of the players $i \in \{A, B\}$ becomes weaker, while player $-i$ becomes stronger. For an alliance to be mutually beneficial, note that even the weaker player's payoff must improve. The existence of mutually-beneficial alliances [50] suggests that player i may improve her payoff simply by weakening herself without making player $-i$ stronger, i.e., by making a budget concession.

To study the strategic role that budget concessions play in three-player General Lotto games, we introduce a modified three-player General Lotto game model, where we add the following preliminary stage that is played prior to Stages 1 and 2 of the nominal three-player General Lotto game:

- *Stage 0:* Each player $i \in \{A, B\}$ selects a portion of her budget $x_i \in [0, X_i]$ to concede. Then, the concession profile $(x_A, x_B) \in [0, X_A] \times [0, X_B]$ becomes common knowledge and binding, and the players' resulting budgets are $X'_A = X_A - x_A$ and $X'_B = X_B - x_B$, respectively.

For ease of exposition, we define the *three-player General Lotto game with budget concessions* as the one-shot game played between players A and B in Stage 0, where each player $i \in \{A, B\}$ selects an action $x_i \in [0, X_i]$ and the players' payoffs coincide with their Stage 2 payoffs assuming that all players 0, A and B play according to a subgame perfect equilibrium (SPE) in Stages 1 and 2,[3] i.e., under a concession profile $(x_A, x_B) \in [0, X_A] \times [0, X_B]$ selected in Stage 0, each player $i \in \{A, B\}$ receives a final payoff

$$\Pi_i^{\text{bc}}(x_i, x_{-i}) := \Pi_i(X_0, X_A - x_A, X_B - x_B, \Phi_A, \Phi_B).$$

In the event that there is more than one SPE in Stages 1 and 2, we may choose an SPE selection rule to ensure the functions Π_i^{bc}, $i \in \{A, B\}$, are well-defined, e.g., the SPE with lowest resulting payoff to player A. Nevertheless, our results for the three-player General Lotto game with budget concessions in the forthcoming Theorem 5.1.1 (i.e., the weak dominance of the action $x_i = 0$ for both players $i \in \{A, B\}$, and the payoff equivalence of all pure strategy Nash equilibria) hold *regardless of the chosen SPE selection rule.*

Note that the players' payoffs under no budget concession (i.e., $(x_A, x_B) = (0, 0)$) are equivalent to their payoffs in the nominal three-player General Lotto game. For a given nominal three-player General Lotto game, we can represent the corresponding three-player General Lotto game with budget concessions as the tuple $G = (\{X_i, \Pi_i^{\text{bc}}\})_{i \in \{A,B\}}$. Figure 5.2b depicts the

[3] The notion of subgame perfect equilibrium extends the notion of Nash equilibrium to multi-stage games, and represents a sequence of player strategies that satisfy the Nash equilibrium condition in within each stage.

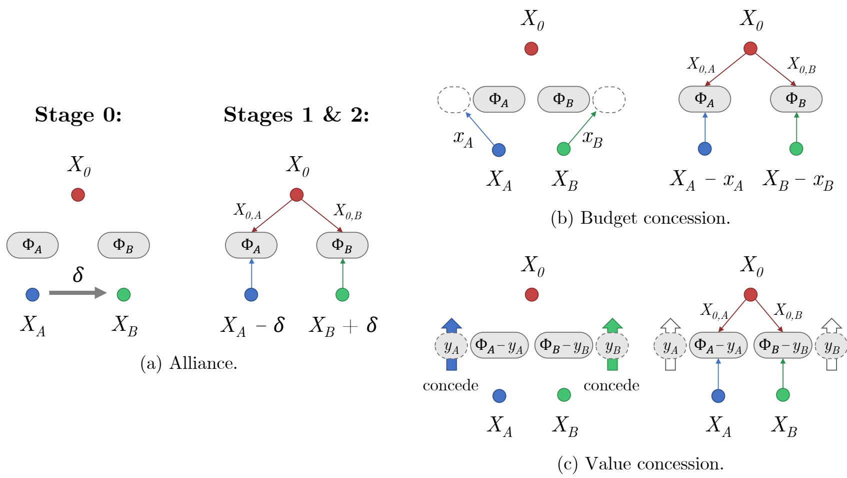

Figure 5.2: *Illustration of the three-player General Lotto game model with alliances proposed in [50], and our proposed variants for studying budget and value concessions.* (a) In the alliances model, player A makes a unilateral budget transfer of $\delta \in [-X_B, X_A]$ forces to player B. (b) In our budget concession model, each player $i \in \{A, B\}$ concedes some budget $x_i \in [0, X_i]$. (c) In our value concession model, each player $i \in \{A, B\}$ concedes some value $y_i \in [0, \Phi_i]$. Any alliances and concessions are selected, and become binding and common knowledge before the subsequent three-player General Lotto game is played.

three-player General Lotto game with budget concessions.

Our first result demonstrates that budget concessions are never advantageous under our three-player General Lotto games with budget concessions game model.

Theorem 5.1.1. *Given any three-player General Lotto game with budget concessions and any SPE selection rule, the choice $x_i = 0$ is a weakly-dominant strategy for each player $i \in \{A, B\}$. That is, for any $(x_A, x_B) \in [0, X_A] \times [0, X_B]$, it holds that*

$$\Pi_i^{\mathrm{bc}}(0, x_{-i}) \geq \Pi_i^{\mathrm{bc}}(x_i, x_{-i}), \quad i = A, B.$$

Futhermore, all pure strategy Nash equilibria are payoff equivalent, i.e., each player $i \in \{A, B\}$ receives payoff equal to $\Pi_i^{\mathrm{bc}}(0, 0)$.

Observe that the above result establishes that each player $i \in \{A, B\}$ prefers not to concede any budget, regardless of player $-i$'s selection in Stage 0. This implies that the payoff improvement in mutually-beneficial alliances stem from the *transfer* of budget between players A and B, and that one of the players $i \in \{A, B\}$ cannot simply weaken herself to achieve similar benefits. In other words – and perhaps surprisingly – the benefits that the weakened player reaps from an alliance rely on making the other player stronger.

5.1.3 Value concessions

As with budget concessions, we propose a variation on the three-player General Lotto game model to study the strategic role of value concessions, which we term *three-player General Lotto game with value concessions*. This variant has the following preliminary stage:

- *Stage 0:* Each player $i \in \{A, B\}$ selects a portion of the value $y_i \in [0, \Phi_i]$ to concede. Then, the values y_A and y_B are awarded to player 0, and the concession profile $(y_A, y_B) \in [0, \Phi_A] \times [0, \Phi_B]$ becomes common knowledge and binding.

Once again, we define the three-player General Lotto game with value concessions as the one-shot game played between A and B in Stage 0, where each player $i \in \{A, B\}$ selects an action $y_i \in [0, \Phi_i]$ and the players' payoffs coincide with their Stage 2 payoffs assuming that all players 0, A and B play according to a subgame perfect equilibrium (SPE) in Stages 1 and 2, i.e., under a concession profile $(y_A, y_B) \in [0, \Phi_A] \times [0, \Phi_B]$ selected in Stage 0, each player $i \in \{A, B\}$ receives a final payoff

$$\Pi_i^{\text{vc}}(y_i, y_{-i}) := \Pi_i(X_0, X_A, X_B, \Phi_A - y_A, \Phi_B - y_B).$$

As with budget concessions, in the event that the players' SPE payoffs in Stages 1 and 2 are non-unique, we can choose one of many SPE selection rules to ensure the functions Π_i^{vc}, $i \in \{A, B\}$, are well-defined. Nonetheless, we show that our main results for the three-player General Lotto game with value concessions in the forthcoming Theorem 5.1.2 (i.e., the existence of beneficial value concessions, the uniqueness (or non-existence) and special structure of the pure strategy Nash equilibrium) hold *regardless of the chosen SPE selection rule.*

Observe that the players' payoffs under no value concessions (i.e., $(y_A, y_B) = (0, 0)$) are equivalent to their payoffs in the nominal three-player General Lotto game. For a given nominal three-player General Lotto game, we represent the corresponding three-player General Lotto game with value concessions as the tuple $G = (\{\Phi_i, \Pi_i^{\text{vc}}\})_{i\in\{A,B\}}$. Figure 5.2c depicts the three-player General Lotto game with value concessions.

In the previous section, we showed that a player cannot improve her payoff by making a budget concession in three-player General Lotto games as the choice $x_i = 0$ is a weakly-dominant strategy for each player $i \in \{A, B\}$. This supports the conventional intuition that concessions cannot represent a valid strategic option in adversarial interactions. In contrast, our next result shows that this intuition is false, as in certain instances there do exist concession profiles of the form $(y_A, y_B) \neq (0, 0)$ that correspond with pure strategy Nash equilibria of the game. We restrict our attention to pure strategy value concessions due to the information structure of the interaction, i.e., the players' Stage 0 choices become common knowledge and binding in the subsequent stages.

Theorem 5.1.2. *Consider the family of all three-player General Lotto games with value concessions $\mathcal{G}$ under any SPE selection rule. The following statements hold:*

1. *For either player $i \in \{A, B\}$, there exist games $G \in \mathcal{G}$ in which $\Pi_i^{\text{vc}}(y_i, 0) > \Pi_i^{\text{vc}}(0, 0)$ for $y_i \in (0, \Phi_i]$, i.e., a non-zero value concession* strictly *improves player i's payoff; and,*

2. *In every game $G \in \mathcal{G}$, either there is a unique pure strategy Nash equilibrium either of the form $(0, 0)$, $(y_A, 0)$ or $(0, y_B)$ with $y_i \in (0, \Phi_i]$, $i \in \{A, B\}$, or there are no pure strategy Nash equilibria.*

Given any two-player General Lotto game with value concessions, the choice $y_i = 0$ is a weakly-dominant strategy for each player $i \in \{A, B\}$.

Observe that the SPE payoffs of players A, B and 0 always sum to $\Phi_A + \Phi_B$ in the three-player General Lotto game with value concessions, and the SPE payoffs of players 0 and i both increase after i makes a beneficial value concession from the nominal game (i.e., in Stage 0, $i \in \{A, B\}$ selects $y_i \in [0, \Phi_i]$ such that $\Pi_i^{\text{vc}}(y_i, 0) > \Pi_i^{\text{vc}}(0, 0)$).[4] It follows that if a player $i \in \{A, B\}$

[4]For any value concession $(y_A, y_B) \neq (0, 0)$, it is straightforward to verify that player 0 can secure strictly

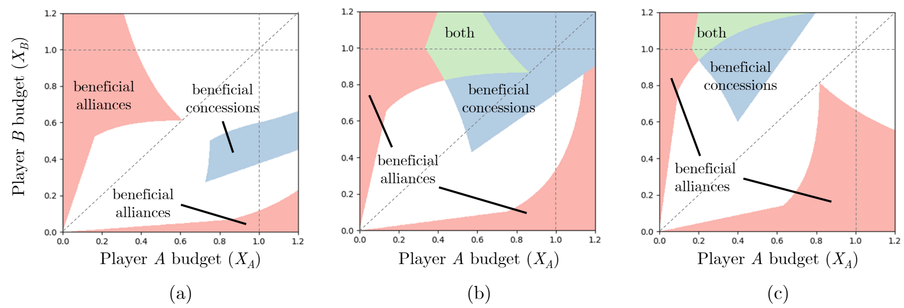

Figure 5.3: *Comparing the existence of mutually-beneficial alliances and beneficial concessions for player A.* We identify the parameters for which player B has a mutually-beneficial alliance (in red), a beneficial value concession (in blue) or both (in green) under normalized player budgets (i.e., $X_0 = 1$), for $X_A, X_B \in [0, 1.2]$, $\Phi_A = 1$ and (a) $\Phi_B = 3/8$, (b) $\Phi_B = 3/4$, and (c) $\Phi_B = 3/2$. The white area indicates where player A has no mutually-beneficial alliance or beneficial value concession. Observe that mutually-beneficial alliances exist in regimes where the ratio Φ_A/X_A is much greater than Φ_B/X_B, whereas beneficial concessions occur in regimes where the two ratios are sufficiently close. Furthermore, our plots suggest that beneficial alliances and concessions co-exist when Φ_A is less than, equal to or only slightly greater than Φ_B, but cease to co-exist when Φ_A is much greater than Φ_B.

makes a beneficial value concession from the nominal game, then player $-i$'s payoff suffers a loss (i.e., $\Pi^{\text{vc}}_{-i}(y_i, 0) < \Pi^{\text{vc}}_{-i}(0, 0)$). Thus, the existence of pure strategy Nash equilibria of the form $(y_A, 0)$ and $(0, y_B)$ with $y_A, y_B > 0$ is interesting, since the non-conceding player's best response strategy is simply to do nothing and suffer the loss in her payoff.

5.1.4 Comparison with alliances

In the previous sections, we analyze the strategic viability of budget and value concessions in the three-player General Lotto game model examined in this work. Next, we position our results against the results in [50] on mutually-beneficial alliances. This permits a comparison between the settings where mechanisms based on coordinated decision making (e.g., alliances) are strategically advantageous, and those where unilateral mechanisms (e.g., concessions) are.

Recall that the authors of [50] identify cases in which *alliances* (i.e., pre-emptive, unilateral budget transfers) are mutually beneficial to players A and B in the three-player General Lotto game. The intuition provided in [50] is that mutually-beneficial alliances occur in instances

more payoff than in the corresponding nominal game even if she continues to use her optimal allocation from Stage 1 of the nominal game.

where the difference between the ratios Φ_A/X_A and Φ_B/X_B is sufficiently large. In contrast, our findings show that beneficial concessions only exist when the ratios Φ_A/X_A and Φ_B/X_B are close (see Lemma 5.3.1 in Section 5.3.1). This comparison suggests that, if there are significant asymmetries in the players' strengths relative to the values of their respective contests, then mechanisms based on centralized decision making, such as alliances, provide strategic advantages; meanwhile, if differences in players' relative strengths are small, then unilateral mechanisms such as concessions prevail.

We confirm this intuition by plotting the parametric regions corresponding to the existence of mutually-beneficial alliances (in red), beneficial concessions (in blue) or both (in green) in Figure 5.3. The player budgets are normalized (i.e., $X_0 = 1$), $X_A, X_B \in [0, 1.2]$, $\Phi_A = 1$ and (a) $\Phi_B = 1/3$, (b) $\Phi_B = 1$, and (c) $\Phi_B = 3$. We immediately observe that it is possible for the two regions to overlap, i.e., there are parameters under which player A has both mutually-beneficial alliances and beneficial pre-commitment. We further observe that this overlap exists only when Φ_A is lower than, equal to or only slightly larger than Φ_B, e.g., as in Figure 5.3b,c. When Φ_A is much greater than Φ_B, e.g,, as in Figure 5.3a, then the two regions are disjoint.

5.2 General Lotto games with pre-allocations

Another possible perspective on strategic pre-emption is that of a heterogeneous allocation of resources, where the agent has access to either *pre-allocated* and *real-time* resources. Observe that such a setting can also be viewed as revealing strategic information about the minimum amount of resources that a player will allocate to a given battlefield.

5.2.1 Model

The *General Lotto game with pre-allocations* (GL-P) is a two-stage game with players A and B, who compete over a set of n battlefields, denoted as $\mathcal{B}$. Each battlefield $b \in \mathcal{B}$ is associated with a known valuation $w_b > 0$, which is common (symmetric) to both players. Player A is endowed with a pre-allocated resource budget $P > 0$ and a real-time resource budget $R_A > 0$.

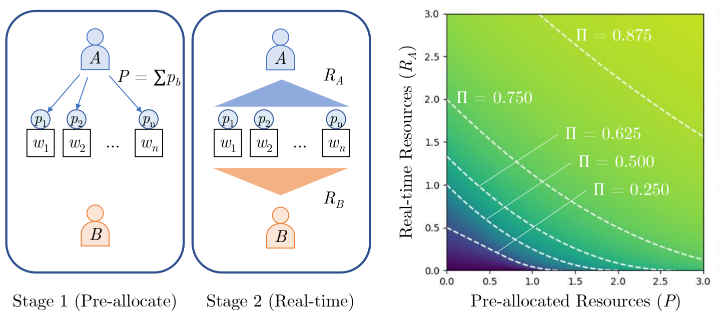

(a) (b)

Performance (Π)	(0, R_A)	(P^{eq}, 0)	P^{eq}/R_A
0.250	(0, 0.500)	(1.333, 0)	2.667
0.500	(0, 1.000)	(2.000, 0)	2.000
0.625	(0, 1.333)	(2.667, 0)	2.000
0.750	(0, 2.000)	(4.000, 0)	2.000
0.875	(0, 4.000)	(8.000, 0)	2.000

(c)

Figure 5.4: (a) The two-stage General Lotto game under consideration. Players A and B compete over n battlefields, whose valuations are given by $\{w_b\}_{b=1}^n$. In Stage 1, player A decides how to deploy P pre-allocated resources to the battlefields. Player B observes the deployment. In Stage 2, the players simultaneously decide how to deploy their real-time resources R_A and R_B, thus engaging in a General Lotto game with favoritism. (b) A contour map of player A's equilibrium payoff in the Stage 2 game under the optimal deployment of her pre-allocated resources S in Stage 1. The dashed lines indicate level curves, i.e. the set of resource pairs $(P, R_A) \in \mathbb{R}_+^2$ that achieve a given desired performance level $\Pi > 0$ (Theorem 5.2.2). Here, we have normalized the battlefield values and player B's budget such that $\sum_{b=1}^n w_b = 1$ and $qR_B = 1$. (c) This table shows the relative effectiveness of pre-allocated to real-time resources, P^{eq}/R_A. Here, P^{eq} is defined as the endowment $(P^{\text{eq}}, 0)$ (i.e. without real-time resources) that achieves the same performance Π as the endowment $(0, R_A)$ (i.e. without pre-allocated resources) for a given R_A. We find that real-time resources are at least twice as effective as pre-allocated resources, and can be arbitrarily more effective in certain settings (Corollary 5.2.1).

Player B is endowed with a real-time resource budget $R_B > 0$, but no pre-allocated resources.[5] The two stages are played as follows:

– *Stage 1:* Player A decides how to allocate her P pre-allocated resources to the battlefields, i.e., it selects a vector $\mathbf{p} = (p_1, \dots, p_n) \in \Delta_n(P) := \{\mathbf{p}' \in \mathbb{R}_+^n : \|\mathbf{p}'\|_1 = P\}$. We term the vector $\mathbf{p}$ as player A's *pre-allocation profile.* No payoffs are derived in Stage 1, and A's choice $\mathbf{p}$ becomes binding and common knowledge.

– *Stage 2:* Players A and B compete in a simultaneous-move sub-game G over $\mathcal{B}$ with their real-time resource budgets R_A, R_B. Here, both players can randomly allocate these resources as long as their expenditure does not exceed their budgets in expectation. Specifically, a strategy for player $i \in \{A, B\}$ is an n-variate (cumulative) distribution F_i over allocations $\mathbf{x}_i \in \mathbb{R}_+^n$ that

[5] Recent computational advances (see, e.g., [51]) permit the study of the scenario where both players are endowed with pre-allocated resources. In this work, we seek to provide *analytical* characterizations of equilibrium payoffs, and, thus, consider the simpler, unilateral pre-allocation setting.

satisfies

$$\mathbb{E}_{\mathbf{x}_i \sim F_i}\left[\sum_{b\in\mathcal{B}} x_{i,b}\right] \leq R_i. \tag{5.1}$$

We use $\mathcal{L}(R_i)$ to denote the set of all strategies F_i that satisfy (5.1). Given that player A chose $\mathbf{p}$ in Stage 1, the expected payoff to player A is given by

$$u_A(\mathbf{p}, F_A, F_B) := \mathbb{E}_{\substack{\mathbf{x}_A\sim F_A,\\ \mathbf{x}_B\sim F_B}}\left[\sum_{b\in\mathcal{B}} w_b \cdot I(x_{A,b}+p_b, qx_{B,b})\right] \tag{5.2}$$

where $I(a,b)=1$ if $a>b$, and $I(a,b)=0$ otherwise for any two numbers $a,b\in\mathbb{R}_+$.[6] In words, player B must overcome player A's pre-allocated resources p_b as well as player A's allocation of real-time resources $x_{A,b}$ in order to win battlefield b. The parameter $q>0$ is the relative quality of player B's real-time resources against player A's resources. When $q<1$ (resp. $q>1$), they are less (resp. more) effective than player A's resources. The payoff to player B is $u_B(\mathbf{p},F_A,F_B) = W - u_A(\mathbf{p},F_A,F_B)$, where we denote $W=\sum_{b\in\mathcal{B}} w_b$.

Stages 1 and 2 of GL-P are illustrated in Figure 5.4a. We denote an instance of GL-P as GL-P$(P,R_A,R_B,\mathbf{w})$, and note that the Stage 2 sub-game (i.e., the game with fixed pre-allocation profile) is an instance of the *General Lotto game with favoritism* [51]. For a given GL-P instance G, we define an equilibrium as any joint strategy profile $(\mathbf{p}^*, F_A^*, F_B^*) \in \Delta_n(P)\times\mathcal{L}(R_A)\times\mathcal{L}(R_B)$ that satisfies

$$\begin{aligned} &u_A(\mathbf{p}^*, F_A^*, F_B^*) \geq u_A(\mathbf{p}, F_A, F_B^*) \text{ and}\\ &u_B(\mathbf{p}^*, F_A^*, F_B^*) \geq u_B(\mathbf{p}^*, F_A^*, F_B) \end{aligned} \tag{5.3}$$

for any $\mathbf{p}\in\Delta_n(P)$, $F_A\in\mathcal{L}(R_A)$ and $F_B\in\mathcal{L}(R_B)$. Notably, player A's strategy consists of her deterministic pre-allocation profile $\mathbf{p}$ in Stage 1, as well as her randomized allocation of real-time resources F_A in Stage 2. It follows from the results in [51] that an equilibrium exists in any GL-P instance G, and that the equilibrium payoffs $\pi_i^*(G) = u_i(\mathbf{p}^*, F_A^*, F_B^*)$, $i\in\{A,B\}$,

[6] The tie-breaking rule (i.e., deciding who wins if $x_{A,b}+p_b = x_{B,b}$) can be assumed to be arbitrary, without affecting any of our results. This property is common in the General Lotto literature, see, e.g., [20, 51].

are unique. For ease of notation, we will use $\pi_i^*(P, R_A, R_B)$, $i \in \{A, B\}$, to denote players' equilibrium payoffs in G when the dependence on the vector $\mathbf{w}$ is clear.

5.2.2 Main results

Theorem 5.2.1. *Consider a GL-P game instance with $P, R_A, R_B > 0$, and $\mathbf{w} \in \mathbb{R}^n_{++}$. The following conditions characterize player A's equilibrium payoff $\pi_A^*(P, R_A, R_B)$:*

1. *If $qR_B < P$, or $qR_B \geq P$ and $R_A \geq \frac{2(qR_B-P)^2}{P+2(qR_B-P)}$, then $\pi_A^*(P, R_A, R_B)$ is*

$$W \cdot \left(1 - \frac{qR_B}{2R_A}\left(\frac{R_A + \sqrt{R_A(R_A+2P)}}{P + R_A + \sqrt{R_A(R_A+2P)}}\right)^2\right). \tag{5.4}$$

2. *Otherwise, $\pi_A^*(P, R_A, R_B)$ is*

$$W \cdot \frac{R_A}{2(qR_B - P)}. \tag{5.5}$$

As a consequence of the above result, we are able to characterize expressions for the level curves of the function $\pi_A^*(P, R_A, R_B)$. That is, for a desired performance level $\Pi \geq 0$ and fixed R_B, we provide the set of all pairs (P, R_A) such that $\pi_A^*(P, R_A, R_B) = \Pi$.

Theorem 5.2.2. *Given any $R_B > 0$ and $\mathbf{w} \in \mathbb{R}^n_{++}$, fix a desired performance level $\Pi \in [0, W]$. The set of all pairs $(P, R_A) \in \mathbb{R}^2_+$ that satisfy $\pi_A^*(P, R_A, R_B) = \Pi$ is given by the following conditions:*

If $0 \leq \Pi < \frac{W}{2}$, then

$$R_A = \begin{cases} \frac{2\Pi}{W}(qR_B - P) & \text{if } P \in \left[0, \frac{(W-2\Pi)qR_B}{W-\Pi}\right) \\ \frac{(qR_BW-(W-\Pi)P)^2}{2qR_B(W-\Pi)W} & \text{if } P \in \left[\frac{(W-2\Pi)qR_B}{W-\Pi}, \frac{WqR_B}{W-\Pi}\right] \end{cases} \tag{5.6}$$

If $\frac{W}{2} \leq \Pi \leq W$, then

$$R_A = \frac{(qR_BW - (W-\Pi)P)^2}{2qR_B(W-\Pi)W}, \text{ if } P \in \left[0, \frac{WqR_B}{W-\Pi}\right] \tag{5.7}$$

If $P > \frac{WqR_B}{W-\Pi}$, then $\pi_A^*(P, R_A, R_B) > \Pi$ for any $R_A \geq 0$.

We can use the result in Theorem 5.2.2 to obtain an expression for the relative effectiveness of pre-allocated and real-time resources when these are deployed in isolation. In the following corollary, we provide this expression, and observe that real-time resources are at least twice as valuable as pre-allocated resources, and can be arbitrarily more valuable in specific settings:

Corollary 5.2.1. *For given $R_A, R_B > 0$, the unique value $P^{\text{eq}} > 0$ such that $\pi_A^*(P^{\text{eq}}, 0, R_B) = \pi_A^*(0, R_A, R_B)$ is characterized by the following expression:*

$$P^{\text{eq}} = \begin{cases} 2R_A & \text{if } R_A > qR_B, \\ \frac{2(qR_B)^2}{2qR_B - R_A} & \text{if } R_A \leq qR_B. \end{cases} \tag{5.8}$$

Notably, the ratio P^{eq}/R_A is lower-bounded by 2, and $P^{\text{eq}}/R_A \to \infty$ as $R_A \to 0^+$.

5.2.3 Optimal investment decisions

In this section, we consider a scenario where player A has an opportunity to make an investment decision regarding its resource endowments. That is, the pair $(P, R_A) \in \mathbb{R}_+^2$ is a strategic choice made by A before the game GL-P$(P, R_A, R_B, \mathbf{w})$ is played. Given a monetary budget $X_A > 0$ for player A, any pair (P, R_A) must belong to the following set of feasible investments:

$$\mathcal{I}(X_A) := \{(P, R_A) : R_A + cP \leq X_A\} \tag{5.9}$$

where $c \geq 0$ is the per-unit cost for purchasing pre-allocated resources, and we assume the per-unit cost for purchasing real-time resources is 1 without loss of generality. We are interested in characterizing player A's optimal investment subject to the above cost constraint, and given player B's resource endowment $R_B > 0$. This is formulated as the following optimization problem:

$$\pi_A^{\text{opt}} := \max_{(P, R_A) \in \mathcal{I}(X_A)} \pi_A^*(P, R_A, R_B). \tag{5.10}$$

In the result below, we derive the complete solution to the optimal investment problem (5.10).

Corollary 5.2.2. *Fix a monetary budget $X_A > 0$, relative per-unit cost $c > 0$, and $R_B > 0$ real-time resources for player B. Then, player A's optimal investment in pre-allocated resources for the optimization problem in* (5.10) *under the linear cost constraint in* (5.9) *is*

$$P^* = \begin{cases} (1 - \frac{c}{2-c})\frac{X_A}{c}, & \textit{if } c < t \\ \in [0, (1 - \frac{c}{2-c})\frac{X_A}{c}], & \textit{if } c = t \\ 0, & \textit{if } c > t \end{cases} \cdot \tag{5.11}$$

where $t := \min\{1, \frac{X_A}{qR_B}\}$. *The optimal investment in real-time resources is* $R_A^* = X_A - cP^*$. *The resulting payoff* π_A^{opt} *to player A is given by*

$$W \cdot \begin{cases} 1 - \frac{qR_B}{2X_A}c(2-c), & \textit{if } c < t \\ 1 - \frac{qR_B}{2X_A}, & \textit{if } c \geq t \textit{ and } \frac{X_A}{qR_B} \geq 1 \\ \frac{X_A}{2qR_B}, & \textit{if } c \geq t \textit{ and } \frac{X_A}{qR_B} < 1 \end{cases} \cdot \tag{5.12}$$

The above solution is obtained by leveraging the level set characterization from Theorem 5.2.2, and the fact that the level sets are strictly decreasing and convex for $\Pi \in (0, W)$. The budget constraint $\mathcal{I}(X_A)$ is a line segment in $\mathbb{R}_+^2$, and we thus seek the level curve that lies tangent to the segment. In cases where the cost c is sufficiently high, no level curve lies tangent to $\mathcal{I}(X_A)$, and, thus, player A invests all of her budget in real-time resources.

5.3 Chapter proofs

5.3.1 Proofs from Section 5.1

Many of the proofs depend on player 0's optimal allocation $(X_{0,A}^*, X_{0,B}^*) \in \mathbb{R}_{\geq 0}^2$ in Stage 1 of the three-player General Lotto game, which was derived in [50]. Observe that this represents

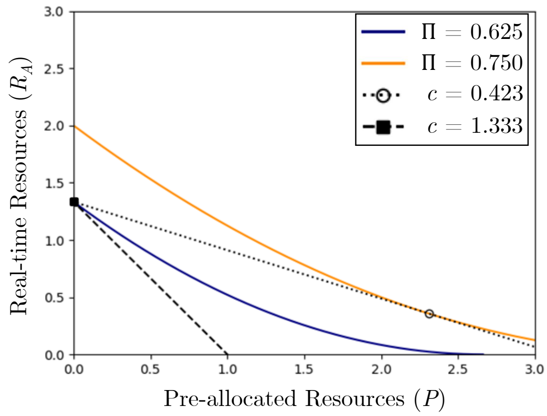

Figure 5.5: The optimal investment $(P^*, R_A^*) \in \mathbb{R}_+^2$ subject to the linear cost constraint in (5.9). Here, we consider the optimal investment problem when $X_A = 4/3$, $qR_B = W = 1$, and $c \in \{0.423, 1.333\}$. Observe that the set of feasible investments $\mathcal{I}(X_A)$ is the line segment connecting $(0, X_A)$ and $(X_A/c, 0)$. The optimal investment lies on the level curve tangent to this line segment. For example, when $c = 0.423$, the optimal investment is $(2.309, 0.357)$ (unfilled circle), as $\mathcal{I}(X_A)$ (dotted, black line) is tangent to the level curve with $\Pi = 0.750$ (solid, orange line). For sufficiently high cost c, $\mathcal{I}(X_A)$ will not be tangent to any level curve, and the optimal investment is $(0, X_A)$. For example, when $c = 1.333$, observe that $\mathcal{I}(X_A)$ (dashed, black line) is not tangent even to the level curve with $\Pi = 0.625$ (solid, blue line), and the optimal investment is $(0, 4/3)$ (filled square).

a sub-game perfect equilibrium decision of player 0 in Stage 1 given the budgets and values arising from the Stage 0 choices of players A and B. We detail their result in the following fact for the reader's convenience:

Fact 5.3.1 (Optimal Stage 1 Allocation [50])**.** *Consider the three-player General Lotto game G with normalized player budgets (i.e., $X_0 = 1$). Let $\mathcal{R}_{1i}(G)$, $\mathcal{R}_{2i}(G)$, $\mathcal{R}_{3i}(G)$ and $\mathcal{R}_4$, $i \in \{A, B\}$, denote the following regions:*

$$\begin{aligned}
&\mathcal{R}_{1i}(G) := \\
&\quad \{(X_i, X_{-i}) \ s.t.\ \Phi_i/\Phi_{-i} > \max\{(X_i)^2, 1\}/(X_i X_{-i})\} \\
&\quad \cup \{(X_i, X_{-i}) \ s.t.\ X_i < 1 \ and \ \Phi_i/\Phi_{-i} = 1/(X_i X_{-i})\}
\end{aligned}$$

$$\begin{aligned}\mathcal{R}_{2i}(G) &:= \{(X_i, X_{-i}) \text{ s.t. } \Phi_i/\Phi_{-i} > X_i/X_{-i} \\ &\qquad \text{and } 0 < 1 - \sqrt{\Phi_i X_i X_{-i}/\Phi_{-i}} \leq X_{-i}\} \\ \mathcal{R}_{3i}(G) &:= \{(X_i, X_{-i}) \text{ s.t. } \Phi_i/\Phi_{-i} \geq X_i/X_{-i} \\ &\qquad \text{and } 1 - \sqrt{\Phi_i X_i X_{-i}/\Phi_{-i}} > X_{-i}\} \\ \mathcal{R}_4(G) &:= \{(X_i, X_{-i}) \text{ s.t. } \Phi_i/\Phi_{-i} = X_i/X_{-i} \\ &\qquad \text{and } X_i + X_{-i} \geq 1\}.\end{aligned}$$

Player 0's optimal allocation $(X^*_{0,A}, X^*_{0,B})$ in Stage 1 is determined in closed-form as follows:

- *If $(X_i, X_{-i}) \in \mathcal{R}_{1i}(G)$, then $X^*_{0,i} = 1$;*
- *If $(X_i, X_{-i}) \in \mathcal{R}_{2i}(G)$, then*
$$X^*_{0,i} = \sqrt{\frac{\Phi_i X_i X_{-i}}{\Phi_{-i}}};$$
- *If $(X_i, X_{-i}) \in \mathcal{R}_{3i}(G)$, then*
$$X^*_{0,i} = \frac{\sqrt{\Phi_i X_i}}{\sqrt{\Phi_i X_i} + \sqrt{\Phi_{-i} X_{-i}}};$$
- *If $(X_i, X_{-i}) \in \mathcal{R}_4(G)$, then select any $X^*_{0,i} \in [\max\{0, 1 - X_{-i}\}, \min\{1, X_i\}]$;*

where $X^*_{0,-i} = 1 - X^*_{0,i}$ in all the above cases.

Though the indices of players A and B are permuted in the region definitions, observe that the expressions for player 0's optimal allocation – as well as the players' payoffs – are identical in $\mathcal{R}_{3B}(G)$ and $\mathcal{R}_{3C}(G)$. Thus, we simplify the notation with $\mathcal{R}_3(G) := \mathcal{R}_{3B}(G) \cup \mathcal{R}_{3C}(G)$.

Recall from the previous section that our proposed models for budget and value concessions are normal-form games under the assumption that players 0, A and B play according to an SPE in Stages 1 and 2. In particular, observe that the players' payoffs at the end of Stage 2 are uniquely determined by player 0's allocation decision in Stage 1, since the equilibrium payoffs in the General Lotto games are unique (Theorem 3.1.1). Crucially, as we summarize in

Fact 5.3.1, player 0's optimal allocation in Stage 1 – and, thus, the players' corresponding SPE payoffs in Stages 1 and 2 – is unique except for when $(X_A, X_B) \in \mathcal{R}_4(G)$. Accordingly, in the forthcoming proofs, we treat the setting when $(X_A, X_B) \in \mathcal{R}_4(G)$ (i.e., the setting when 0 has multiple optimal allocations) explicitly, and show that the results in Theorem 5.1.1 and 5.1.2 hold irrespective of the imposed SPE selection rule.

Proof of Theorem 5.1.1 For any three-player General Lotto game with budget concessions G, the proof amounts to showing that player A's payoff is nonincreasing in x_A for any concession profile $(x_A, x_B) \in [0, X_A] \times [0, X_B]$ when (X_A, X_B) is in any of the regions $\mathcal{R}_j(G)$, $j \in \{1A, 1B, 2A, 2B, 3, 4\}$, identified in Fact 5.3.1. This is sufficient as the same reasoning applies symmetrically to player B. Throughout, we denote the nominal three-player General Lotto game (i.e., the game without any budget concessions) as G, and the three-player General Lotto game with budget concessions corresponding to G under a given budget concession profile (x_A, x_B) as $G'(x_A, x_B)$.

We begin with our treatment of settings where players 0, A and B have non-unique SPE payoffs in Stages 1 and 2, corresponding with settings where player 0 has more than one possible optimal allocation in Stage 1.

– *Non-unique SPE payoffs:* Recall from Fact 5.3.1 that players' SPE payoffs in Stages 1 and 2 are non-unique only when $(X_A - x_A, X_B - x_B) \in \mathcal{R}_4(G'(x_A, x_B))$. For any of these SPEs, we show that – regardless of the SPE selection rule – A prefers to concede $x < x_A$ if $x_A > 0$ and prefers to concede x_A if $x_A = 0$. According to Fact 5.3.1, when $(X_A - x_A, X_B - x_B) \in \mathcal{R}_4(G'(x_A, x_B))$, any allocation satisfying $X^*_{0,A} \in [\max\{0, 1 - (X_B - x_B)\}, \min\{1, X_A - x_A\}]$ and $X^*_{0,B} = 1 - X^*_{0,A}$ is an optimal allocation for player 0. We show that (i) if A concedes x with $0 < x < x_A$, then 0 allocates at most $\max\{0, 1 - (X_B - x_B)\}$ to the battlefields in front $\mathcal{B}_A$, and (ii) if A concedes $x > x_A$, then 0 allocates at least $\min\{1, X_A - x\}$ to front $\mathcal{B}_A$. It follows that A prefers to concede $x < x_A$ since A's Stage 2 payoff is nonincreasing in the ratio $X^*_{0,A}/(X_A - x)$ (Theorem 3.1.1).

In scenario (i), observe that any budget concession x with $0 \leq x < x_A$ would satisfy either $(X_A - x, X_B - x_B) \in \mathcal{R}_{1B}(G'(x, x_B))$ or $(X_A - x, X_B - x_B) \in \mathcal{R}_{2B}(G'(x, x_B))$, since $X_A - x_A + X_B - x_B \geq 1$ must hold if $(X_A - x_A, X_B - x_B) \in \mathcal{R}_4(G'(x, x_B))$. If $(X_A - x, X_B - x_B) \in \mathcal{R}_{1B}(G'(x, x_B))$, then $X^*_{0,B} = 1$ and $X^*_{0,A} = 1 - X^*_{0,B} = 0$. Moreover, if $(X_A - x, X_B - x_B) \in \mathcal{R}_{2B}(G'(x, x_B))$, then $X^*_{0,B} \geq X_B - x_B$ and, thus, $X^*_{0,A} \leq 1 - (X_B - x_B)$.

Likewise, in scenario (ii), $x > x_A$ satisfies either $(X_A - x, X_B - x_B) \in \mathcal{R}_{1A}(G'(x, x_B))$, $(X_A - x, X_B - x_B) \in \mathcal{R}_{2A}(G'(x, x_B))$, or $(X_A - x, X_B - x_B) \in \mathcal{R}_3(G'(x, x_B))$. If $(X_A - x, X_B - x_B) \in \mathcal{R}_{1A}(G'(x, x_B))$, then $X^*_{0,A} = 1$. Furthermore, if $(X_A - x, X_B - x_B) \in \mathcal{R}_{2A}(G'(x, x_B) \cup \mathcal{R}_3(G'(x, x_B))$ then $X^*_{0,A} > X_A - x$.

Next, we consider the setting when player 0's optimal allocation in Stage 1 is unique, and, thus, the players' corresponding SPE payoffs in Stages 1 and 2 are unique.

– *Unique SPE payoffs:* We begin with the scenario $(X_A - x_A, X_B - x_B) \in \mathcal{R}_{1B}(G'(x_A, x_B))$. Recall that, in this scenario, player 0 commits no budget to the battlefields in front $\mathcal{B}_A$. Thus, player A's payoff before the concession is Φ_A, the highest possible payoff. Furthermore, if $(X_A - x_A, X_B - x_B) \in \mathcal{R}_{1B}(G'(x_A, x_A))$, then $(X_A, X_B - x_B) \in \mathcal{R}_{1B}(G'(0, x_B))$ must also hold, since the expression $1 - \sqrt{\Phi_B(X_A - x)(X_B - x_B)/\Phi_A}$ is increasing in x.

In all other regions, we show that player A's payoff $\Pi^{\mathrm{bc}}_A(x, x_B)$ is strictly decreasing in x by checking the first partial derivative with respect to $x \geq 0$:

If $(X_A - x, X_B - x_B) \in \mathcal{R}_{1A}(G'(x, x_B))$ and $X_A - x > 1$, then

$$\frac{\partial}{\partial x}\Phi_A\left[1 - \frac{1}{2(X_A - x)}\right] = -\frac{\Phi_A}{2(X_A - x)^2} < 0.$$

If $(X_A - x, X_B - x_B) \in \mathcal{R}_{1A}(G'(x, x_B))$ and – conversely – $X_A - x \leq 1$, then

$$\frac{\partial}{\partial x}\frac{\Phi_A(X_A - x)}{2} = -\frac{\Phi_A}{2} < 0.$$

If $(X_A - x, X_B - x_B) \in \mathcal{R}_{2A}(G'(x, x_B))$, then

$$\frac{\partial}{\partial x}\frac{\Phi_A(X_A - x)}{2\sqrt{\frac{\Phi_A(X_A - x)(X_B - x_B)}{\Phi_B}}} = -\frac{\Phi_A}{4\sqrt{\frac{\Phi_A(X_A - x)(X_B - x_B)}{\Phi_B}}} < 0.$$

If $(X_A - x, X_B - x_B) \in \mathcal{R}_{2B}(G'(x, x_B))$, then

$$\frac{\partial}{\partial x}\Phi_A\left[1-\frac{1-\sqrt{\frac{\Phi_B(X_A-x)(X_B-x_B)}{\Phi_A}}}{2(X_A-x)}\right]$$
$$= -\frac{\Phi_A\left[2-\sqrt{\frac{\Phi_B(X_A-x)(X_B-x_B)}{\Phi_A}}\right]}{4(X_A-x)^2},$$

which is strictly negative as the condition

$$1-\sqrt{\Phi_B(X_A-x)(X_B-x_B)/\Phi_A} \geq 0$$

must hold if $(X_A - x, X_B - x_B) \in \mathcal{R}_{2B}(G')$. Finally, if $(X_A - x, X_B - x_B) \in \mathcal{R}_3(G'(x, x_B))$, then

$$\frac{\partial}{\partial x}\frac{\Phi_A(X_A-x)}{2\frac{\sqrt{\Phi_A(X_A-x)}}{\sqrt{\Phi_A(X_A-x)}+\sqrt{\Phi_B(X_B-x_B)}}} = -\frac{\Phi_A}{2}-\frac{\Phi_A\sqrt{\Phi_B(X_B-x_B)}}{4\sqrt{\Phi_A(X_A-x)}},$$

which is strictly negative. Since there is no payoff improvement associated with making a budget concession in any of the regions (including $\mathcal{R}_4$), it follows that making no budget concession maximizes player A's payoff. This concludes the proof. □

Proof of Theorem 5.1.2 The proof is presented as two lemmas, which we summarize below:

1. In Lemma 5.3.1, we provide necessary (under mild assumptions on the SPE selection rule) and sufficient conditions on the three-player General Lotto games in which either player $i \in \{A, B\}$ can improve her payoff by making a value concession; and,

2. In Lemma 5.3.2, we use the result in Lemma 5.3.1 to show that there is at most one pure strategy Nash equilibrium that coincides with a pure strategy concession profile of the form $(0, 0)$, $(y_A, 0)$ or $(0, y_B)$ in the General Lotto game with value concessions.

Observe that the claim in Theorem 5.1.2 follows after combining these two lemmas, as there exist instances in which value concessions can improve a player's payoff (Lemma 5.3.1), and –

if there exist any pure strategy Nash equilibria – the unique pure strategy Nash equilibrium in the game with value concessions is of the form $(0,0)$, $(y_A,0)$ or $(0,y_B)$ (Lemma 5.3.2).

The claims of the forthcoming lemmas, and several arguments within the proofs, are dependent on a parameter $\epsilon \to 0^+$ that we use to place the point (X_A, X_B) strictly in the interior of a particular region $\mathcal{R}_j$ to take derivatives and make payoff comparisons. Note that, typically, the use of such an argument would require a definition of approximate best-response strategies, and approximate equilibrium to formally establish convergence and the solution concept. However, we only use these arguments to rule out the existence of certain equilibrium structures. In fact, from the proof of Lemma 5.3.2, we find that the only emergent pure strategy Nash equilibria do not require such arguments at all. For this reason, and for conciseness, we omit any further discussion on this topic.

Lemma 5.3.1. *Consider any three-player General Lotto game G with normalized player budgets (i.e., $X_0 = 1$), and $(X_A, X_B) \in \mathcal{R}_j(G)$ such that $j \in \{1A, 1B, 2A, 2B, 3\}$. Let $y_1^* = \Phi_A - \Phi_B X_A/X_B$, $y_2^* = \Phi_A - \Phi_B X_A X_B/(2-4X_A)^2$, $y_3^* = \Phi_A - \Phi_B X_A X_B$ and $y_4^* = \Phi_A - \Phi_B X_A X_B/(1-X_A)^2$. Further, let $\epsilon \to 0^+$. It holds that player A can improve her payoff by unilaterally deviating from the value concession profile $(0,0)$ if and only if*

$$\Pi_A^{\text{vc}}(y_A^{\text{opt}}, 0) > \Pi_A^{\text{vc}}(0,0),$$

where the value $y_A^{\text{opt}} \in [0, \Phi_A]$ is defined according to the following conditions:

(i) If $(X_A, X_B) \in \mathcal{R}_{1A}(G) \cup \mathcal{R}_{2A}(G)$ and $X_B \geq 1$, then $y_A^{\text{opt}} = y_1^* + \epsilon$;

(ii) If $(X_A, X_B) \in \mathcal{R}_{1A}(G) \cup \mathcal{R}_{2A}(G)$, $X_B < 1$, $X_A + X_B \geq 1$ and

$$\frac{\partial}{\partial y}\Pi_A^{\text{vc}}(y,0)\bigg|_{y=y_1^*+\epsilon} \leq 0,$$

then $y_A^{\text{opt}} = y_1^* + \epsilon$;

(iii) *If* $(X_A, X_B) \in \mathcal{R}_{1A}(G) \cup \mathcal{R}_{2A}(G)$, $X_B < 1$, $X_A + X_B \geq 1$ *and*

$$\frac{\partial}{\partial y}\Pi_A^{\mathrm{vc}}(y,0)\bigg|_{y=y_1^*+\epsilon} > 0,$$

then $y_A^{\mathrm{opt}} = \min\{y_2^*, y_3^*\}$;

(iv) *If* $(X_A, X_B) \in \mathcal{R}_{1A}(G) \cup \mathcal{R}_{2A}(G) \cup \mathcal{R}_3(G)$, $X_B + X_C < 1$ *and*

$$\frac{\partial}{\partial y}\Pi_A^{\mathrm{vc}}(y,0)\bigg|_{y=y_4^*+\epsilon} > 0,$$

then $y_A^{\mathrm{opt}} = \min\{y_2^*, y_3^*\}$;

(v) *If* $(X_A, X_B) \in \mathcal{R}_{2B}(G)$ *and*

$$\frac{\partial}{\partial y}\Pi_A^{\mathrm{vc}}(y,0)\bigg|_{y=0} > 0,$$

then $y_A^{\mathrm{opt}} = \min\{y_2^*, y_3^*\}$;

(vi) *Otherwise,* $y_A^{\mathrm{opt}} = 0$.

The same conditions hold symmetrically for player B.

Proof. Throughout the proof, we will use G to denote the nominal three-player General Lotto game (with no concessions), and use $G'(y)$ to denote the three-player General Lotto game with value concessions under the concession profile $(y, 0)$. We divide the proof in two parts. In Part 1, we demonstrate if and when player A can use a value concession to improve her payoff when the point (X_A, X_B) falls in each of the regions $\mathcal{R}_j(G)$, $j \in \{1A, 1B, 2A, 2B, 3, 4\}$. Then, in Part 2, we identify the optimal value concession value $y^{\mathrm{opt}} \neq 0$, and thus the maximum payoff that player A can achieve under a value concession. Thus, the necessary and sufficient conditions in the claim amount to verifying that player A's maximum achievable payoff using the value concession $y^{\mathrm{opt}} \neq 0$ is strictly greater than player A's payoff under no concession (i.e., in the nominal game).

Part 1. Observe that, for sufficiently large value concession $y \in [0, \Phi_A]$, the point (X_A, X_B) can "transition" from $\mathcal{R}_j(G)$ to $\mathcal{R}_k(G'(y))$, for $j \neq k \in \{1A, 1B, 2A, 2B, 3, 4\}$, as the median line $X_B = \Phi_B X_A/(\Phi_A - y)$ which divides regions $\mathcal{R}_{1B}(G'(y)) \cup \mathcal{R}_{2B}(G'(y))$ and regions $\mathcal{R}_{1C}(G'(y)) \cup \mathcal{R}_{2C}(G'(y))$ rotates counter-clockwise about the origin as y is increased, and the value $1 - \sqrt{(\Phi_A - y)X_A X_B/\Phi_B}$ (resp., $1 - \sqrt{\Phi_B X_A X_B/(\Phi_A - y)}$) is increasing (resp., decreasing) in y. For example, a point $(X_A, X_B) \in \mathcal{R}_{1A}(G) \cup \mathcal{R}_{2A}(G)$ will transition to $\mathcal{R}_4(G'(y))$ and then to $\mathcal{R}_{1B}(G'(y)) \cup \mathcal{R}_{2B}(G'(y))$ as y is increased. However, note that a point $(X_A, X_B) \in \mathcal{R}_3(G'(y))$ will never transition to or from $\mathcal{R}_4(G'(y))$ as y is increased since $X_A + X_B < 1$ in $\mathcal{R}_3$, and $X_A + X_B \geq 1$ in $\mathcal{R}_4$.

As in the proof of Theorem 5.1.1, we first resolve the setting where players' SPE payoffs in Stages 1 and 2 are non-unique, i.e., $(X_A, X_B) \in \mathcal{R}_4(G'(y))$ for some $y \geq 0$. We show that we need not explicitly consider the transition of the point (X_A, X_B) through the region $\mathcal{R}_4$ in our study of unilateral best response strategies, since player A either has a different concession strategy $y' \neq y$ such that $(X_A, X_B) \notin \mathcal{R}_4(G'(y'))$ and her payoff is improved, or the A's payoff $\Pi^{\text{vc}}(y, 0)$ is continuous in the transition from $\mathcal{R}_4(G'(y))$ to $\mathcal{R}_{1B}(G'(y)) \cup \mathcal{R}_{2B}(G'(y))$.

– *Non-unique SPE payoffs:* We begin by showing that player A achieves at least the same payoff by conceding $y' = y + \epsilon$, $\epsilon \to 0^+$, when $(X_A, X_B) \in \mathcal{R}_4(G'(y))$, $y \geq 0$. Recall that the players' payoffs are only ever non-unique when (X_A, X_B) is in the region $\mathcal{R}_4$, because any allocation that satisfies $\max\{0, 1 - X_B\} \leq X^*_{0,A} \leq \min\{X_A, 1\}$ and $X^*_{0,B} = 1 - X^*_{0,A}$ is optimal for player 0. When $X_B \geq 1$, observe that $\max\{0, 1 - X_B\} = 0$ and thus $X^*_{0,A} \geq 0$. Additionally, if $(X_A, X_B) \in \mathcal{R}_4(G'(y))$ and $X_B \geq 1$, then it follows that $(X_A, X_B) \in \mathcal{R}_{1B}(G'(y'))$. Recall from Fact 5.3.1 that $X^*_{0,A} = 0$ when (X_A, X_B) is in the region $\mathcal{R}_{1B}$. Since the allocation $X^*_{0,A}$ when (X_A, X_B) is in $\mathcal{R}_4(G'(y))$ is at least the same as in $\mathcal{R}_{1B}(G(y))$, it immediately follows that $\Pi^{\text{vc}}_A(y', 0) \geq \Pi^{\text{vc}}_A(y, 0)$. Similarly, if $(X_A, X_B) \in \mathcal{R}_4(G'(y))$ and $X_B < 1$, then $X^*_{0,A} \geq \max\{0, 1 - X_B\} = 1 - X_B$ and $(X_A, X_B) \in \mathcal{R}_{2B}(G'(y'))$. Since $X^*_{0,B} > X_B$ – and, thus, $X^*_{0,A} < 1 - X_B$ – in $\mathcal{R}_{2B}$, it immediately follows that $\Pi^{\text{vc}}_A(y', 0) > \Pi^{\text{vc}}_A(y, 0)$. Thus, when $(X_A, X_B) \in \mathcal{R}_4(G'(y))$, the function $\Pi^{\text{vc}}_A(y, 0)$ is either continuous in y as the point (X_A, X_B)

transitions from $\mathcal{R}_4(G'(y))$ to $\mathcal{R}_{1B}(G'(y')) \cup \mathcal{R}_{2B}(G'(y'))$, or it is discontinuous in this transition and $\Pi_A^{\text{vc}}(y, 0)$ strictly increases.

From the above, it follows that we need not consider $(X_A, X_B) \in \mathcal{R}_4(G'(y))$ for the remainder of the proof (i.e., we need not consider the setting where players' SPE payoffs in Stages 1 and 2 are non-unique), since – depending on the SPE selection rule – either player A can strictly improve her payoff by conceding $y' = y + \epsilon$, $\epsilon \to 0^+$, such that $(X_A, X_B) \in \mathcal{R}_{1B}(G'(y')) \cup \mathcal{R}_{2B}(G'(y'))$, or the transition is continuous in A's payoff, i.e., $\Pi_A^{\text{vc}}(y, 0) = \Pi_A^{\text{vc}}(y', 0)$. Additionally, we have shown that Conditions (i)–(iii) in the claim hold for when $(X_A, X_B) \in \mathcal{R}_4(G)$.

– *Unique SPE payoffs*: We continue by making some important observations about the first and second partial derivatives in the regions $\mathcal{R}_j$, $j \in \{1A, 1B, 2A, 2B, 3\}$. The first partial derivative of player A's payoff with respect to $y > 0$ is always strictly negative when (X_A, X_B) is in the regions $\mathcal{R}_{1A}(G'(y))$, $\mathcal{R}_{1B}(G'(y))$, $\mathcal{R}_{2A}(G'(y))$ or $\mathcal{R}_3(G'(y))$:

If $(X_A, X_B) \in \mathcal{R}_{1A}(G'(y))$ and $X_A > 1$, then

$$\frac{\partial}{\partial y}\Pi_A^{\text{vc}}(y, 0) = -\frac{1}{2X_A},$$

else, if $(X_A, X_B) \in \mathcal{R}_{1A}(G'(y))$ and $X_A \leq 1$, then

$$\frac{\partial}{\partial y}\Pi_A^{\text{vc}}(y, 0) = -\frac{X_A}{2}.$$

If $(X_A, X_B) \in \mathcal{R}_{1B}(G'(y))$, then

$$\frac{\partial}{\partial y}\Pi_A^{\text{vc}}(y, 0) = -1.$$

If $(X_A, X_B) \in \mathcal{R}_{2A}(G'(y))$, then

$$\frac{\partial}{\partial y}\Pi_A^{\text{vc}}(y, 0) = -\frac{X_A}{4\sqrt{\frac{(\Phi_A - y)X_A X_B}{\Phi_B}}}.$$

If $(X_A, X_B) \in \mathcal{R}_3(G'(y))$, then

$$\frac{\partial}{\partial y}\Pi_A^{\mathrm{vc}}(y,0) = -\frac{X_A}{2}\left[1 + \frac{1}{2}\frac{\sqrt{(\Phi_A - y)X_A}}{\sqrt{\Phi_B X_B}}\right].$$

Furthermore, the second partial derivative of player A's payoff with respect to $y > 0$ is always negative when $(X_A, X_B) \in \mathcal{R}_{2B}(G'(y))$:

$$\frac{\partial^2}{\partial y^2}\Pi_A^{\mathrm{vc}}(y,0) = -\frac{\sqrt{\frac{\Phi_B X_A X_B}{\Phi_A - y}}}{8(\Phi_A - y)X_A} < 0.$$

Now we are ready to examine the possible transitions between the regions following a value concession, and the corresponding impact on player A's payoff.

Consider the scenario when $(X_A, X_B) \in \mathcal{R}_3(G)$. Observe that from region $\mathcal{R}_3(G)$, the only possible transitions as we increase y are to regions $\mathcal{R}_{1B}(G'(y))$ and $\mathcal{R}_{2B}(G'(y))$ since the values $1 - \sqrt{(\Phi_A - y)X_A X_B/\Phi_B}$ and $1 - \sqrt{\Phi_B X_A X_B/(\Phi_A - y)}$ are increasing and decreasing in y, respectively. Furthermore, the players' payoffs in the transition from region $\mathcal{R}_3(G)$ to region $\mathcal{R}_{2B}(G'(y))$ are continuous but not necessarily smooth, i.e., the first partial derivatives with respect to y need not be continuous. Since the first partial derivative of player A's payoff with respect to y is always negative for $(X_A, X_B) \in \mathcal{R}_3(G'(y))$, it follows that all beneficial value concessions in this setting must involve a transition of (X_A, X_B) to either $\mathcal{R}_{1B}(G'(y))$ or $\mathcal{R}_{2B}(G'(y))$.

Next, consider the scenario when either $(X_A, X_B) \in \mathcal{R}_{1A}(G)$ or $(X_A, X_B) \in \mathcal{R}_{2A}(G)$. As we increase y in this setting, note that if $X_A > 1$ then (X_A, X_B) must be in $\mathcal{R}_{1A}(G)$ and will transition to either $\mathcal{R}_{1B}(G'(y))$ or $\mathcal{R}_{2B}(G'(y))$ without first transitting through $\mathcal{R}_{2A}(G'(y))$, and that if $X_B > 1$ then (X_A, X_B) will transit to $\mathcal{R}_{1B}(G'(y))$ without passing through $\mathcal{R}_{2B}(G'(y))$. Further note that, in the scenarios where either $(X_A, X_B) \in \mathcal{R}_{1A}(G)$ or $(X_A, X_B) \in \mathcal{R}_{2A}(G)$ and the point (X_A, X_B) transitions through $\mathcal{R}_3(G'(y))$ as we increase y (i.e., when $X_A + X_B \leq 1$), the partial derivative of player A's payoff with respect to y must remain negative at least until it reaches the boundary between $\mathcal{R}_3(G'(y))$ and $\mathcal{R}_{2B}(G'(y))$. Thus, there can be no beneficial value concession such that (X_A, X_B) transitions from either $\mathcal{R}_{1A}(G)$ or $\mathcal{R}_{2A}(G)$ to $\mathcal{R}_3(G'(y))$. Likewise, there is no beneficial value concession such that (X_A, X_B) transits from $\mathcal{R}_{1A}(G)$ to $\mathcal{R}_{2A}(G'(y))$. Thus, for a value concession y to improve A's payoff when (X_A, X_B) is in either $\mathcal{R}_{1A}(G)$ or $\mathcal{R}_{2A}(G)$, if $X_B > 1$, then (X_A, X_B) must transition to $\mathcal{R}_{1B}(G'(y))$, and, if $X_B \leq 1$, then (X_A, X_B) must transition to either $\mathcal{R}_{1B}(G'(y))$ or $\mathcal{R}_{2B}(G'(y))$. Furthermore, for a value concession y to improve A's payoff when $(X_A, X_B) \in \mathcal{R}_3(G)$, (X_A, X_B) must transition to either $\mathcal{R}_{1B}(G'(y))$ or $\mathcal{R}_{2B}(G'(y))$.

Finally, consider the scenarios when either $(X_A, X_B) \in \mathcal{R}_{1B}(G)$ or $(X_A, X_B) \in \mathcal{R}_{2B}(G)$. Note that if $(X_A, X_B) \in \mathcal{R}_{1B}(G)$, then (X_A, X_B) remains in $\mathcal{R}_{1B}(G'(y))$ as y is increased since $1 - \sqrt{\Phi_B X_A X_B / (\Phi_A - y)}$ is decreasing in y. Similarly, $(X_A, X_B) \in \mathcal{R}_{2B}(G)$ will either remain in $\mathcal{R}_{2B}(G'(y))$ or transition to $\mathcal{R}_{1B}(G'(y))$ as y is increased. We showed above that the first partial derivative of player A's payoff with respect to y when $(X_A, X_B) \in \mathcal{R}_{1B}(G'(y))$ is negative, and so there can be no beneficial value concession if $(X_A, X_B) \in \mathcal{R}_{1B}(G)$.

Part 2. In the previous part of the proof, we observed that transitions between regions can occur following value concessions. In particular, we demonstrate that a value concession y cannot improve A's payoff if (X_A, X_B) is in either $\mathcal{R}_{1A}(G'(y))$, $\mathcal{R}_{2A}(G'(y))$, or $\mathcal{R}_3(G'(y))$, or $(X_A, X_B) \in \mathcal{R}_{1B}(G)$. We will use further insights garnered from the previous part in the remainder of the proof.

Consider the scenario where $X_B \geq 1$ and (X_A, X_B) is in either $\mathcal{R}_{1A}(G)$ or $\mathcal{R}_{2A}(G)$. Observe that in this setting, the value concession $y_1^* + \epsilon$ satisfies $(X_A, X_B) \in \mathcal{R}_{1C}(G'(y_1^* + \epsilon))$. For $y < y_1^*$, (X_A, X_B) remains in either $\mathcal{R}_{1A}(G'(y))$ or $\mathcal{R}_{2A}(G'(y))$ and, thus, player A cannot derive any benefit by the previous part of the lemma. Further, player A receives strictly lower payoff for the concession of $y > y_1^* + \epsilon$ than the concession of $y_1^* + \epsilon$ because the partial derivative of her payoff with respect to y is strictly negative in $\mathcal{R}_{1B}(G'(y))$. Since $X_B \geq 1$, it must be that $X_A + X_B > 1$ and so (X_A, X_B) cannot transit to $\mathcal{R}_3$ for any y. It follows that the optimal value concession value is $y_A^{\text{opt}} = y_1^* + \epsilon$ in this setting, as in Condition (i) and (iv) of the claim.

Next, consider the scenario where (X_A, X_B) is in either $\mathcal{R}_{1A}(G)$ or $\mathcal{R}_{2A}(G)$, $X_B < 1$ and

$X_A + X_B \geq 1$ (i.e., (X_A, X_B) still cannot transit to $\mathcal{R}_3$ for any y). Here, the value concession $y_1^* + \epsilon$ satisfies $(X_A, X_B) \in \mathcal{R}_{2B}(G'(y_1^* + \epsilon))$. Once again, (X_A, X_B) remains in $\mathcal{R}_{1A}(G'(y))$ or $\mathcal{R}_{2A}(G'(y))$ after a concession of any value $y < y_1^*$, which cannot provide payoff improvements to player A. Next, recall from the proof of the previous part that the second partial derivative of player A's payoff with respect to y is strictly negative in $\mathcal{R}_{2B}(G'(y))$ (i.e., player A's payoff is concave down in y). This means that if the first partial derivative of player A's payoff with respect to y is strictly negative or zero at $y = y_1^* + \epsilon$, then the optimal value concession value is $y^{\text{opt}} = y_1^* + \epsilon$ as the derivative will remain nonpositive in $\mathcal{R}_{2B}(G'(y))$ as well as in $\mathcal{R}_{1B}(G'(y))$. Observe that this outcome corresponds with Condition (ii) of the claim. However, if the first partial derivative of player A's payoff with respect to y is strictly positive when $y = y_1^* + \epsilon$, then there are two possibilities: (1) that there is a value y for which player A's payoff is maximized (i.e., partial derivative with respect to y is zero) in the interior of region $\mathcal{R}_{2B}(G'(y))$, or (2) that the first partial derivative remains positive at the boundary between $\mathcal{R}_{1B}(G'(y))$ and $\mathcal{R}_{2B}(G'(y))$. The first partial derivative of player A's payoff with respect to y is zero after conceding the value $y \geq 0$ that satisfies $4X_A = 2 - \sqrt{\Phi_B X_A X_B/(\Phi_A - y)}$ which is $y = \Phi_A - \Phi_B X_A X_B/(2 - 4X_A)^2$, precisely the definition of y_2^*. Furthermore, (X_B, X_C) is at the boundary between $\mathcal{R}_{1B}(G'(y))$ and $\mathcal{R}_{2B}(G'(y))$ after conceding the value $y \geq 0$ that satisfies $1 - \sqrt{\Phi_B X_A X_B/(\Phi_A - y)} = 0$ which is $y = \Phi_A - \Phi_B X_A X_B$, precisely the definition of y_3^*. Thus, if the first partial derivative of player A's payoff with respect to y is strictly positive when $y = y_1^* + \epsilon$, then the optimal value concession value is $y^{\text{opt}} = y_2^*$ if $y_2^* < y_3^*$ (Possibility 1), or $y^{\text{opt}} = y_3^*$ if $y_2^* \geq y_3^*$ (Possibility 2). Observe that this is equivalent to writing $y_A^{\text{opt}} = \min\{y_2^*, y_3^*\}$ as in Condition (iii) of the claim.

Next, consider the scenario where $X_A + X_B < 1$ and (X_A, X_B) is in either $\mathcal{R}_{1A}(G)$, $\mathcal{R}_{2A}(G)$ or $\mathcal{R}_3(G)$. Observe that the point (X_A, X_B) is at the boundary between $\mathcal{R}_{2B}(G'(y))$ and $\mathcal{R}_3(G'(y))$ after conceding the value $y \geq 0$ that satisfies $1 - \sqrt{\Phi_B X_A X_B/(\Phi_A - y)} = X_A$ which is $y = \Phi_A - \Phi_B X_A X_B/(1 - X_A)^2$, precisely the definition of y_4^*. Since the point (X_A, X_B) transits through $\mathcal{R}_3(G'(y))$, the partial derivative of player A's payoff with respect to y remains strictly negative for all $y < y_4^*$. Furthermore, since the players' payoffs are continuous as we transit from $\mathcal{R}_3(G'(y))$ to $\mathcal{R}_{2C}(G'(y))$, there can be no benefit to transitting only to the boundary between $\mathcal{R}_3(G'(y))$ and $\mathcal{R}_{2C}(G'(y))$. Recall that the second partial derivative with respect to y of player A's payoff in $\mathcal{R}_{2B}(G'(y))$ is strictly negative. Thus, if the first partial derivative of player A's payoff with respect to y is nonpositive after conceding $y = y_4^* + \epsilon$, then the partial derivative will remain nonpositive in $\mathcal{R}_{2B}(G'(y))$ as well as in $\mathcal{R}_{1B}(G'(y))$. However, if the derivative is strictly positive after conceding $y = y_4^* + \epsilon$, then we consider Possibilities 1 and 2 as in the above paragraph, and obtain an optimal value concession value of $y_A^{\text{opt}} = \min\{y_2^*, y_3^*\}$ as in Condition (iv) of the claim.

Finally, consider the scenario where (X_A, X_B) is in $\mathcal{R}_{2B}(G)$. Here, if the first partial derivative of player B's payoff with respect to y is strictly negative or zero when $y = 0$, then there is no beneficial value concession as the partial derivative will remain nonpositive in $\mathcal{R}_{2B}(G'(y))$ as well as in $\mathcal{R}_{1B}(G'(y))$. However, if the first partial derivative of player A's payoff with respect to y is strictly positive when $v = 0$, then we consider Possibilities 1 and 2 once again, an optimal value concession value of $y_A^{\text{opt}} = \min\{y_2^*, y_3^*\}$ as in Condition (v) of the claim. $\square$ $\square$

Lemma 5.3.2. *The three-player General Lotto game with value concessions has at most one pure strategy Nash equilibrium that coincides with a pure strategy concession profile. If such an equilibrium exists, then it must be of the form* $(0, 0)$*,* $(y_A, 0)$ *or* $(0, y_B)$*.*

Proof. This result draws heavily from the necessary and sufficient conditions for the existence of beneficial value concessions identified in Lemma 5.3.1.

We begin by considering scenarios where the players' SPE payoffs in Stages 1 and 2 are non-unique, and show that – regardless of the SPE selection rule – no pure strategy Nash equilibrium can exist in these scenarios.

– *Non-unique SPE payoffs*: Recall from the analogous discussion on "Non-unique SPE payoffs" in Part 1 of the proof of Lemma 5.3.1 that one of the players $i \in \{A, B\}$ can strictly improve her payoff when $(X_A, X_B) \in \mathcal{R}_4$ under a given concession profile $(y_A, y_B) \in [0, \Phi_A) \times$

$[0, \Phi_B)$ by conceding $y_i' = y_i + \epsilon$, $\epsilon \to 0^+$, if either $X_{-i} < 1$, or $X_{-i} \geq 1$ and the player 0's allocation in the selected SPE for Stages 1 and 2 satisfies $X^*_{0,i} > 0$. This leaves only the setting where $X_{-i} \geq 1$ and $X^*_{0,i} = 0$ as a possible scenario for candidate equilibria when $(X_A, X_B) \in \mathcal{R}_4$. However, applying the same reasoning symmetrically, it follows that the concession $y'_{-i} = y_{-i}+\epsilon$ is a strictly better response for the player $-i$ in this case, since $X^*_{0,-i} = 1 > 0$ must hold. Finally, note that neither player $i \in \{A, B\}$ can improve her payoff by conceding $y_i = \Phi_i$, and so an appropriate $\epsilon \to 0^+$ must exist. Thus, there can be no pure strategy Nash equilibrium where $(X_A, X_B) \in \mathcal{R}_4$ under any concession profile $(y_A, y_B) \in [0, \Phi_A] \times [0, \Phi_B]$.

As we have resolved the scenario where the players' SPE payoffs in Stages 1 and 2 are non-unique, we only consider scenarios where the players' SPE payoffs in Stages 1 and 2 are unique for the remainder of the proof.

Before continuing, we make two observations that will aid in proving the claim:

O1) If (X_A, X_B) falls in any of the regions $\mathcal{R}_j$, $j \in \{1A, 2A, 3\}$, (resp., $j \in \{1B, 2B, 3\}$) under a pure strategy Nash equilibrium (y^*_A, y^*_B), then it must be that $y^*_A = 0$ (resp., $y^*_B = 0$), since we showed that $\partial/(\partial y_A)\Pi^{\text{vc}}_A(y_A, y_B)$ (resp., $\partial/(\partial y_B)\Pi^{\text{vc}}_B(y_A, y_B)$) is strictly negative in $\mathcal{R}_j$. This implies that at most one player $i \in \{A, B\}$ will make a non-zero concession in any pure strategy Nash equilibrium, i.e., any pure strategy Nash equilibrium is of the form $(0, 0)$, $(y^*_A, 0)$ or $(0, y^*_B)$ where y^*_A and y^*_B are the optimal concessions identified in Lemma 5.3.1 when (X_A, X_B) is in each region $\mathcal{R}_j$, $j \in \{1A, 1B, 2A, 2B, 3\}$.

O2) By Lemma 5.3.1, if player A (resp., B) has a beneficial unilateral deviation from the profile $(0, 0)$, then her optimal such deviation places the point (X_A, X_B) in one of the regions $\mathcal{R}_j$, $j \in \{1B, 2B\}$ (resp., $\{1A, 2A\}$). Since at most one player $i \in \{A, B\}$ can make a non-zero concession in a given pure strategy Nash equilibrium (Observation O1), if $y^*_i > 0$ in a given pure strategy Nash equilibrium, then the player $-i$'s best response must be not to concede. By contrapositive, if player $-i$ has a concession $y_{-i} > 0$ such that $\Pi^{\text{vc}}_{-i}(y^*_i, y_{-i}) > \Pi^{\text{vc}}_{-i}(y^*_i, 0)$, then there is no pure strategy Nash equilibrium associated with player i's optimal unilateral deviation.

Following Observations O1 and O2, the remainder of the proof amounts to showing that any given instance of the normal form game has at most one pure strategy Nash equilibrium. To do so we demonstrate that when the point (X_A, X_B) falls within any of the regions $\mathcal{R}_j$, $j \in \{1A, 1B, 2A, 2B, 3\}$, identified in Fact 5.3.1, it must either be that at most one of the players $i \in \{A, B\}$ has a beneficial unilateral deviation from the nominal concession profile $(y_A, y_B) = (0, 0)$, or that if both players have a beneficial unilateral deviation, then for at least one of the players $i \in \{A, B\}$, the player $-i$ has a concession $y_{-i} > 0$ such that $\Pi^{\text{vc}}_{-i}(y_i^*, y_{-i}) > \Pi^{\text{vc}}_{-i}(y_i^*, 0)$.

– *Unique SPE payoffs*: We consider the remaining regions in Parts (i) and (ii), below. Though we may focus on the proof from player A's perspective, we stress that completeness of the proof is contingent on applying the same reasoning symmetrically to player B.

Part (i). In this part, we consider the instances with $X_A + X_B \geq 1$. By Lemma 5.3.1, when $(X_A, X_B) \in \mathcal{R}_{1A}(G) \cup \mathcal{R}_{2A}(G)$ and $X_B \geq 1$, if player A's optimal unilateral deviation from the nominal profile is non-zero, then it must be $y_A^{\text{opt}} = y_1^* + \epsilon$, $\epsilon \to 0^+$, such that $(X_A, X_B) \in \mathcal{R}_{1B}$ under the concession profile $(y_A^{\text{opt}}, 0)$. Similarly, when $(X_A, X_B) \in \mathcal{R}_{1A}(G) \cup \mathcal{R}_{2A}(G)$, $X_B < 1$, $X_A + X_B \geq 1$ and $\partial/(\partial y_A)\Pi^{\text{vc}}_A(y_A, 0) \leq 0$ at $y_A = y_1^* + \epsilon$, if player A's optimal value unilateral deviation is non-zero, then it must be $y_A^{\text{opt}} = y_1^* + \epsilon$, $\epsilon \to 0^+$, such that $(X_A, X_B) \in \cup\mathcal{R}_{2B}$ under the concession profile $(y_A^{\text{opt}}, 0)$. In either case, player B can best respond to such a concession by conceding $y_B = \epsilon'$ – for sufficiently small $\epsilon' > \epsilon$ – such that $(X_A, X_B) \in \mathcal{R}_{1A} \cup \mathcal{R}_{2A}$ under the concession profile $(y_A^{\text{opt}}, \epsilon')$ and $\Pi^{\text{vc}}_B(y_A^{\text{opt}}, \epsilon') > \Pi^{\text{vc}}_B(y_A^{\text{opt}}, 0)$. Thus, by Observation O2, it follows that when $(X_A, X_B) \in \mathcal{R}_{1A}(G) \cup \mathcal{R}_{2A}(G)$ and either $X_B \geq 1$, or $X_B < 1$ and $X_A + X_B \geq 1$, if player A has a beneficial unilateral deviation, then there can be no pure strategy Nash equilibrium corresponding to A's optimal unilateral deviation from the nominal profile. Hence, only one pure strategy Nash equilibrium – corresponding with B's optimal unilateral deviation from the nominal profile – can exist in this setting.

Next, consider the region $(X_A, X_B) \in \mathcal{R}_{1A}(G) \cup \mathcal{R}_{2A}(G)$, $X_B < 1$ and $X_A + X_B \geq 1$, when $\partial/(\partial y_A)\Pi^{\text{vc}}_A(y_A, 0) > 0$ at $y_A = y_1^* + \epsilon$, $\epsilon \to 0^+$. We show that if both players A and B have a beneficial unilateral deviation from the profile $(0, 0)$, then player A's optimal value concession is $y_A^{\mathrm{opt}} = y_1^* + \epsilon$ (i.e., $\partial/(\partial y_A)\Pi_A^{\mathrm{vc}}(y_A, 0) \le 0$ at $y_A = y_1^* + \epsilon$) and, thus, results in a contradiction. First, observe that $\partial/(\partial y_A)\Pi_A^{\mathrm{vc}}(y_A, 0) > 0$ at $y_A = y_1^* + \epsilon$ if and only if $X_B < 2 - 4X_A$. Further, when $(X_A, X_B) \in \mathcal{R}_{2A}(G)$, player B has a beneficial unilateral deviation from $(0, 0)$ if and only if $\partial/(\partial y_B)\Pi_B^{\mathrm{vc}}(0, y_C) > 0$ at $y_C = 0$, from which it follows that:

$$\begin{aligned} 0 <& -1 - \frac{1}{4}\sqrt{\frac{\Phi_A X_A}{\Phi_B X_B}} + \frac{1}{2X_B} \\ &= \frac{1}{2X_B}\Big[-2X_B - \frac{1}{2}\sqrt{\frac{\Phi_A X_A X_B}{\Phi_B}} + 1\Big] \\ &\le \frac{1}{2X_B}\Big[\frac{-3}{2}X_B + \frac{1}{2}\Big], \end{aligned}$$

where the first inequality follows from $\partial/(\partial y_B)\Pi_B^{\mathrm{vc}}(0, y_C) > 0$ at $y_C = 0$, and the last inequality holds by the necessary conditions for $(X_A, X_B) \in \mathcal{R}_{2A}(G)$. This means that when player B has a beneficial unilateral deviation from $(0, 0)$ and $(X_A, X_B) \in \mathcal{R}_{2A}(G)$ it must hold that $X_B < 1/3$. Note that the set of $(X_A, X_B) \in \mathbb{R}^2_{\ge 0}$ with $X_A + X_B \ge 1$ and $X_B < 2 - 4X_A$ is strictly disjoint from the set of $(X_A, X_B) \in \mathbb{R}^2_{\ge 0}$ with $X_A + X_B \ge 1$ and $X_B < 1/3$. Thus, if $(X_A, X_B) \in \mathcal{R}_{2A}(G)$, $X_A + X_B \ge 1$ and both players A and B have beneficial unilateral deviation from $(0, 0)$, then $\partial/(\partial y_A)\Pi_A^{\mathrm{vc}}(y_A, 0) \le 0$ which is a contradiction. Thus, B cannot have a beneficial unilateral deviation from the nominal concession profile when $(X_A, X_B) \in \mathcal{R}_{1A}(G) \cup \mathcal{R}_{2A}(G)$, $X_B < 1$, $X_A + X_B \ge 1$, and $\partial/(\partial y_A)\Pi_A^{\mathrm{vc}}(y_A, 0) > 0$ at $y_A = y_1^* + \epsilon$, $\epsilon \to 0^+$.

Part (ii). In the final part of the proof, we consider $(X_A, X_B) \in \mathcal{R}_{2A}(G) \cup \mathcal{R}_{2B}(G) \cup \mathcal{R}_3(G)$ with $0 \le X_A + X_B < 1$. For any pair $X_A, X_B \ge 0$ in this regime, we show that at most one player $i \in \{A, B\}$ has a beneficial unilateral deviation from $(0, 0)$. For $(X_A, X_B) \in \mathcal{R}_{2A}(G) \cup \mathcal{R}_3(G)$ with $0 \le X_A + X_B < 1$, a necessary condition for player A to have a beneficial unilateral deviation from $(0,0)$ is $\partial/(\partial y_A)\Pi_A^{\mathrm{vc}}(y_A,0) > 0$ at $y_A = y_4^* + \epsilon$, from which it follows that:

$$\begin{aligned}
0 &< \frac{1}{2X_A} - \frac{\sqrt{(1-X_A)^2}}{2X_A} + \frac{\sqrt{(1-X_A)^2}}{4X_A} - 1 \\
&= \frac{1}{2X_A} - \frac{1-X_A}{2X_A} + \frac{1-X_A}{4X_A} - 1 \\
&= \frac{1}{4X_A} - \frac{3}{4},
\end{aligned}$$

where the first inequality follows from $\partial/(\partial y_A)\Pi_A^{\mathrm{vc}}(y_A,0) > 0$ at $y_A = y_4^* + \epsilon$, and the equality in the second line follows from $0 \leq X_A + X_B < 1$ and $X_A, X_B > 0$. Thus, when player A has a beneficial value concession and $(X_A, X_B) \in \mathcal{R}_{2A}(G) \cup \mathcal{R}_3(G)$ with $0 \leq X_A + X_B < 1$, it must hold that $X_A < 1/3$. We showed in Part (i) that $X_A < 1/3$ must also hold whenever player A has a beneficial unilateral deviation from $(0,0)$ and $(X_A, X_B) \in \mathcal{R}_{2B}(G)$.[7] Thus, when $0 \leq X_A + X_B < 1$, we have shown that if both players A and B have beneficial value concessions simultaneously, then $X_A, X_B \in [0, 1/3]$ must hold.

Player A's optimal value concession when $0 \leq X_A + X_B < 1$ is $y_A^{\mathrm{opt}} = \min\{y_2^*, y_3^*\}$. It immediately follows from the definitions of y_2^* and y_3^* that $y_A^{\mathrm{opt}} = y_2^*$ for $X_A \in [1/4, 3/4]$ and $y_A^{\mathrm{opt}} = y_3^*$, otherwise. We divide $(X_A, X_B) \in [0, 1/3] \times [0, 1/3]$ into three subregions as follows:

–$X_A, X_B \in [1/4, 1/3]$: For each $i \in \{A, B\}$, it is straightforward to verify that after her optimal value concession $v_i^{\mathrm{opt}} = \Phi_i - \Phi_{-i} X_A X_B/(2-4X_i)^2$, it holds that $(X_A, X_B) \in \mathcal{R}_{2(-i)}$ and player i receives payoff $\Phi_{-i} X_{-i}/(4 - 8X_i)$. If $(X_A, X_B) \in \mathcal{R}_3(G)$, then player $i \in \{A, B\}$ has nominal payoff:

$$\Pi_i = \Phi_i X_i + \sqrt{\Phi_A \Phi_B X_A X_B}.$$

If both players' optimal value concessions are beneficial unilateral deviations from $(0,0)$ in this case then it must hold that:

$$2\sqrt{\Phi_A \Phi_B X_A X_B} < (\Phi_A X_A + \Phi_B X_B)\Big[\frac{1}{4 - 8X_B} - 1\Big].$$

[7] We actually showed that $X_B < 1/3$ whenever player B has a beneficial value concession and $(X_A, X_B) \in \mathcal{R}_{2A}(G)$, but this statement follows by symmetry.

However, for $x \in [1/4, 1/3]$, it holds that $4 - 8x > 4 - 8/3 > 1$ and, thus, the right-hand side is strictly less than zero, which is a contradiction. This implies that both players A and B cannot simultaneously have beneficial value concessions in this case. If $(X_A, X_B) \in \mathcal{R}_{2B}(G)$, player A's nominal payoff is

$$\Pi_A = \Phi_A - \frac{\Phi_A}{2X_A} + \frac{1}{2}\sqrt{\frac{\Phi_A\Phi_B X_B}{X_A}},$$

and player B's nominal payoff is

$$\Pi_B = \frac{1}{2}\sqrt{\frac{\Phi_A\Phi_B X_B}{X_A}}.$$

If $\Phi_A \leq \Phi_B$, then player B's optimal value concession is a beneficial unilateral deviation only if

$$\begin{aligned}
&\sqrt{\frac{\Phi_A\Phi_B X_B}{X_A}} < \frac{\Phi_A X_A}{2 - 4X_B} \\
\implies &\Phi_A\sqrt{\frac{X_B}{X_A}} < \Phi_A\frac{X_A}{2 - 4X_B} \\
\implies &\frac{3}{4} < \frac{1}{2},
\end{aligned}$$

where the second line holds because $\Phi_A \leq \Phi_B$, and the third line holds because $X_A, X_B \in [1/4, 3/4]$. This is a contradiction, which means player B's optimal value concession is not a beneficial unilateral deviation from $(0, 0)$ in this case. If $\Phi_A > \Phi_B$, then player A's optimal value concession is a beneficial unilateral deviation only if

$$\begin{aligned}
&\Phi_A - \frac{\Phi_A}{2X_A} + \frac{1}{2}\sqrt{\frac{\Phi_A\Phi_B X_B}{X_A}} < \frac{\Phi_B X_B}{4 - 8X_A} \\
\implies &\frac{\Phi_A}{2} < \frac{\Phi_A X_B}{4 - 8X_A} \\
\implies &\frac{1}{2} < \frac{1}{2},
\end{aligned}$$

where the second line holds because $0 < 1 - \sqrt{\Phi_B X_A X_B/\Phi_A} \leq X_A$ holds when $(X_A, X_B) \in \mathcal{R}_{2B}(G)$ and $\Phi_A > \Phi_B$, and the third line holds because $X_A, X_B \in [1/4, 3/4]$. This is a contradiction, which means player A cannot have a beneficial value concession in this case.

This reasoning holds symmetrically for $(X_A, X_B) \in \mathcal{R}_{2A}(G)$.

–$X_A, X_B \in [0, 1/4]$: For each $i \in \{A, B\}$, it is straightforward to verify that after her optimal value concession $y_i^{\text{opt}} = \Phi_i - \Phi_{-i} X_A X_B$, it holds that $(X_A, X_B) \in \mathcal{R}_{2(-i)}(G'(y_i^{\text{opt}}))$ and player i receives payoff $\Phi_{-i} X_A X_B$. If $(X_A, X_B) \in \mathcal{R}_3(G)$ and both players A and B have beneficial unilateral deviations from $(0, 0)$, it must hold that

$$2\sqrt{\Phi_A \Phi_B X_A X_B} < \Phi_A X_A (X_B - 1) + \Phi_B X_B (X_A - 1).$$

However, the right-hand side must be negative since $0 < X_A, X_B < 1$ which means there is a contradiction and both players cannot have unilateral deviations simultaneously. If $(X_A, X_B) \in \mathcal{R}_{2B}(G)$, then player A has a beneficial unilateral deviation only if

$$\begin{aligned} \Phi_A - \frac{\Phi_A}{2X_A} + \frac{1}{2}\sqrt{\frac{\Phi_A \Phi_B X_B}{X_A}} &< \Phi_B X_A X_B \\ \implies \frac{\Phi_A}{2} &< \Phi_A X_A X_B, \end{aligned}$$

where the second line holds because

$$0 \leq 1 - \sqrt{\Phi_B X_A X_B / \Phi_A} < X_A$$

in $\mathcal{R}_{2B}(G)$. This is a contradiction since $X_A X_B \leq 1/16 < 1/2$ for $X_A, X_B \in [0, 1/4]$ which means A has no beneficial unilateral deviations in this case. This reasoning holds symmetrically for $(X_A, X_B) \in \mathcal{R}_{2A}(G)$.

–$X_A \in [1/4, 1/3], X_B \in [0, 1/4]$ OR $X_A \in [0, 1/4]$, $X_B \in [1/4, 1/3]$: We only consider the former, as the proof follows symmetrically for the latter. Observe that $y_A^{\text{opt}} = \Phi_A - \Phi_B X_A X_B/(2-4X_A)^2$ while $y_B^{\text{opt}} = \Phi_B - \Phi_A X_A X_B$, and that players A and B's payoffs after their optimal value concessions are $\Phi_B X_B/(4 - 8X_A)$ and $\Phi_A X_A X_B$, respectively. If $(X_A, X_B) \in \mathcal{R}_3(G)$ and both players A and B have beneficial unilateral deviations from $(0,0)$, it must hold that

$$2\sqrt{\Phi_A\Phi_B X_A X_B} < \Phi_A X_A(X_B-1) + \frac{\Phi_B X_B}{4-8X_A} - \Phi_B X_B.$$

As in the above, this leads to a contradiction since $X_B < 1$ and $4-8X_A \geq 4-8/3 > 1$ imply that the right-hand side is negative. Thus, both players cannot simultaneously have beneficial unilateral deviations in this case. If $(X_A, X_B) \in \mathcal{R}_{2A}(G)$, then B has a beneficial unilateral deviation from $(0,0)$ only if

$$\Phi_B - \frac{\Phi_B}{2X_B} + \frac{1}{2}\sqrt{\frac{\Phi_B\Phi_A X_A}{X_B}} < \Phi_A X_A X_B,$$

which we know leads to a contradiction from the above. If $(X_A, X_B) \in \mathcal{R}_{2B}(G)$ and $\Phi_A \leq \Phi_B$, then B has a beneficial value concession only if

$$\begin{aligned} &\frac{1}{2}\sqrt{\frac{\Phi_A\Phi_B X_B}{X_A}} < \Phi_A X_A X_B \\ \implies &\sqrt{X_B} < X_A X_B, \end{aligned}$$

where the second line holds because $\Phi_A \leq \Phi_B$ and $X_A \in [0, 1/4]$. This is a contradiction as $0 < X_A, X_B < 1/3$. If $(X_A, X_B) \in \mathcal{R}_{2B}(G)$ and $\Phi_A > \Phi_B$, then A has a benificial unilateral deviation from $(0,0)$ only if

$$\begin{aligned} &\Phi_A - \frac{\Phi_A}{2X_A} + \frac{1}{2}\sqrt{\frac{\Phi_A\Phi_B X_B}{X_A}} < \frac{\Phi_B X_B}{4-8X_A} \\ \implies &\frac{\Phi_A}{2} < \frac{\Phi_A X_B}{4-8X_A}, \end{aligned}$$

where the second line follows because $0 \leq 1 - \sqrt{\Phi_B X_A X_B/\Phi_A} < X_A$ holds when $(X_A, X_B) \in \mathcal{R}_{2B}(G)$ and because $\Phi_A > \Phi_B$. This leads to a contradiction since $X_B/(2-4X_A) < 3/8 < 1$.

Observing that we have exhausted all possible points $(X_A, X_B) \in \mathbb{R}^2_{\geq 0}$ concludes the proof.

□ □

5.3.2 Proof of Theorem 5.2.1

Method to derive equilibria of second-stage subgame The recent work of Vu and Loiseau [51] provides a general method to derive an equilibrium of the second stage subgame from the GL-P game, which is termed a General Lotto game with favoritism (GL-F). In a GL-F game, the pre-allocation vector $\mathbf{p}$ is an exogenous parameter. We denote an instance of as GL-F$(\mathbf{p}, R_A, R_B)$. The method to calculate an equilibrium involves solving the following system[8] of two equations for two unknowns $(\kappa_A, \kappa_B) \in \mathbb{R}^2_{++}$:

$$R_A = \sum_{b=1}^{n} \frac{[h_b(\kappa_A, \kappa_B) - p_b]^2}{2qw_b\kappa_B}, \ R_B = \sum_{b=1}^{n} \frac{h_b^2(\kappa_A, \kappa_B) - p_b^2}{2qw_b\kappa_A} \tag{5.13}$$

where $h_b(\kappa_A, \kappa_B) := \min\{qw_b\kappa_B, w_b\kappa_A + p_b\}$ for $b \in \mathcal{B}$. The above equations correspond to the budget constraint (5.1) for both players. There always exists a solution $(\kappa_A^*, \kappa_B^*) \in \mathbb{R}^2_{++}$ to this system [51], and corresponds to the following equilibrium payoffs.

Lemma 5.3.3 (Adapted from [51])**.** *Suppose $(\kappa_A^*, \kappa_B^*) \in \mathbb{R}^2_{++}$ solves* (5.13)*. Let $\mathcal{B}_1 := \{b \in \mathcal{B} : h_b(\kappa_A^*, \kappa_B^*) = qw_b\kappa_B\}$ and $\mathcal{B}_2 = \mathcal{B} \backslash \mathcal{B}_1$. Then there is a corresponding equilibrium (F_A^*, F_B^*) of the game GL-F$(\mathbf{p}, R_A, R_B)$ where player A's equilibrium payoff is given by*

$$\begin{aligned} \pi_A(\mathbf{p}, R_A, R_B) := &\sum_{b \in \mathcal{B}_1} w_b \left[1 - \frac{q\kappa_B^*}{2\kappa_A^*}\left(1 - \frac{p_i^2}{(qw_b\kappa_B)^2}\right)\right] \\ &+ \sum_{b \in \mathcal{B}_2} w_b \frac{\kappa_A^*}{2q\kappa_B^*} \end{aligned} \tag{5.14}$$

and the equilibrium payoff to player B is $\pi_B(\mathbf{p}, R_A, R_B) = W - \pi_A(\mathbf{p}, R_A, R_B)$.

The equilibrium strategies are characterized by marginal distributions detailed in [51].

Proof of Theorem 5.2.1 The proof follows two parts: In Part 1, we establish that, for given $P, R_A, R_B > 0$ and $\mathbf{w} \in \mathbb{R}^n_{++}$, $\mathbf{p}^* = \frac{P}{W}\mathbf{w}$ is player A's optimal pre-allocation profile in Stage

[8] The problem setting considered in their method is more general, admitting possibly negative pre-allocations $p_b < 0$ (i.e. favoring player B), asymmetries in players' battlefield valuations $w_b > 0$, and different resource effectiveness parameters q_b for each battlefield. However, exact closed-form solutions under heterogeneous values $\mathbf{w}$, arbitrary pre-allocations $\mathbf{p}$, and effectiveness parameters q_b are generally unattainable.

1 of GL-P. Then, in Part 2, we derive an explicit expression for player A's payoff in Stage 2 under the optimal pre-allocation profile $\mathbf{p}^*$ derived in Part 1. Throughout the proof, we use $\pi_i(\mathbf{p}, R_A, R_B)$, $i \in \{A, B\}$, to denote the players' payoffs in the Stage 2 sub-game for fixed pre-allocation profile $\mathbf{p} \in \Delta_n(P)$. Recall that the Stage 2 sub-game amounts to a General Lotto game with favoritism GL-F$(\mathbf{p}, R_A, R_B)$.

– **Part 1:** The proof amounts to showing that $\mathbf{p}^* = \frac{P}{W}\mathbf{w}$ is a global maximizer of player A's equilibrium payoff $\pi_A^*(\mathbf{p}, R_A, R_B)$ for $\mathbf{p} \in \Delta_n(P)$. For the following analysis, we define $\mathbb{T}_n := \{\mathbf{z} \in \mathbb{R}^n : \sum_{b=1}^n z_b = 0$ as the tangent space of $\Delta_n(P)$. The lemma below first establishes that $\mathbf{p}^*$ is a local maximizer when either $\mathcal{B}_1 = \mathcal{B}$ or $\mathcal{B}_2 = \mathcal{B}$.

Lemma 5.3.4. *The pre-allocation $\mathbf{p}^* = \frac{P}{W}\mathbf{w}$ is a local maximizer of $\pi_A(\mathbf{p}, R_A, R_B)$ over $\mathbf{p} \in \Delta_n(P)$, for any $P, R_A, R_B > 0$.*

Proof. From Lemma 5.3.3 and the definition of $h_b(\kappa_A, \kappa_B)$ in (5.13), we observe that the solution to (5.13) under the pre-allocation s^* is always in one of two completely symmetric cases: 1) $\mathcal{B}_1 = \mathcal{B}$; or 2) $\mathcal{B}_2 = \mathcal{B}$. Thus, we need to show s^* is a local maximizer in both cases.

<u>**Case 1 ($\mathcal{B}_1 = \mathcal{B}$):**</u> For $\mathbf{p} \in \Delta_n(P)$, the system (5.13) is written

$$R_A = \sum_{b=1}^{n} \frac{(qw_b\kappa_B - p_b)^2}{2qw_b\kappa_B} \text{ and } R_B = \sum_{b=1}^{n} \frac{(qw_b\kappa_B)^2 - p_b^2}{2qw_b\kappa_A} \tag{5.15}$$
$$\text{where } 0 < qw_b\kappa_B - p_b \leq \kappa_A \text{ holds } \forall b \in \mathcal{B}.$$

It yields an algebraic solution

$$\begin{aligned} q\kappa_B^* &= \frac{1}{W}\left[P + R_A + \sqrt{(P+R_A)^2 - W\|\mathbf{p}\|_{\mathbf{w}}^2}\right] \\ \kappa_A^* &= \frac{(P+R_A)q\kappa_B^* - \|\mathbf{p}\|_{\mathbf{w}}^2}{qR_B}. \end{aligned} \tag{5.16}$$

where $\|\mathbf{p}\|_{\mathbf{w}}^2 = \sum_{b=1}^n \frac{p_b^2}{w_b}$. This solution needs to satisfy the set of conditions $0 < qw_b\kappa_B - p_b \leq \kappa_A$ $\forall b \in \mathcal{B}$, but the explicit characterization of these conditions is not needed to show that s^* is a local maximum. Indeed, first observe that the expression for $q\kappa_B^*$ is required to be real-valued, which we can write as the condition

$$\mathbf{p} \in R^{(1n)} := \left\{ \mathbf{p} \in \Delta_n(P) : \|\mathbf{p}\|_{\mathbf{w}}^2 < \frac{(P + R_A)^2}{W} \right\}. \tag{5.17}$$

We thus have a region $R^{(1n)}$ for which the expression of player A's equilibrium payoff (derived using Lemma 5.3.3) is well-defined:

$$\pi_A^{(1n)}(\mathbf{p}) := W \left(1 - \frac{qR_B}{f(\|\mathbf{p}\|_{\mathbf{w}})} \left(1 - \frac{W\|\mathbf{p}\|_{\mathbf{w}}^2}{(P+R_A+f(\|\mathbf{p}\|_{\mathbf{w}}))^2} \right) \right) \tag{5.18}$$

where $f(\|\mathbf{p}\|_{\mathbf{w}}) := \sqrt{(P + R_A)^2 - W\|\mathbf{p}\|_{\mathbf{w}}^2}$. The partial derivatives are calculated to be

$$\frac{\partial \pi_A^{(1n)}}{\partial p_b}(\mathbf{p}) = \frac{p_b}{w_b} \cdot \frac{2W^2 qR_B}{f(\|\mathbf{p}\|_{\mathbf{w}})(P + R_A + f(\|\mathbf{p}\|_{\mathbf{w}}))^2} \tag{5.19}$$

A critical point of $\pi_A^{(1n)}$ must satisfy $\mathbf{z}^\top \nabla \pi_A^{(1n)}(\mathbf{p}) = 0$ for any $\mathbf{z} \in \mathbb{T}_n$. Indeed for any $\mathbf{p} \in R^{(1n)}$, we calculate

$$\begin{aligned} (\mathbf{p} - \frac{P}{W}\mathbf{w})^\top \nabla \pi_A^{(1n)}(\mathbf{p}) &= g(\|\mathbf{p}\|_{\mathbf{w}}) \cdot \left(\|\mathbf{p}\|_{\mathbf{w}}^2 - \frac{P^2}{W} \right) \\ &\geq 0 \end{aligned} \tag{5.20}$$

where $g(\|\mathbf{p}\|_{\mathbf{w}}) := \frac{2W^2 qR_B}{f(\|\mathbf{p}\|_{\mathbf{w}})(P+R_A+f(\|\mathbf{p}\|_{\mathbf{w}}))^2} > 0$ for any $\mathbf{p} \in R^{(1n)}$. The inequality above is met with equality if and only if $\mathbf{p} = \mathbf{p}^*$. This is due to the fact that $\min_{\mathbf{p} \in \Delta_n(P)} \|\mathbf{p}\|_{\mathbf{w}}^2 = \|\mathbf{p}^*\|_{\mathbf{w}}^2 = \frac{P^2}{W}$. Thus, $\mathbf{p}^*$ is the unique maximizer of $\pi_A^{(1n)}(\mathbf{p})$ on $R^{(1n)}$.

<u>**Case 2 ($\mathcal{B}_2 = \mathcal{B}$):**</u> For $\mathbf{p} \in \Delta_n(P)$, the system is written as

$$R_A = \sum_{b=1}^{n} \frac{(w_b \kappa_A)^2}{2qw_b \kappa_B} \text{ and } R_B = \sum_{b=1}^{n} \frac{(w_b \kappa_A - p_b)^2 - (p_b)^2}{2qw_b \kappa_A},$$

where $qw_b\kappa_B - p_b > w_b\kappa_A$ holds for all $b \in \mathcal{B}$. This readily yields the algebraic solution:

$$q\kappa_B^* = \frac{2}{W} \frac{(qR_B - P)^2}{R_A} \text{ and } \kappa_A^* = \frac{2}{W}(qR_B - P). \tag{5.21}$$

For this solution to be valid, the following conditions are required:

- $\kappa_A^*, q\kappa_B^* \in \mathbb{R}_{++}$: This requires that $qR_B - P > 0$.
- $qw_b\kappa_B^* - p_b > w_b\kappa_A^*$ for all $b \in \mathcal{B}$: This requires that

$$\frac{2}{W}\frac{(qR_B - P)^2}{R_A} - \frac{2}{W}(qR_B - P) - \max_b\{\frac{p_b}{w_b}\} > 0.$$

The left-hand side is quadratic in $qR_B - P$, and, thus, requires either

$$qR_B - P < \frac{2/W - \sqrt{\frac{4}{W^2} + 4\max_b\{\frac{p_b}{w_b}\}\frac{2}{WR_A}}}{4/(WR_A)}, \text{ or}$$
$$qR_B - P > \frac{2/W + \sqrt{\frac{4}{W^2} + 4\max_b\{\frac{p_b}{w_b}\}\frac{2}{WR_A}}}{4/(WR_A)}.$$

The former cannot hold since the numerator on the right-hand side is strictly negative, but $\kappa_A^*, q\kappa_B^* \in \mathbb{R}_{++}$ requires $qR_B - P > 0$. Thus, the latter must hold, which simplifies to the condition

$$qR_B - P > \frac{R_A}{2}\left[1 + \sqrt{1 + \frac{2W}{R_A}\max_b\{\frac{p_b}{w_b}\}}\right]. \tag{5.22}$$

Clearly, (5.22) is more restrictive than $qR_B - P > 0$, and, thus, dictates the boundary of Case 2.

For any $\mathbf{p} \in \Delta_n(P)$ such that all battlefields are in Case 2, the expression for player A's payoff in (5.14) simplifies to

$$\pi_A(\mathbf{p}, R_A, R_B) = \sum_{b=1}^{n} w_b \frac{\kappa_A^*}{2q\kappa_B^*} = \frac{W}{2}\frac{R_A}{qR_B - P},$$

where we use the expression for $q\kappa_B^*$ and κ_A^* in (5.21). Observe that player A's payoff is constant in the quantity $\mathbf{p}$. Thus, for any $\mathbf{p}$ that satisfies (5.22), it holds that all battlefields are in Case 2, and that player A's payoff is the above. We conclude the proof noting that, for given quantities R_A and P, if there exists any $\mathbf{p} \in \Delta_n(P)$ such that (5.22) is satisfied, then $\mathbf{p}^* = \mathbf{w} \cdot (P/W)$ must also satisfy (5.22), since $||\mathbf{p}||_\infty \geq ||\mathbf{p}^*||_\infty$ and the right-hand side in (5.22) is increasing in $||\mathbf{p}||_\infty$. □

Next, we prove that the function $\pi_A(\mathbf{p}, R_A, R_B)$ is maximized by $\mathbf{p}^* = \frac{P}{W}\mathbf{w}$. We showed in Lemma 5.3.4 that $\mathbf{p}^*$ is a local maximizer over $\mathbf{p} \in \Delta_n(P)$ when either $\mathcal{B}_1 = \mathcal{B}$ or $\mathcal{B}_2 = \mathcal{B}$. It remains to be shown that player A cannot achieve a higher payoff for $\mathbf{p} \in \Delta_n(P)$ that results in both sets $\mathcal{B}_1$ and $\mathcal{B}_2$ being nonempty. Throughout the proof, we will use the short-hand notation $W_j = \sum_{b\in\mathcal{B}_j} w_b$, $P_j = \sum_{b\in\mathcal{B}_j} p_b$ and $\mathbf{p}_j = (p_b)_{b\in\mathcal{B}_j}$, $j = 1, 2$, for conciseness.

For $\mathbf{p} \in \Delta_n(P)$, we have that

$$\begin{aligned} X_A &= \sum_{b\in\mathcal{B}_1} \frac{(qw_b\kappa_B - p_b)^2}{2qw_b\kappa_B} + \sum_{b\in\mathcal{B}_2} \frac{(w_b\kappa_A)^2}{2qw_b\kappa_B}, \\ X_B &= \sum_{b\in\mathcal{B}_1} \frac{(qw_b\kappa_B)^2 - (p_b)^2}{2qw_b\kappa_A} + \sum_{b\in\mathcal{B}_2} \frac{(w_b\kappa_A + p_b)^2 - (p_b)^2}{2qw_b\kappa_A}, \end{aligned}$$

where $0 < qw_b\kappa_B - p_b \leq w_b\kappa_A$ holds for all $b \in \mathcal{B}_1$, and $qw_b\kappa_B - p_b > w_b\kappa_A$ holds for all $b \in \mathcal{B}_2$. The system of equations readily gives the expression:

$$\begin{aligned} W_1(q\kappa_B)^2 + W_2(\kappa_A)^2 &= 2q\kappa_B(X_A + P_2) - ||\mathbf{p}_1||_w^2 \\ &= 2\kappa_A(qX_B - P_2) + ||\mathbf{p}_1||_w^2, \end{aligned} \tag{5.23}$$

where recall that $||\mathbf{p}_1||_w^2 = \sum_{b\in\mathcal{B}_2}[(p_b)^2/w_b]$. The solution to the above system of equations is

$$\begin{aligned} q\kappa_B^* &= \frac{C_1H_2 \pm \sqrt{(C_2)^2H_1H_2}}{W_1(C_2)^2 + W_2(C_1)^2}, \\ \kappa_A^* &= \frac{C_2H_1 \pm \sqrt{(C_1)^2H_1H_2}}{W_1(C_2)^2 + W_2(C_1)^2}, \end{aligned} \tag{5.24}$$

where we introduce the short-hand notation $C_1 = R_A + P_1$, $C_2 = qR_B - P_2$, $H_1 = (C_1)^2 - W_1||\mathbf{p}_1||_w^2$ and $H_2 = (C_2)^2 + W_2||\mathbf{p}_1||_w^2$, for conciseness. We consider only the scenario where $\pm = +$ in (5.24), since the expression for κ_A^* is strictly negative when $\pm = -$. Simply observe that $C_1 > 0$, $(C_1)^2 > H_1$, $0 < (C_2)^2 < H_2$ and, thus, that either (i) $H_1 > 0$, $C_2 > 0$ and $0 < C_2H_1 < C_1\sqrt{H_1H_2}$, (ii) $H_1 < 0$, $C_2 < 0$ and $0 < C_2H_1 = |C_2||H_1| < C_1\sqrt{|H_1||H_2|}$, or (iii) only one of H_1 or C_2 is negative, in which case $C_2H_1 < 0$.

Substituting (5.24) into (5.14) and simplifying, we obtain

$$\pi_A(\mathbf{p}, R_A, R_B) = W_1 + \frac{\sqrt{H_1 H_2} - C_1 C_2}{||\mathbf{p}_1||_w^2}, \tag{5.25}$$

and the partial derivatives of $\pi_A(\mathbf{p}, R_A, R_B)$ with respect to p_b for $b \in \mathcal{B}_1$ and $b \in \mathcal{B}_2$, respectively, are:

$$\begin{aligned} \frac{\partial}{\partial p_b}\pi_A\bigg|_{b\in\mathcal{B}_1} &= \frac{-p_b/w_b}{(||\mathbf{p}_1||_w^2)^2\sqrt{H_1 H_2}}(C_1\sqrt{H_2} - C_2\sqrt{H_1})^2 \\ &\quad + \frac{1}{||\mathbf{p}_1||_w^2\sqrt{H_1}}(C_1\sqrt{H_2} - C_2\sqrt{H_1}) \\ \frac{\partial}{\partial p_b}\pi_A\bigg|_{b\in\mathcal{B}_2} &= \frac{1}{||\mathbf{p}_1||_w^2\sqrt{H_2}}(C_1\sqrt{H_2} - C_2\sqrt{H_1}). \end{aligned} \tag{5.26}$$

We first consider critical points $\mathbf{p}$ strictly in the interior of $\Delta_n(P)$, and resolve the points on the boundary later. One necessary condition for a critical point is that $\partial\pi_A/(\partial p_b) - \partial\pi_A/(\partial p_c) = 0$ for all $b \in \mathcal{B}_1$ and $c \in \mathcal{B}_2$. Firstly, observe that $C_1 > \sqrt{H_1}$ and $\sqrt{H_2} > C_2$, and, thus, it must be that $C_1\sqrt{H_2} - C_2\sqrt{H_1} > 0$. We can thus divide the expression $\partial\pi_A/(\partial p_b) = \partial\pi_A/(\partial p_c)$ on both sides by $C_1\sqrt{H_2} - C_2\sqrt{H_1}$ and rearrange to obtain

$$(p_b/w_b)(C_1\sqrt{H_2} - C_2\sqrt{H_1}) = ||\mathbf{p}_1||_w^2(\sqrt{H_2} - \sqrt{H_1}) > 0.$$

Observe that the left-hand side is strictly greater than zero, and, thus, the right-hand side must be as well. This immediately requires $\sqrt{H_2} - \sqrt{H_1} > 0$, since $||\mathbf{p}_1||_w^2 > 0$. Re-arranging the above expression, note that we also require

$$\sqrt{H_1}[C_2(p_b/w_b) - ||\mathbf{p}_1||_w^2] = \sqrt{H_2}[C_1(p_b/w_b) - ||\mathbf{p}_1||_w^2].$$

Since we have just shown that $\sqrt{H_2} > \sqrt{H_1}$ must hold, it follows that each $b \in \mathcal{B}_1$ satisfies either (i) $C_2(p_b/w_b) - ||\mathbf{p}_1||_w^2 < C_1(p_b/w_b) - ||\mathbf{p}_1||_w^2 < 0$; or (ii) $C_2(p_b/w_b) - ||\mathbf{p}_1||_w^2 > C_1(p_b/w_b) - ||\mathbf{p}_1||_w^2 > 0$. Observe that $C_1(p_b/w_b) > ||\mathbf{p}_1||_1^2$ must hold for $b' \in \arg\max_{b\in\mathcal{B}_1} p_b/w_b$, and thus b' must satisfy scenario (ii) and $C_2 > C_1$ (or, equivalently, $qR_B - P > R_A$). This last inequality then implies that scenario (ii) must be satisfied for all $b \in \mathcal{B}_1$.

We have shown that, in order for $\partial\pi_A/(\partial p_b) - \partial\pi_A/(\partial p_c) = 0$ to hold for all $b \in \mathcal{B}_1$ and $c \in \mathcal{B}_2$, a critical point $\mathbf{p}$ must satisfy

$$\frac{p_b}{w_b} = \bar{p} := \frac{\sqrt{H_2} - \sqrt{H_1}}{C_1\sqrt{H_2} - C_2\sqrt{H_1}}||\mathbf{p}_1||_w^2,$$

for each $b \in \mathcal{B}_1$. Expanding this expression, and solving for $\bar{p}$ explicitly, we obtain the following two possible (real) solutions for $\bar{p}$:

$$\bar{p} = 0 \text{ or } \bar{p} = \frac{2(qR_B - P)(qR_B - R_A - P)}{WR_A},$$

where we use $P_1 = W_1\bar{p}$, $P_2 = P - P_1$, and $||\mathbf{p}_1||_w^2 = W_1(\bar{p})^2$. As $\bar{p} = 0$ is inadmissible, we consider the latter expression for $\bar{p}$. After inserting this expression for $\bar{p}$ into the right-hand side of (5.22), where $\max_b\{p_b/w_b\} = \bar{p}$, we obtain

$$\begin{aligned}&\frac{R_A}{2}\left[1 + \sqrt{1 + \frac{2W}{R_A}\bar{p}}\right] \\ &= \frac{R_A}{2} + qR_B - P - \frac{R_A}{2} = qR_B - P,\end{aligned}$$

which follows since we showed above that $qR_B - P > R_A$ must hold. Thus, the only critical point sits at the boundary of the region where all battlefields are in Case 2, since decreasing $\bar{p}$ even slightly will satisfy the condition in (5.22). We can further verify that the payoff at this critical point is equal to the constant payoff in the region where all battlefields are in Case 2, but omit this for conciseness.

We conclude the proof by resolving the scenario where $\mathbf{p}$ lies on the boundaries of $\Delta_n(P)$. Observe that the conditions on $q\kappa_B^*$ and κ_A^* immediately imply that $p_b/w_b > p_c/w_c$ for any $b \in \mathcal{B}_1$ and $c \in \mathcal{B}_2$. Thus, on the boundaries of $\Delta_n(P)$, it must either be that all battlefields with $p_b = 0$ (and possibly more) are in Case 2, or that all battlefields in $\mathcal{B}$ are in Case 1 (which is covered by Lemma 5.3.4). In the scenario where all battlefields with $p_b = 0$ are in Case 2, note that the necessary condition $(\partial/(\partial p_i) - \partial/(\partial p_j))\pi_A \geq 0$ for $i \in \arg\min_{b\in\mathcal{B}_1}\{p_b/w_b\}$ and $j \in \arg\max_{b\in\mathcal{B}_1}\{p_b/w_b\}$ only holds with equality if $p_b/w_b = P_1/W_1$ for all $b \in \mathcal{B}_1$. If $P_1/W_1 < \bar{p}$, then the inequality in (5.22) is satisfied implying that all battlefields are in Case 2, and Lemma 5.3.4 shows that $\mathbf{p}^* = \mathbf{w}(P/W)$ must correspond with the same payoff to player A. Otherwise, if $P_1/W_1 = \bar{p}$, then we showed above that the global maximum sits at the boundary where all battlefields are in Case 2 and $\mathbf{p}^* = \mathbf{w}(P/W)$ achieves the same payoff. Finally, if $P_1/W_1 > \bar{p}$, then, from (5.26), we know that $\partial\pi_A/(\partial p_b) - \partial\pi_A/(\partial p_c) < 0$ must hold for all $b \in \mathcal{B}_1$ and $c \in \mathcal{B}_2$, since the choice $p_b/w_b = \bar{p}$ satisfies $\partial\pi_A/(\partial p_b) - \partial\pi_A/(\partial p_c) = 0$, and $\partial\pi_A/(\partial p_b)$ is decreasing with respect to p_b/w_b while $\partial\pi_A/(\partial p_c)$ is constant. This violates a necessary condition for a critical point, and implies that A's payoff is increasing in the direction of decreasing p_b and increasing p_c, as expected. □

– **Part 2:** In the proof of Lemma 5.3.4, we provide the closed-form solutions to the system of equations (5.13) for the symmetric case $\mathcal{B}_1 = \mathcal{B}$ (resp. $\mathcal{B}_2 = \mathcal{B}$) in (5.16) (resp. (5.21)). In the following analysis, we derive conditions on the underlying parameters for which these closed-form solutions of (5.13) exist and satisfy the corresponding constraints on $\kappa_A^*, q\kappa_B^* > 0$, and find that these two cases encompass all possible game instances GL-P$(P, R_A, R_B, \mathbf{w})$.

Case 1 ($\mathcal{B}_1 = \mathcal{B}$): Substituting $\mathbf{p} = (P/W)\cdot\mathbf{w}$ into (5.16) and simplifying, we obtain

$$\begin{aligned} q\kappa_B^* &= \frac{1}{W}\left[P + R_A + \sqrt{R_A(R_A+2P)}\right] \\ \kappa_A^* &= \frac{(P+R_A)q\kappa_B^* - P^2/W}{qR_B}. \end{aligned} \tag{5.27}$$

Next, we verify that this solution satisfies the conditions $0 < q\kappa_B^* - P/W \leq \kappa_A^*$.

- $q\kappa_B^* - P/W > 0$: This holds by inspection.
- $q\kappa_B^* - P/W \leq \kappa_A$: We can write this condition as

$$qR_B - P \leq R_A + \frac{PR_A}{R_A + \sqrt{R_A(R_A+2P)}} \tag{5.28}$$

We note that whenever $qR_B - P < 0$, this condition is always satisfied. When $qR_B - P \geq 0$, this condition does not automatically hold, and an equivalent expression of (5.28) is given by

$$R_A \geq \frac{2(qR_B - P)^2}{P + 2(qR_B - P)}. \tag{5.29}$$

Observe that $R_A = \frac{2(qR_B-P)^2}{P+(qR_B-P)}$ satisfies (5.28) with equality, and is in fact the only real solution (one can reduce it to a cubic polynomial in R_A).

When these conditions hold, the equilibrium payoff $\pi_A(P, R_A, R_B)$ (computed from Lemma 5.3.3) is given by the expression (5.4).

Case 2 ($\mathcal{B}_2 = \mathcal{B}$): Substituting $\mathbf{p} = (P/W) \cdot \mathbf{w}$ into (5.21) and simplifying, we obtain

$$\kappa_A^* = \frac{2(qR_B - P)}{W} \text{ and } q\kappa_B^* = \frac{2(qR_B - P)^2}{WR_A}. \tag{5.30}$$

The solution satisfies the condition $0 < \kappa_A^* < q\kappa_B^* - P/W$ if and only if $qR_B - P > 0$ and $R_A > \frac{2(qR_B-P)^2}{P+(qR_B-P)}$. When this holds, the equilibrium payoff (from Lemma 5.3.3) is $\pi_A(P, R_A, R_B) = W \cdot \frac{R_A}{2(qR_B-P)}$. □

Part II

Preference-Based Mechanisms

Chapter 6

Introduction

6.1 Model: Congestion games with taxes

A congestion game (also known as, atomic congestion game) consists of a set of players $N = \{1, \ldots, n\}$ and a set of resources $\mathcal{R}$. Each player $i \in N$ selects an action a_i from a corresponding set of feasible actions $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$. The cost that a player experiences for selecting a given resource $r \in \mathcal{R}$ depends only on the total number of players selecting r, and is denoted as $C_r : \{1, \ldots, n\} \to \mathbb{R}$. Given an assignment $a = (a_1, \ldots, a_n) \in \mathcal{A}$, where $\mathcal{A} = \Pi_{i=1}^{N} \mathcal{A}_i$, each player $i \in N$ experiences a cost equal to the sum over costs on resources $r \in a_i$. Correspondingly, the system cost is measured by the sum of the players costs, i.e.,

$$\mathrm{SC}(a) = \sum_{i \in N} \sum_{r \in a_i} C_r(|a|_r) \tag{6.1}$$

where $|a|_r$ denotes the number of players selecting resource r in assignment a. Observe that a congestion game can be represented as a tuple $G = (N, \mathcal{R}, \{\mathcal{A}_i\}_{i \in N}, \{C_r\}_{r \in \mathcal{R}})$.

We denote by $\mathcal{G}^{n,\mathcal{L}}$ the family of all congestion game instances with a maximum number of players n, where all resource cost functions $\{\ell_r\}_{r \in \mathcal{R}}$ belong to a common family of resource cost functions $\mathcal{L}$. To ease the notation, we will use $\mathcal{G}^n$ to refer to the family $\mathcal{G}^{n,\mathcal{L}}$ when the dependence on $\mathcal{L}$ is clear.

6.1.1 Taxes

In the study of taxes, each resource $r \in \mathcal{R}$ is associated with a tax function $\tau_r : \{1, \ldots, n\} \to \mathbb{R}$ (positive or negative). In this case, each player $i \in N$ incurs a cost involving both the resource costs it experiences and the imposed taxes:

$$C_i(a) = \sum_{r \in a_i} \Big[\ell_r(|a|_r) + \tau_r(|a|_r) \Big]. \tag{6.2}$$

We consider taxes that only influence the players' costs and do not factor into the social cost. Scenarios where taxes are incorporated into the social cost have also been studied in, e.g., [72, 73].

When players selfishly choose their actions to minimize their incurred costs, an emergent outcome is often described by a *pure Nash equilibrium*. A pure Nash equilibrium is an assignment $a^{\mathrm{ne}} \in \mathcal{A}$ such that $C_i(a^{\mathrm{ne}}) \leq C_i(a_i, a^{\mathrm{ne}}_{-i})$ for all $a_i \in \mathcal{A}_i$ and all $i \in N$, where a'_i, a_{-i} denotes the assignment obtained when player i plays action a'_i and the remaining players continue to play their actions in a. Observe that a system designer can influence the set of pure Nash equilibria through the choice of tax functions τ_r, $r \in \mathcal{R}$. Accordingly, with abuse of notation, we augment the tuple representation of a congestion game as $G = (N, \mathcal{R}, \{\mathcal{A}_i\}_{i \in N}, \{\ell_r\}_{r \in \mathcal{R}}, \{\tau_r\}_{r \in \mathcal{R}})$ which incorporates the imposed taxes on the resources.

We consider the use of *local* taxes to improve the equilibrium performance. Local taxes only use information about the resource cost function ℓ_r to compute the tax function τ_r on any given resource $r \in \mathcal{R}$. The restriction to local taxes is a natural requirement, especially in settings where scalable and computationally simple rules are desirable. This structure is also commonly utilized in the existing literature, e.g., Pigouvian taxes [74]. Accordingly, we define a local taxation rule as a map from the set of admissible resource costs $\mathcal{L}$ to taxes, i.e., under a given local taxation rule T, the tax function associated with each resource $r \in \mathcal{R}$ is given by $\tau_r = T(\ell_r)$. For any given family of congestion games $\mathcal{G}$ without taxes, we denote by $\mathcal{G}_T$ the corresponding modified family of congestion games with taxes $\tau_r = T(\ell_r)$ on each edge $r \in \mathcal{R}$.

When convenient, we will use the notation of agent cost generating functions $F_r(k) = \ell_r(k) + \tau_r(k),\ k = 1, \ldots, n$, to denote the cost experienced by an agent on a given resource $r \in \mathcal{R}$. Under this modified notation, we may also refer to the family of games $\mathcal{G}_{\mathcal{P}}$ which contains all congestion games induced by a (possibly infinite) set $\mathcal{P}$ of admissible system-cost, agent-cost function pairs, i.e., $\mathcal{G}_{\mathcal{P}}$ is the set of all games G that satisfy $\{C_r, F_r\} \in \mathcal{P}$ for all $r \in \mathcal{R}$.

6.1.2 Performance measures

For any given family of congestion game instances $\mathcal{G}$, we may choose to measure the equilibrium performance using two commonly-studied metrics that we term *Price of Anarchy* and *Price of Stability*, respectively defined as

$$\mathrm{PoA}(\mathcal{G}) = \sup_{G \in \mathcal{G}} \max_{a \in \mathrm{NE}(G)} \frac{\mathrm{SC}(a)}{\mathrm{MinCost}(G)}, \tag{6.3}$$

$$\mathrm{PoS}(\mathcal{G}) = \sup_{G \in \mathcal{G}} \min_{a \in \mathrm{NE}(G)} \frac{\mathrm{SC}(a)}{\mathrm{MinCost}(G)}, \tag{6.4}$$

where $\mathrm{MinCost}(G)$ denotes the minimum achievable social cost for instance G as defined in (6.1) and $\mathrm{NE}(G)$ denotes the set of all pure Nash equilibria in G. It is important to note that the set $\mathrm{NE}(G)$ must be non-empty for any congestion game G, and, thus, that the Price of Anarchy and Price of Stability are well-defined. This holds since all congestion games are *potential games* [75]. A potential game is any game G for which there exists a function $\Phi : \mathcal{A} \to \mathbb{R}$ such that

$$\Phi(a) - \Phi(a_i, a_{-i}) = C_i(a) - C_i(a_i', a_{-i}), \quad \forall a_i \in \mathcal{A}_i, \forall i \in N, \forall a \in \mathcal{A}. \tag{6.5}$$

In particular, congestion games admit the following potential function [75]:

$$\Phi(a) = \sum_{r \in \cup a_i} \sum_{k=1}^{|a|_r} [\ell_r(k) + \tau_r(k)]. \tag{6.6}$$

Observe that the Price of Anarchy provides guarantees on the performance of *any* equilibrium in the set of games while the Price of Stability offers performance guarantees for the *best* equilibrium of any instance in the set. By definition, the metrics must satisfy $\text{PoA}(\mathcal{G}) \geq \text{PoS}(\mathcal{G}) \geq 1$ for any family $\mathcal{G}$.

6.1.3 Welfare maximization games

We may also consider *distributed welfare games* [76], which are the welfare maximization analogue to generalized congestion games. In these games, there is a set of agents $N = \{1, \dots, n\}$ and a finite set of resources $\mathcal{R}$, where each agent $i \in N$ has an associated set of admissible actions $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$ and each resource $r \in \mathcal{R}$ is associated with a resource welfare function $W_r : \{1, \dots, n\} \to \mathbb{R}_{\geq 0}$ and a utility generating function $F_r : \{1, \dots, n\} \to \mathbb{R}$. The system welfare and agent utility functions under a given collective action $a \in \mathcal{A}$ are defined as

$$W(a) = \sum_{r \in \cup a_i} W_r(|a|_r), \qquad U_i(a) = \sum_{r \in a_i} U_r(|a|_r).$$

As with congestion games, we will consider families of distributed welfare games $\mathcal{G}_{\mathcal{W},\mathcal{U}}$ which contain all distributed welfare games that satisfy $W_r \in \mathcal{W}$ and $U_r = \mathcal{U}(W_r)$ for all $r \in \mathcal{R}$, where $\mathcal{U} : \mathcal{W} \to \mathbb{R}^n$ is referred to as the utility rule.

Definition 6.1.1 (Basis)**.** *The set of welfare rules $\mathcal{W}$ is spanned from a set of basis welfare rules $\{w^1, \dots, w^m\}$ [1] if $\mathcal{W}$ can be characterized by any non-negative combination of the basis welfare rules $\{w^1, \dots, w^m\}$, i.e.,*

$$\mathcal{W} = \{w : w = \sum_{j=1}^{m} \alpha^j w^j,\ \forall \alpha^1, \dots, \alpha^m \geq 0\}. \tag{6.7}$$

We remark that the performance guarantees are identical regardless of if we consider the welfare set $\mathcal{W} = \{w^1, \dots, w^m\}$ or consider $\mathcal{W}$ to be the set spanned from the basis $\{w^1, \dots, w^m\}$.

[1] We assume, without loss of generality, that $w^j(1) = 1$.

Furthermore, we will restrict attention to welfare rules that are *submodular*, or informally, welfare rules that admit a notion of decreasing marginal-returns that are commonplace in many objectives relevant to engineered systems.

Definition 6.1.2 (Submodularity)**.** *A welfare rule w is submodular if $w(j+1) \geq w(j)$ for all $j \in \mathbb{N}$ and $w(j+1) - w(j)$ is non-increasing in j.*

Note that in distributed welfare games, the Price of Anarchy and Price of Stability are respectively defined as the minimum and maximum ratio between the social welfare $W(\cdot)$ at a pure Nash equilibrium, and the maximum achievable social welfare in the game. Thus, in contrast with congestion games, the Price of Anarchy and Price of Stability of any distributed welfare game satisfies $\mathrm{PoA}(G) \leq \mathrm{PoS}(G) \leq 1$.

6.2 Applications

Congestion games suitably model systems that involve the interactions of strategic users and an underlying shared infrastructure. The typical example is that of drivers sharing a road network. A major difficulty in designing such systems is that one must account for each user's decision making process in order to guarantee good overall system performance. The detrimental effects of selfish user behaviour on the performance of these systems have been observed in a variety of contexts, including unfair allocation of essential goods and services [77, 78], overexploitation of natural resources [79, 80] and congestion in internet and road-traffic networks [81, 82]. A widely studied approach for influencing the system performance is the use of taxes, which can come in the form of rewards or penalties. Examples of taxes include taxes levied on users whose decisions have a negative impact on the system performance, and rebates given to users for making decisions aligned with the greater good.

6.3 Related work

The Price of Anarchy was first introduced by [83] as a performance metric to characterize the equilibrium efficiency in games. The first exact characterization of the Price of Anarchy in congestion games was derived independently by [84] and [85] for affine congestion games without taxes. These results were later generalized to all polynomial congestion games without taxes by [86].

Characterizations of the Price of Anarchy without taxes naturally led to the study of taxes to improve worst-case efficiency guarantees. The design of taxes to optimize equilibrium efficiency guarantees falls under the broader literature on coordination rules introduced by [87]. Within the context of congestion games, [88] derive local and global congestion-independent rules that minimize the Price of Anarchy for linear resource costs. For polynomial congestion games, [72] consider the class of rules that use only information about a social optimum of each instance. Among taxation rules of this specialized class, they derive the best achievable Price of Anarchy guarantees in polynomial congestion games, as well as a methodology for computing the optimal taxes. [89] generalize these results beyond polynomial congestion games and show that the efficiency of optimal, polynomially-computable taxes (using *global* information) matches the corresponding bound on the hardness of approximation. [90] derive *local* taxes that minimize the Price of Anarchy in any class of congestion games, which are shown to have similar efficiency guarantees as the rules using global information from [72, 88, 89]. [91] derive upper bounds on the Price of Anarchy associated with the marginal cost rule in polynomial congestion games, which were later refined and generalized in [90].

Aside from the Price of Anarchy, another interesting metric that has been the subject of extensive analysis is the Price of Stability. The Price of Stability was defined by [92] (though its study dates back even earlier to, for example, [93]), who provide an exact characterization of this metric for a specialized class of congestion games. The exact Price of Stability for linear congestion games without taxes was derived by [94] and [95], followed by an exact characterization for all polynomial congestion games without taxes by [96]. [97] study the Price of Anarchy and Price of Stability in affine congestion games under various taxation rules (for example, altruism, congestion independent taxes).

Chapter 7

Smoothness and primal-dual methods

7.1 Roughgarden's smoothness

Though it is a powerful performance metric, direct computation of the PoA is extremely difficult. In fact, even the problems of computing the minimum achievable system cost and of computing a pure Nash equilibrium for a given congestion game are both associated with pessimistic hardness results even in restricted settings [98, 99, 89]. For this reason, researchers have developed analytical techniques aimed at tractably characterizing the PoA over various classes of games. One such technique that is widely used in the existing literature is Roughgarden's (λ, μ)-smoothness argument, formally defined in [100]. A cost minimization game G is termed (λ, μ)-smooth if the following two conditions are met:

(i) For all $a \in \mathcal{A}$, we have $\sum_{i=1}^{n} J_i(a) \geq C(a)$;

(ii) For all $a, a' \in \mathcal{A}$, there exist $\lambda > 0$ and $\mu < 1$ such that

$$\sum_{i \in N} J_i(a'_i, a_{-i}) \leq \lambda C(a') + \mu C(a). \tag{7.1}$$

If a game G is (λ, μ)-smooth, then the PoA of game G is upper bounded by

$$\mathrm{PoA}(G) \leq \frac{\lambda}{1-\mu}.$$

Observe that if all the games in a class $\mathcal{G}$ are shown to be (λ, μ)-smooth, then the PoA of the class $\mathrm{PoA}(\mathcal{G})$ is also upper bounded by $\lambda/(1-\mu)$. We refer to the best upper bound obtainable using a smoothness argument on a given class of games $\mathcal{G}$ as the *Robust Price of Anarchy* (RPoA), i.e.,

$$\begin{aligned} \mathrm{RPoA}(\mathcal{G}) := \inf_{\lambda \geq 0, \mu \leq 1} \frac{\lambda}{1-\mu} \quad & \text{subject to:} \\ & (7.1) \text{ holds } \forall G \in \mathcal{G}. \end{aligned} \tag{7.2}$$

Observe that, given a family of games $\mathcal{G}$, the RPoA is the optimal value of a fractional program (which can be reformulated, in this case, to a linear program). It is important to note that the RPoA represents only an upper bound on the PoA. More specifically, for any class of (λ, μ)-smooth cost minimization games $\mathcal{G}$, it holds that $\mathrm{PoA}(\mathcal{G}) \leq \mathrm{RPoA}(\mathcal{G})$, where it could be that $\mathrm{PoA}(\mathcal{G}) < \mathrm{RPoA}(\mathcal{G})$ as we show in the examples below:

Example 7.1.1 (Congestion games without taxes)**.** *For the family of all congestion games $\mathcal{G}$ with affine resource cost functions (i.e., in each game $G \in \mathcal{G}$, each $r \in \mathcal{R}$ has $\ell_r(x) = a_r \cdot x + b_r$ with $a_r, b_r \geq 0$), the optimal smoothness parameters are $\lambda = 5/3$ and $\mu = 1/3$, such that* $\mathrm{RPoA}(\mathcal{G}) = 5/2$, *which is equal to* $\mathrm{PoA}(\mathcal{G})$ *[84, 85].*

Example 7.1.2 (Congestion games under Pigouvian taxes)**.** *For the family of all congestion games $\mathcal{G}$ with affine resource cost functions under Pigouvian (marginal cost) taxes (i.e., $\tau_r(x) = \ell_r(x) \cdot x - \ell_r(x-1) \cdot (x-1)$ for all $r \in \mathcal{R}$), the optimal smoothness parameters are $\lambda = 17/5$ and $\mu = 2/5$, such that* $\mathrm{RPoA}(\mathcal{G}) = 17/3$. *However, using the techniques provided in the forthcoming Theorem 7.4.1, we compute the true PoA to be* $\mathrm{PoA}(\mathcal{G}) = 3$ *which is nearly 50% smaller. Thus, (λ, μ)-smoothness does not give a tight characterization of the PoA in this case.*

The above examples establish that the RPoA and PoA do not always match. Crucially, it immediately follows that any analytical approach for quantifying and/or optimizing the PoA that is based on smoothness (see, e.g., [86, 101, 102]) is *inadequate* for settings where the corresponding smoothness bounds are loose. Based on this observation, in the forthcoming section, we introduce a novel notion of smoothness that improves upon the PoA bound provided by the RPoA.

7.2 Generalized smoothness

In this section, we provide a generalization of the smoothness framework, termed *generalized smoothness*. We will then proceed to show how this new framework provides tighter efficiency bounds and covers a broader spectrum of problem settings than the original smoothness framework.

Definition 7.2.1 (Generalized smoothness)**.** *The cost minimization game G is (λ, μ)-generalized smooth if, for any two allocations $a, a' \in \mathcal{A}$, there exist $\lambda > 0$ and $\mu < 1$ satisfying,*

$$\sum_{i=1}^{n} \Big[J_i(a'_i, a_{-i}) - J_i(a) \Big] + C(a) \leq \lambda C(a') + \mu C(a). \tag{7.3}$$

Note that we maintain the notation of (λ, μ) as in the original notion of smoothness for ease of comparison. In the specific case when $\sum_{i=1}^{n} J_i(a) = C(a)$ for all $a \in \mathcal{A}$, observe that the smoothness conditions in (7.3) are equivalent to the original smoothness conditions in (7.1). As with (7.2), we define the *Generalized Price of Anarchy* (GPoA) of a class of cost minimization games $\mathcal{G}$ as the best upper bound obtainable using a generalized smoothness argument, i.e.,

$$\mathrm{GPoA}(\mathcal{G}) := \inf_{\lambda>0, \mu<1} \left\{ \frac{\lambda}{1-\mu} \text{ s.t. } (7.3) \text{ holds } \forall G \in \mathcal{G} \right\}. \tag{7.4}$$

With our first result we show that (i) PoA bounds under the generalized smoothness framework follow in the same way as the original smoothness framework without the restriction that $\sum_{i=1}^{n} J_i(a) \geq C(a)$ for all $a \in \mathcal{A}$, and (ii) the generalized smoothness framework provides stronger bounds on the PoA than the original smoothness framework whenever both are defined. For clarity, the proof of (i) follows trivially from [100]. The novelty of the result is in (ii), which establishes an ordering between RPoA, GPoA and PoA.

Proposition 7.2.1. *For any (λ, μ)-generalized smooth game G, the following statements hold:*

(i) The PoA of G is upper bounded as $\mathrm{PoA}(G) \leq \lambda/(1-\mu)$.

(ii) If the game G is (λ, μ)-smooth, then $\mathrm{RPoA}(G) \geq \mathrm{GPoA}(G) \geq \mathrm{PoA}(G)$. *Furthermore, if* $\sum_{i=1}^{n} J_i(a) > C(a)$ *holds for all $a \in \mathcal{A}$, then* $\mathrm{RPoA}(G) > \mathrm{GPoA}(G) \geq \mathrm{PoA}(G)$.

Further comparisons between the RPoA and GPoA can be made, as summarized by the following observations:

– *Observation #1*: The PoA and GPoA are shift-, and scale-invariant, i.e., for any given $\gamma > 0$ and $(\delta_1, \ldots, \delta_n) \in \mathbb{R}^n$,

$$\mathrm{PoA}((N, \mathcal{A}, C, \{J_i\})) = \mathrm{PoA}((N, \mathcal{A}, C, \{\gamma J_i + \delta_i\})),$$
$$\mathrm{GPoA}((N, \mathcal{A}, C, \{J_i\})) = \mathrm{GPoA}((N, \mathcal{A}, C, \{\gamma J_i + \delta_i\})).$$

Neither of these properties hold in general for the RPoA, i.e., for any given $\gamma > 0$ and $(\delta_1, \ldots, \delta_n) \in \mathbb{R}^n$,

$$\mathrm{RPoA}((N, \mathcal{A}, C, \{J_i\})) \neq \mathrm{RPoA}((N, \mathcal{A}, C, \{\gamma J_i + \delta_i\})),$$

except when $\gamma = 1$ and $(\delta_1, \ldots, \delta_n) = 0$.

– *Observation #2*: The RPoA is optimized by budget-balanced agent cost functions, i.e., $\sum_{i \in N} J_i(a) = C(a)$ for all $a \in \mathcal{A}$. In general, this does not hold for the PoA and GPoA.

– *Observation #3*: For a given cost minimization game G, we define an average coarse-correlated equilibrium as a probability distribution $\sigma \in \Delta(\mathcal{A})$ satisfying, for all $a' \in \mathcal{A}$,

$$\mathbb{E}_{a \sim \sigma}\left[\sum_{i=1}^{N} J_i(a)\right] \leq \mathbb{E}_{a \sim \sigma}\left[\sum_{i=1}^{N} J_i(a'_i, a_{-i})\right]. \tag{7.5}$$

Note that the set of average coarse correlated equilibria contains all of the game's pure Nash equilibria, mixed Nash equilibria, correlated equilibria and coarse correlated equilibria [100]. The GPoA tightly characterizes the average coarse correlated equilibrium performance of any cost minimization game G, and, thus, of any class of cost minimization games $\mathcal{G}$. The proof follows identically to the result by [103] that proves this claim for the RPoA under an alternative definition of average coarse correlated equilibrium. Unsurprisingly, these two definitions of average coarse correlated equilibrium match for games with $\sum_{i=1}^{n} J_i(a) = C(a)$, for all $a \in \mathcal{A}$.

The above observations are stated without proof for brevity, but can easily be verified by the reader.

So far, we have presented two different smoothness bounds aimed at quantifying the PoA: (i) RPoA, and (ii) GPoA. The result in Theorem 7.2.1 shows that, though the generalized smoothness conditions amount to a minor variation on the original smoothness conditions, the GPoA always provides better (i.e., tighter) bounds on the PoA than the RPoA. Recall that our primary reasoning for considering smoothness bounds like RPoA and GPoA is that computing the PoA directly is a difficult problem, so we wish to consider a surrogate metric that is both simpler to characterize and sufficiently representative of the PoA instead. Though we show that the GPoA is more representative of the PoA than the RPoA, it remains to be shown that either one of these bounds is actually simpler to compute than the PoA. To resolve these two concerns, we wish to address the following questions:

1. *Does optimizing the GPoA ever coincide with optimizing the PoA?*; and, if so, then

2. *Are there tractable techniques for optimizing the GPoA?*

In the next chapter, we will identify a broad class of games for which we can answer both these questions in the affirmative.

7.3 Smoothness from a primal-dual perspective

A recent and interesting perspective on smoothness examines such arguments from a primal-dual perspective [104, 103]. More specifically, if we formulate the computation of the Price of Anarchy (or any other performance measure of interest) as an optimization problem, the relevant smoothness conditions correspond with the constraints of the dual problem. Adapting an example from [103], observe that the worst-case coarse-correlated equilibrium efficiency of a cost-minimization game can be exactly computed as the optimal value of the following optimization problem:

$$\begin{aligned} &\underset{\sigma}{\text{maximize}} \ \sum_{a\in\mathcal{A}} \left[\sigma(a)\cdot\frac{\mathrm{SC}(a)}{\mathrm{MinCost}(G)}\right] \text{ subject to:} \\ &\sum_{a\in\mathcal{A}} \sigma(a)\cdot[C_i(a)-C_i(a_i',a_{-i})] \le 0,\ \forall a_i'\in\mathcal{A}_i,\ \forall i\in N, \\ &\sum_{a\in\mathcal{A}} \sigma(a) = 1, \quad \sigma(a)\ge 0,\ \forall a\in\mathcal{A}. \end{aligned} \tag{7.6}$$

Observe that the objective in the above optimization problem is the ratio between the expected social cost under the probability distribution σ and the minimum-achievable cost in the game G and that the constraints in the second and third lines ensure that σ is a coarse-correlated equilibrium and a valid probability distribution, respectively. Note that we can equivalently write the above optimization problem as the following linear program:

$$\begin{aligned} &\underset{\sigma}{\text{maximize}} \ \sum_{a\in\mathcal{A}} [\sigma(a)\cdot\mathrm{SC}(a)] \text{ subject to:} \\ &\sum_{a\in\mathcal{A}} \sigma(a)\cdot[C_i(a)-C_i(a_i',a_{-i})] \le 0,\ \forall a_i'\in\mathcal{A}_i,\ \forall i\in N, \\ &\sum_{a\in\mathcal{A}} \sigma(a) = 1/\mathrm{MinCost}(G), \quad \sigma(a)\ge 0,\ \forall a\in\mathcal{A}. \end{aligned} \tag{7.7}$$

where we normalize the distribution σ with $\mathrm{MinCost}(G)$ to simplify the objective. The dual of the above linear program is

$$\underset{z_1,\ldots,z_n \geq 0, p}{\text{minimize}} \; p \text{ subject to:}$$

$$\mathrm{SC}(a) - p \cdot \mathrm{MinCost}(G) + \sum_{i=1}^{n} \left[z_i \cdot [C_i(a_i', a_{-i}) - C_i(a)] \right] \leq 0, \quad \forall a_i' \in \mathcal{A}_i,\ \forall i \in N,\ \forall a \in \mathcal{A}. \tag{7.8}$$

Note that the optimal value of the above dual program exactly characterizes the worst-case coarse-correlated equilibrium efficiency of the game G. This follows from the equivalence of the original optimization problem and the primal linear program and from strong duality in linear programming. Note that the optimal value of the above dual program is already an upper bound on the Price of Anarchy, since any pure strategy Nash equilibrium of the game G is also a coarse-correlated equilibrium of G. After applying several restrictions to the above dual problem – namely, setting $z_1 = \cdots = z_n = z$, considering only those games with $\sum_{i=1}^{n} C_i(a) \geq \mathrm{SC}(a)$ for all $a \in \mathcal{A}$ and replacing $\mathrm{MinCost}(G)$ with $\mathrm{SC}(a')$ – and after the change of variables ..., the objective of the resulting program is to minimize $\lambda/(1-\mu)$, and the constraints are identical to Roughgarden's (λ, μ)-smoothness conditions in (7.1). Similarly, after applying the restrictions $z_1 = \cdots = z_n = z$ and replacing $\mathrm{MinCost}(G)$ with $\mathrm{SC}(a')$, and after the change of variables ..., the constraints in the above dual program are identical to the generalized smoothness condition in (7.3). From this perspective, it is unsurprising that generalized smoothness offers improved bounds on the Price of Anarchy than Roughgarden's smoothness since it relies on only a subset of the restrictions to the dual program for computing the worst-case coarse-correlated equilibrium efficiency.

From the above line of reasoning, several important observations can be made:

1. Under the appropriate restriction to cost minimization games that satisfy $\sum_{i=1}^{n} C_i(a) \geq \mathrm{SC}(a)$ for all $a \in \mathcal{A}$, Roughgarden's (λ, μ)-smoothness conditions in (7.1) characterize a strict upper bound on the Price of Anarchy for cost minimization games with $\sum_{i=1}^{n} C_i(a) > \mathrm{SC}(a)$ for all $a \in \mathcal{A}$. This is because all the constraints in the dual program

above will be loose, and the optimal value of the dual program is itself an upper bound on the Price of Anarchy.

2. Following a similar line of reasoning to the one above, we can leverage primal-dual arguments to obtain smoothness conditions for various other performance measures, including the Price of Stability. The general recipe is simple: First, formulate a primal program that characterizes an upper bound on the performance measure of interest; then, take the constraints of the dual program as your smoothness conditions.

3. For a given family of games $\mathcal{G}$, an upper-bound on the equilibrium efficiency can be characterized by applying the same smoothness condition to every game $G \in \mathcal{G}$ simultaneously, as in (7.2) or (7.4). In other words, the constraints in the dual program must now be satisfied for every game $G \in \mathcal{G}$.

The final observation above on computing equilibrium efficiency guarantees for a family of games is only helpful if the family of games in question has a modest number of games. For example, in the extreme case, a family of games may have infinitely many games, such as the family of affine congestion games in Example 7.1.1. In such a case, we cannot solve a linear program with infinitely many constraints. In the next section, we show how game parameterizations can be used to tractably obtain equilibrium efficiency guarantees, even for arbitrarily large families of games.

7.4 Balancing tightness and tractability

In the previous sections, we saw that applying a smoothness condition to a family of games $\mathcal{G}$, especially families containing infinitely many games, can be difficult, and that introducing game parameterizations can help reduce the computational burden of characterizing the corresponding efficiency guarantees. In this section, we present several popular game parameterizations that will be used in the proofs. The best game parameterization for a particular application will balance the tightness of the provided bounds with the tractability of the corresponding optimization problem.

For the purposes of this section, we restrict our attention to the family of congestion games $\mathcal{G}$ with a maximum number of players n, each with up to k actions. Each resource has resource-cost function equal to a non-negative linear combination over m basis cost functions; formally, every resource $r \in \mathcal{R}$ has m coefficients $\alpha_r^1, \ldots, \alpha_r^m \geq 0$ such that $C_r(x) = \sum_{j=1}^{m} [\alpha_r^j \cdot c^j(x)]$ for every game $G \in \mathcal{G}$, where $c^1, \ldots, c^m$ are basis cost functions.

7.4.1 Tightest but not tractable

Given the maximum number of users n, consider a game parameterization corresponding with a $(2^{kn}) \times (kn)$ table $\mathcal{R}$ whose rows are all the unique permutations of kn-long binary vectors $\mathbf{r} \in \{0,1\}^{kn}$. Under this parameterization, each row $\mathbf{r} \in \mathcal{R}$ corresponds with a different resource in the parameterized congestion game, and, collectively, the rows encode the users' actions $a_i^{\text{w-ne}}, a_i^{\text{b-ne}}, a_i^{\text{opt}}$ as follows:

Consider the resource e corresponding with the row $\mathbf{r} \in \mathcal{R}$, and let $a_i^{(K)}$ denote player i's K-th action. For any $i \in \{1, \ldots, n\}$, if $\mathbf{r}_{Kn+i} = 1$, then $e \in a_i^{(K)}$, else $e \notin a_i^{(K)}$. The coefficients in the basis representation of the 2^{kn} resource cost functions will be the decision variables of our final linear program (i.e., there are $(2^{kn}) \times m$ decision variables). This means that we can represent any game with $(2^{kn}) \times m$ parameters where n is the maximum number of users, and m is the number of basis cost functions m. Thus, any linear program formulated under this parameterization is not tractable for large n. Nonetheless, lower bounds for small n can be computed within a reasonable amount of time, as provided in Figure 9.3.

Under this parameterization, observe that we retain the granularity required to independently identify each player, and each of that player's actions. With that granularity, we can encode elaborate linear programs. However, this granularity comes at the cost of tractability; a primal program based on this parameterization has $2^{kn} \times m$ decision variables, which generally coincides with the number of smoothness conditions.

7.4.2 Looser but tractable

For a given congestion game $G \in \mathcal{G}_{\Delta(\mathcal{P})}$, our game parameterization is defined as follows for allocations $a, a' \in \mathcal{A}$: For every resource $r \in \mathcal{R}$, we define integers $x_r, y_r, z_r \geq 0$ where $x_r = |a|_r$ is the number of agents that select r in a, $y_r = |a'|_r$ is the number of agents that select r in a' and $z_r = |\{i \in N \text{ s.t. } r \in a_i\} \cap \{i \in N \text{ s.t. } r \in a_i'\}|$ is the number of agents that select r in both a and a'. Note that $1 \leq x_r + y_r - z_r \leq n$ and $z_r \leq \min\{x_r, y_r\}$ must hold for all $r \in \mathcal{R}$. For all $x, y, z \geq 0$ such that $1 \leq x + y - z \leq n$ and $z \leq \min\{x, y\}$, and all $j = 1, \ldots, m$, we define the parameters

$$\theta(x, y, z, j) = \sum_{r \in \mathcal{R}(x,y,z)} \alpha_r^j, \tag{7.9}$$

where $\mathcal{R}(x, y, z) = \{r \in \mathcal{R} \text{ s.t. } (x_r, y_r, z_r) = (x, y, z)\}$, and $\alpha_r^j \geq 0$, $j = 1, \ldots, m$, are the coefficients in the representation of the resource-cost, agent-cost function pair $\{C_r, F_r\}$. Although the parameterization into values $\theta(x, y, z, j) \geq 0$ is of size $\mathcal{O}(mn^3)$, we show in Step 2 of the forthcoming proof of Theorem 7.4.1 that only $\mathcal{O}(mn^2)$ parameters are needed in the computation of the PoA. Furtheremore, in Step 3 of the proof of Theorem 7.4.1, we show that this parameterization is still sufficiently tight to be used for characterizing the Price of Anarchy of any family of games $\mathcal{G}_{\mathcal{P}}$ under any set of function pairs $\mathcal{P}$

7.4.3 Tight, tractable PoA in generalized congestion games

The broader literature on congestion games is often interested in characterizing the PoA associated with the family of all congestion games, $\mathcal{G}_{\mathcal{P}}$, under a specified set (possibly infinite) of admissible resource-cost, agent-cost function pairs $\mathcal{P}$. Recall that for each game $G \in \mathcal{G}_{\mathcal{P}}$, it must be that each $r \in \mathcal{R}$ satisfies $\{C_r, F_r\} \in \mathcal{P}$. For ease of notation, we may choose to denote the family $\mathcal{G}_{\mathcal{P}}$ simply as $\mathcal{G}$ when the dependence on the set $\mathcal{P}$ is clear. Our next result shows that the GPoA provides a tight bound on the PoA associated with any family of generalized congestion games $\mathcal{G}_{\mathcal{P}}$.

Theorem 7.4.1. *For any set of resource-cost, agent-cost function pairs $\mathcal{P}$ and positive integer*

n, let $\mathcal{G}$ denote the family of all generalized congestion games with a maximum of n agents in which each resource $r \in \mathcal{R}$ satisfies $\{C_r, F_r\} \in \mathcal{P}$. It holds that $\mathrm{PoA}(\mathcal{G}) = \mathrm{GPoA}(\mathcal{G})$.

Theorem 7.4.1 highlights that the GPoA represents a tight bound on the PoA for the family $\mathcal{G}_{\mathcal{P}}$ under any set of resource-cost, agent-cost function pairs $\mathcal{P}$. Therefore, for this broad class of problems, there is no loss in characterizing the PoA using the generalized smoothness bound. However, it remains to be shown whether $\mathrm{GPoA}(\mathcal{G}_{\mathcal{P}})$ can be quantified efficiently.

In many commonly studied settings, we can leverage the structure of the set $\mathcal{P}$ to efficiently quantify the GPoA of the family of all generalized congestion games under $\mathcal{P}$. For instance, in the forthcoming proof of Theorem 7.4.1, we show that the $\mathrm{GPoA}(\mathcal{G}_{\mathcal{P}})$ in (7.4) can be computed efficiently when the number of resource-cost, agent-cost function pairs in $\mathcal{P}$ and the maximum number of agents are finite. Moreover, the GPoA may be computable even when the size of $\mathcal{P}$ is not finite. Specifically, let $\mathcal{P} = \{\{C^1, F^1\}, \dots, \{C^m, F^m\}\}$ denote any (finite) set of m resource-cost, agent-cost function pairs, and let $\Delta(\mathcal{P})$ denote the set of all resource-cost, agent-cost function pairs $\{C, F\}$ that can be represented as

$$C(k) = \sum_{j=1}^{m} \alpha^j \cdot C^j(k), \quad k = 1, \dots, n,$$
$$F(k) = \sum_{j=1}^{m} \alpha^j \cdot F^j(k), \quad k = 1, \dots, n,$$

with $\alpha^1, \dots, \alpha^m \geq 0$.[1] In the proof of Theorem 7.4.1, we establish that the GPoA of the family $\mathcal{G}_{\Delta(\mathcal{P})}$ is equal to the GPoA of $\mathcal{G}_{\mathcal{P}}$, and, thus, can also be computed as the solution of a tractable linear program. We formally state this observation in the following corollary, where we define $C^j(0) = F^j(0) = F^j(n+1) = 0$, for $j = 1, \dots, m$, for ease of notation:

Corollary 7.4.1. *For any set of resource-cost, agent-cost function pairs $\mathcal{P}$ and positive integer n, let $\mathcal{G}_{\mathcal{P}}$ and $\mathcal{G}_{\Delta(\mathcal{P})}$ denote the families of all generalized congestion games with a maximum of n agents under their specified sets of function pairs. Then, it holds that* $\mathrm{PoA}(\mathcal{G}_{\mathcal{P}}) = \mathrm{PoA}(\mathcal{G}_{\Delta(\mathcal{P})})$.

[1] For example, the family of affine congestion games (studied in, e.g., [84, 85]) is equivalent to the family $\mathcal{G}_{\Delta(\mathcal{P})}$ with $m = 2$, where $\{C^1(k), F^1(k)\} = \{k, 1\}$ and $\{C^2(k), F^2(k)\} = \{k^2, k\}$, $k = 1, \dots, n$.

Further, let ρ^{opt} *be the optimal value of the following (tractable) linear program:*

$$\begin{aligned}\rho^{\text{opt}} = &\underset{\nu\in\mathbb{R}_{\geq 0},\rho\in\mathbb{R}}{\text{maximize}} \quad \rho \quad \text{subject to:}\\ &C(y) - \rho C(x) + \nu[(x-z)F(x) - (y-z)F(x+1)] \geq 0,\\ &\qquad\qquad \forall\{C,F\}\in\mathcal{P}, \quad \forall(x,y,z)\in\mathcal{I}_{\mathcal{R}}(n),\end{aligned} \tag{7.10}$$

where $\mathcal{I}_{\mathcal{R}}(n)$ *is defined in the forthcoming* (7.14)*. Then, it holds that* $\text{PoA}(\mathcal{G}_{\mathcal{P}}) = \text{GPoA}(\mathcal{G}_{\mathcal{P}}) = 1/\rho^{\text{opt}}$.

The linear program in (7.10) has two decision variables and $\mathcal{O}(mn^2)$ constraints. Thus, for m and n finite, there are computationally efficient approaches for characterizing the PoA of generalized congestion games.

7.5 Chapter proofs

7.5.1 Proofs from Section 7.2

Proof of Proposition 7.2.1. For the proof of statement (i), observe that, for all $a^{\text{ne}} \in \text{NE}(G)$ and $a^{\text{opt}} \in \mathcal{A}$,

$$\begin{aligned}C(a^{\text{ne}}) &\leq \sum_{i=1}^{n}\left[J_i(a_i^{\text{opt}}, a_{-i}^{\text{ne}}) - J_i(a^{\text{ne}})\right] + C(a^{\text{ne}})\\ &\leq \lambda\, C(a^{\text{opt}}) + \mu\, C(a^{\text{ne}}).\end{aligned} \tag{7.11}$$

The inequalities hold by the Nash equilibrium condition and (7.3), respectively. Rearranging gives the result.

The remainder of the proof focuses on statement (ii). Since the condition $\sum_{i=1}^n J_i(a) \geq C(a)$ for all $a \in \mathcal{A}$ implies that any pair of (λ, μ) satisfying (7.1) necessarily satisfies (7.3), we note that the GPoA is less than or equal to the RPoA, i.e., $\text{RPoA}(G) \geq \text{GPoA}(G) \geq \text{PoA}(G)$.

Note that for any game $G = (N, \mathcal{A}, C, \{J_i\})$ with $\sum_{i=1}^n J_i(a) > C(a)$ for all $a \in \mathcal{A}$ there must exist a uniform scaling factor $0 < \gamma < 1$ such that $\sum_{i=1}^n \gamma J_i(a) \geq C(a)$, but for which the PoA remains the same, i.e., for $G' = (N, \mathcal{A}, C, \{J_i'\})$ where $J_i' = \gamma J_i$, it holds that $\text{PoA}(G') =$

PoA(G). The PoA remains the same despite the rescaling, because the Nash equilibrium condition for each player is unaffected by a positive scaling factor (i.e., $\mathrm{NE}(G) = \mathrm{NE}(G')$), and because the optimal cost remains unchanged since the scaling does not impact the system cost. Further, one can verify from (7.1) that $\mathrm{RPoA}(G) > \mathrm{RPoA}(G')$, and thus $\mathrm{RPoA}(G) > \mathrm{RPoA}(G') \geq \mathrm{PoA}(G') = \mathrm{PoA}(G)$. Finally, we know that $\mathrm{GPoA}(G')$ is less than or equal to $\mathrm{RPoA}(G')$ and can verify from (7.3) that $\mathrm{GPoA}(G) = \mathrm{GPoA}(G')$. Thus, $\mathrm{RPoA}(G) > \mathrm{RPoA}(G') \geq \mathrm{GPoA}(G') = \mathrm{GPoA}(G) \geq \mathrm{PoA}(G)$. □

7.5.2 Proof of Theorem 7.4.1

The following informal outline of the proof is directly followed by the formal proof, which follows a similar structure:

– *Step 1* : We define our game parameterization, which represents any generalized congestion game $G \in \mathcal{G}$ with $\mathcal{O}(mn^3)$ parameters $\theta(x, y, z, j) \geq 0$ corresponding with basis pairs $\{(C^j, F^j)\}$, $j = 1, \ldots, m$, and triplets $x, y, z \in \{0, \ldots, n\}$ such that $1 \leq x + y - z \leq n$ and $z \leq \min\{x, y\}$.

– *Step 2* : For any family of generalized congestion games $\mathcal{G}$, we observe that an upper bound on the GPoA can be computed as a fractional program with $\mathcal{O}(mn^2)$ constraints under the game parameterization presented in Step 1.

– *Step 3* : Following a change of variables, we observe that the linear program in (7.10) is equivalent to the fractional program from Step 2. We then provide a game $G \in \mathcal{G}$ with PoA equal to the upper bound on the GPoA, implying that $\mathrm{PoA}(G) \geq \mathrm{GPoA}(\mathcal{G})$. Since $\mathrm{PoA}(G) \leq \mathrm{PoA}(\mathcal{G}) \leq \mathrm{GPoA}(\mathcal{G})$, it must then be that $\mathrm{PoA}(G) = \mathrm{PoA}(\mathcal{G}) = \mathrm{GPoA}(\mathcal{G})$ for the family of generalized congestion games $\mathcal{G}$, concluding the proof.

Proof. It is straightforward to show that any game in the family $\mathcal{G}_{\Delta(\mathcal{P})}$ is (strategically) equivalent to a game in $\mathcal{G}_{\mathcal{P}}$ (with potentially a much larger resource set), and, thus, that $\mathrm{PoA}(\mathcal{G}_{\mathcal{P}}) = \mathrm{PoA}(\mathcal{G}_{\Delta(\mathcal{P})})$. For example, consider a game $G \in \mathcal{G}_{\Delta(\mathcal{P})}$ with rational coefficients $\alpha_r^1, \ldots, \alpha_r^m \geq 0$, $r \in \mathcal{R}$. Let LCD denote the lowest common denominator across all coefficients $\alpha_r^1, \ldots, \alpha_r^m \geq 0$, $r \in \mathcal{R}$, and observe that $\mathrm{LCD} \cdot \alpha_r^j$ is an integer for each $j \in \{1, \ldots, m\}$ and $r \in \mathcal{R}$. Thus, G

is equivalent to the game $G' \in \mathcal{G}_{\mathcal{P}}$ where we replace each resource r in game G with $\mathrm{LCD} \cdot \alpha_r^j$ resources with function pair $\{C^j, F^j\}$, for each type $j \in \{1, \ldots, m\}$. As this amounts to a uniform rescaling of the resource-cost and agent-cost functions, the PoA remains unchanged. In the case of irrational coefficients, we can approximate these from above or below (as required) to arbitrary precision using rational numbers, and then use the above approach to obtain an equivalent game in $\mathcal{G}_{\mathcal{P}}$. The above reasoning is further elaborated in [90, 100].

For any given set $\mathcal{P}$, the remainder of the proof shows that $\mathrm{PoA}(\mathcal{G}_{\Delta(\mathcal{P})}) = \mathrm{GPoA}(\mathcal{G}_{\Delta(\mathcal{P})})$. The proof is shown in three steps, as summarized in the informal outline.

– *Step 1* : For a given game $G \in \mathcal{G}_{\Delta(\mathcal{P})}$, our game parameterization is defined as follows for allocations $a, a' \in \mathcal{A}$: For every resource $r \in \mathcal{R}$, we define integers $x_r, y_r, z_r \geq 0$ where $x_r = |a|_r$ is the number of agents that select r in a, $y_r = |a'|_r$ is the number of agents that select r in a' and $z_r = |\{i \in N \text{ s.t. } r \in a_i\} \cap \{i \in N \text{ s.t. } r \in a_i'\}|$ is the number of agents that select r in both a and a'. Note that $1 \leq x_r + y_r - z_r \leq n$ and $z_r \leq \min\{x_r, y_r\}$ must hold for all $r \in \mathcal{R}$. For all $x, y, z \geq 0$ such that $1 \leq x + y - z \leq n$ and $z \leq \min\{x, y\}$, and all $j = 1, \ldots, m$, we use the parameterization in (7.9) to represent the coefficiencts of the resource-cost, agent-cost function pair $\{C_r, F_r\}$. Although the parameterization into values $\theta(x, y, z, j) \geq 0$ is of size $\mathcal{O}(mn^3)$, we show in Step 2 that only $\mathcal{O}(mn^2)$ parameters are needed in the computation of the PoA.

– *Step 2* : For any generalized congestion game $G \in \mathcal{G}_{\Delta(\mathcal{P})}$, we denote an optimal allocation as a^{opt}, and a Nash equilibrium as a^{ne}, i.e. $a^{\mathrm{ne}} \in \mathrm{NE}(G)$ such that $\mathrm{PoA}(G) \geq C(a^{\mathrm{ne}})/C(a^{\mathrm{opt}})$. We observe that using the above definitions of (x_r, y_r, z_r) for $a = a^{\mathrm{ne}}$ and $a' = a^{\mathrm{opt}}$, it follows that

$$\sum_{i=1}^{n} J_i(a_i^{\mathrm{opt}}, a_{-i}^{\mathrm{ne}}) = \sum_{r \in \mathcal{R}} \Big[(y_r - z_r)F_r(x_r + 1) + z_r F_r(x_r)\Big].$$

Informally, if an agent $i \in N$ selects a given resource $r \in \mathcal{R}$ in both a_i^{ne} and a_i^{opt}, then by deviating from a_i^{ne} to a_i^{opt}, the agent does not add to the load on r, i.e., $|a_i^{\mathrm{opt}}, a_{-i}^{\mathrm{ne}}|_r = |a^{\mathrm{ne}}|_r = x_r$. However, if $r \in a_i^{\mathrm{opt}}$ and $r \notin a_i^{\mathrm{ne}}$, then $|a_i^{\mathrm{opt}}, a_{-i}^{\mathrm{ne}}|_r = |a^{\mathrm{ne}}|_r + 1 = x_r + 1$.

Recall that for all $r \in \mathcal{R}$, it must hold that $z_r \leq \min\{x_r, y_r\}$, and $1 \leq x_r + y_r - z_r \leq n$. We

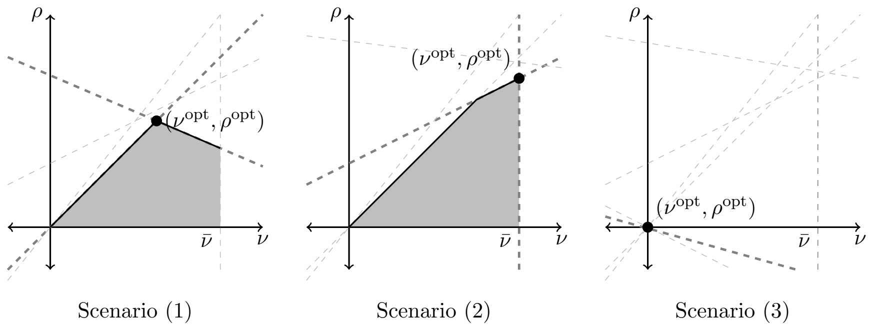

Figure 7.1: *The three different scenarios in which optimal solutions* $(\nu^{\rm opt}, \rho^{\rm opt})$ *to* (7.15) *can exist.* We illustrate the reasoning behind each of the three scenarios for optimal solutions $(\nu^{\rm opt}, \rho^{\rm opt})$ to the linear program in (7.15). Since the objective of (7.15) is to maximize ρ, the optimal values will be at the (upper) boundary of the feasible set, illustrated with a solid, bolded line in each of the examples above. Additionally, the optimal solution $(\nu^{\rm opt}, \rho^{\rm opt})$ is marked by a solid, black dot in the illustrations above. In Scenario (1), on the left, $(\nu^{\rm opt}, \rho^{\rm opt})$ lie on the intersection of a boundary line with positive slope and a boundary line with non-positive slope. In Scenario (2), centre, $(\nu^{\rm opt}, \rho^{\rm opt})$ lie on the intersection of a boundary line with positive slope at $\nu = \bar{\nu}$, which is defined in (7.16). In Scenario (3), on the right, there exists a halfplane boundary line with nonpositive slope and ρ-intercept equal to zero, and so $(\nu^{\rm opt}, \rho^{\rm opt}) = (0, 0)$. Using the parameters corresponding to the halfplanes on which the pair $(\nu^{\rm opt}, \rho^{\rm opt})$ lays, we can construct games $G \in \mathcal{G}$ with $\mathrm{PoA}(G) = 1/\rho^{\rm opt}$ in each of these scenarios.

define the set of triplets $\mathcal{I}(n) \subseteq \{0, 1, \ldots, n\}^3$ as

$$\mathcal{I}(n) := \{(x, y, z) \in N^3 \,|\, 1 \le x + y - z \le n \text{ and } z \le \min\{x, y\}\}, \tag{7.12}$$

and $\gamma(\mathcal{G}_{\Delta(\mathcal{P})})$ as the value of the following fractional program:

$$\begin{aligned} \gamma(\mathcal{G}_{\Delta(\mathcal{P})}) := \inf_{\lambda>0, \mu<1} \quad & \frac{\lambda}{1-\mu} \quad \text{subject to:} \\ (z-x)F^j(x) + (y-z)F^j(x+1) + C^j(x) &\le \lambda C^j(y) + \mu C^j(x), \\ \forall j = 1, \ldots, m, \quad & \forall (x, y, z) \in \mathcal{I}(n). \end{aligned} \tag{7.13}$$

Observe that, by (7.9), the generalized smoothness condition in (7.3) can be rewritten for a given generalized congestion game as

$$\begin{aligned} &\sum_{\mathcal{I}(n)} \sum_{j=1}^{m} \left[(x-z)F^j(x) - (y-z)F^j(x+1) + C^j(x)\right] \theta(x, y, z, j) \\ &\le \sum_{\mathcal{I}(n)} \sum_{j=1}^{m} \left[\lambda C^j(y) + \mu C^j(x)\right] \theta(x, y, z, j) \end{aligned}$$

It must then hold that for any pair (λ, μ) in the feasible set of the fractional program in (7.13), all games $G \in \mathcal{G}_{\Delta(\mathcal{P})}$ are (λ, μ)-generalized smooth, i.e., $\gamma(\mathcal{G}_{\Delta(\mathcal{P})}) \geq \mathrm{GPoA}(\mathcal{G}_{\Delta(\mathcal{P})})$. This is because the generalized smoothness condition for generalized congestion games can be expressed as a weighted sum with positive coefficients over a subset of the constraints in (7.13).

To conclude Step 2 of the proof, we show that it is sufficient to define $\gamma(\mathcal{G}_{\Delta(\mathcal{P})})$ in (7.13) over the reduced set of constraints corresponding to $j \in \{1, \dots, m\}$ and triplets in $\mathcal{I}_{\mathcal{R}}(n) \subseteq \mathcal{I}(n)$, where $\mathcal{I}(n)$ is defined as in (7.12) and

$$\begin{aligned} \mathcal{I}_{\mathcal{R}}(n) :=& \{(x, y, z) \in \mathcal{I}(n) \text{ s.t. } x + y - z = n\} \\ &\cup \{(x, y, z) \in \mathcal{I}(n) \text{ s.t. } (x - z)(y - z)z = 0\}. \end{aligned} \tag{7.14}$$

For each $j \in \{1, \dots, m\}$ and any $(x, y, z) \in \mathcal{I}(n)$, observe that the constraint in (7.13) is equivalent to $yF^j(x+1) - xF^j(x) + z[F^j(x) - F^j(x+1)] \leq \lambda C^j(y) + (\mu - 1)C^j(x)$. If $F^j(x+1) \geq F^j(x)$, the strictest condition on λ and μ corresponds to the lowest value of z. Thus, $z = \max\{0, x + y - n\}$, and either $(x - z)(y - z)z = 0$ or $x + y - z = n$. Otherwise, if $F^j(x+1) < F^j(x)$, then the largest value of z is strictest, i.e., $z = \min\{x, y\}$ which satisfies $(x - z)(y - z)z = 0$.

– *Step 3* : In order to derive the game instances with PoA matching $\gamma(\mathcal{G}_{\Delta(\mathcal{P})})$, it is convenient to perform the following change of variables: $\nu(\lambda, \mu) := 1/\lambda$ and $\rho(\lambda, \mu) := (1 - \mu)/\lambda$. For ease of notation, we will refer to the new variables simply as ν and ρ, respectively, i.e., $\nu = \nu(\lambda, \mu)$ and $\rho = \rho(\lambda, \mu)$. For each $j \in \{1, \dots, m\}$ and each $(x, y, z) \in \mathcal{I}_{\mathcal{R}}(n)$, it is straightforward to verify that the constraints in (7.13) can be rewritten in terms of ν and ρ as

$$C^j(y) - \rho C^j(x) + \nu[(x - z)F^j(x) - (y - z)F^j(x+1)] \geq 0.$$

Thus, the value $\gamma(\mathcal{G}_{\Delta(\mathcal{P})})$ must be equal to $1/\rho^{\text{opt}}$, where ρ^{opt} is the value of the following linear program:

$$\begin{aligned}
&\rho^{\text{opt}} = \underset{\nu\in\mathbb{R}_{\geq 0},\rho\in\mathbb{R}}{\text{maximize}} \quad \rho \quad \text{subject to:}\\
&C^j(y)-\rho C^j(x)+\nu[(x-z)F^j(x)-(y-z)F^j(x+1)] \geq 0, \\
&\qquad\qquad \forall j = 1,\dots,m, \quad \forall (x,y,z)\in\mathcal{I}_{\mathcal{R}}(n).
\end{aligned} \tag{7.15}$$

It is important to note here that while $\gamma(\mathcal{G}_{\Delta(\mathcal{P})})$ is the infimum of a fractional program (see, e.g., (7.13)), the value ρ^{opt} can be computed as a maximum because the feasible set is bounded and closed. Firstly, since $\gamma(\mathcal{G}_{\Delta(\mathcal{P})})$ is an upper bound on the PoA, its inverse (i.e., ρ) must be in the bounded and closed interval $[0,1]$. Additionally, one can verify that ν is not only bounded from below by 0, but also from above by the quantity

$$\begin{aligned}
&\bar{\nu} := \min_{j\in\{1,\dots,m\}} \underset{(x,y,z)\in\mathcal{I}_{\mathcal{R}}(n)}{\text{minimize}} \frac{C^j(y)}{(y-z)F^j(x+1)-(x-z)F^j(x)}\\
&\text{s.t. } (x-z)F^j(x)-(y-z)F^j(x+1) < 0, C^j(x) = 0,
\end{aligned} \tag{7.16}$$

which comes from the constraints in (7.15) corresponding to triplets $(x,y,z)\in\mathcal{I}_{\mathcal{R}}(n)$ such that $C^j(x)=0$ and $(x-z)F^j(x)-(y-z)F^j(x+1)<0$. Such a value must exist, as we define $C^j(0)=0$. One can verify that any $j\in\{1,\dots,m\}$ and $(x,y,z)\in\mathcal{I}_{\mathcal{R}}(n)$ such that $C^j(x)=0$ and $(x-z)F^j(x)-(y-z)F^j(x+1)\geq 0$ correspond to constraints that are satisfied trivially in (7.15) since $\nu\geq 0$, by definition, and $C^j(y)\geq 0$ for all $y=0,1,\dots,n$, by assumption.

We denote with $\mathcal{H}^j(x,y,z)$ the halfplane of (ν,ρ) values that satisfy the constraint corresponding to $j\in\{1,\dots,m\}$ and $(x,y,z)\in\mathcal{I}_{\mathcal{R}}(n)$, i.e.,

$$\begin{aligned}
&\mathcal{H}^j(x,y,z) := \Big\{(\nu,\rho)\in\mathbb{R}_{\geq 0}\times\mathbb{R} \text{ s.t.}\\
&\rho \leq \frac{C^j(y)}{C^j(x)} + \frac{1}{C^j(x)}\nu\left[(x-z)F^j(x)-(y-z)F^j(x+1)\right]\Big\}.
\end{aligned}$$

The set of feasible (ν,ρ) is the intersection of these $m\times|\mathcal{I}_{\mathcal{R}}(n)|$ halfplanes. Since the objective is to maximize ρ, any solution $(\nu^{\text{opt}},\rho^{\text{opt}})$ to the linear program in (7.15) must be on the (upper) boundary of the feasible set. We argue below that a solution $(\nu^{\text{opt}},\rho^{\text{opt}})$ can only exist in one of the three following scenarios: (1) at the intersection of two halfplanes' boundaries, where

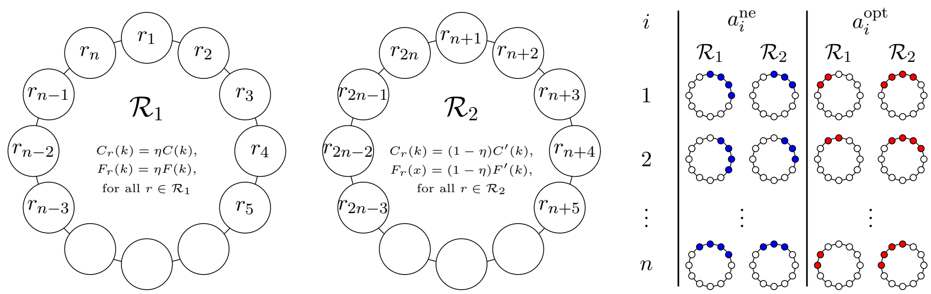

Figure 7.2: *The game instance construction G consisting of n agents, and two disjoint cycles $\mathcal{R}_1$ and $\mathcal{R}_2$, as described in the proof of Theorem 7.4.1, Step 2 for Scenarios (1) and (2).* Consider the family of games $\mathcal{G}_{\Delta(\mathcal{P})}$, where n is the maximum number of agents and $\mathcal{P}$ is the set of basis functions pairs, and suppose that $(\nu^{\text{opt}}, \rho^{\text{opt}})$ satisfy the conditions of Scenarios (1) or (2). Further, suppose that the parameters for which (7.17) and (7.18) hold are $C, F, C', F' \in \mathcal{P}$, $(x, y, z) = (4, 2, 0), (x', y', z') = (3, 4, 2) \in \mathcal{I}_{\mathcal{R}}(n)$ and $\eta \in [0, 1]$. In the above figure, we illustrate the game $G \in \mathcal{G}_{\Delta(\mathcal{P})}$ such that $\text{PoA}(G) = \text{PoA}(G') = 1/\rho^{\text{opt}}$ according to the reasoning for constructing game instances in Scenarios (1) and (2). Observe that each resource $r \in \mathcal{R}_1$ has $C_r(k) = \eta C(k)$, and $F_r(k) = \eta F(k)$, whereas each resource $r \in \mathcal{R}_2$ has $C_r(k) = (1-\eta)C'(k)$, and $F_r(x) = (1-\eta)F'(k)$, for all $k \in \{1, \ldots, n\}$. Each agent $i \in N$ has two actions a_i^{ne} and a_i^{opt}, as defined in the table on the right. Observe that every resource in $\mathcal{R}_1$ is selected by 4 agents in the allocation $a^{\text{ne}} = (a_1^{\text{ne}}, \ldots, a_n^{\text{ne}})$, and 3 agents in $a^{\text{opt}} = (a_1^{\text{opt}}, \ldots, a_n^{\text{opt}})$, where no agent $i \in N$ has a common resource between its actions a_i^{ne} and a_i^{opt}, i.e., $x_r = 4 = x$, $y_r = 3 = y$, and $z_r = 0 = z$ for all $r \in \mathcal{R}_1$. Similarly, $x_r = 3 = x'$, $y_r = 4 = y'$, and $z_r = 2 = z'$, for each resource $r \in \mathcal{R}_2$.

one halfplane has boundary line with positive slope, and the other has boundary line with nonpositive slope; (2) on a halfplane boundary line with positive slope at $\nu^{\text{opt}} = \bar{\nu}$; or (3) at $(\nu^{\text{opt}}, \rho^{\text{opt}}) = (0, 0)$.

We denote with $\partial\mathcal{H}^j(x, y, z)$ the boundary line of the halfplane $\mathcal{H}^j(x, y, z)$, i.e., the set of $(\nu, \rho) \in \mathbb{R}_{\geq 0} \times \mathbb{R}$ such that the inequality in the definition of $\mathcal{H}^j(x, y, z)$ holds with equality. Observe that the boundary lines of halfplanes corresponding to the choice $y = z = 0$ have ρ-intercept equal to zero and slope $xF^j(x)/C^j(x)$. If $F^j(x) \leq 0$ for any $j \in \{1, \ldots, m\}$ and $x \in \{1, \ldots, n\}$, then an optimal pair (ν, ρ) is trivially at the origin, i.e., $(\nu^{\text{opt}}, \rho^{\text{opt}}) = (0, 0)$ (i.e., Scenario (3) above). Note that the ρ-intercept of any halfplane boundary cannot be below 0, as we only consider cost functions such that $C^j(k) \geq 0$ for all k and all j. Otherwise, the maximum value of ρ occurs at the intersection of a boundary line with positive slope and a boundary line with nonpositive slope (i.e., Scenario (1) above) or on a boundary line with positive slope at $\nu = \bar{\nu}$ (i.e., Scenario (2) above). We illustrate this reasoning in Figure 7.1.

Observe that for Scenarios (1) and (2), the pair $(\nu^{\rm opt}, \rho^{\rm opt})$ is at the intersection of two boundary lines, which we denote as $\partial\mathcal{H}^j(x,y,z)$ and $\partial\mathcal{H}^{j'}(x',y',z')$. The parameters $j,j' \in \{1,\dots,m\}$ and $(x,y,z),(x',y',z') \in \mathcal{I}_{\mathcal{R}}(n)$ satisfy the following:

$$\begin{aligned}\rho^{\rm opt}C^j(x) - C^j(y) &= \nu^{\rm opt}[(x{-}z)F^j(x) - (y{-}z)F^j(x{+}1)],\\ \rho^{\rm opt}C^{j'}(x'){-}C^{j'}(y') &= \nu^{\rm opt}[(x'{-}z')F^{j'}(x'){-}(y'{-}z')F^{j'}(x'{+}1)],\end{aligned} \tag{7.17}$$

because $(\nu^{\rm opt}, \rho^{\rm opt})$ is on both boundary lines. Further, there must exist $\eta \in [0,1]$ such that

$$\begin{aligned}0 =&\eta\left[(x{-}z)F^j(x) - (y{-}z)F^j(x{+}1)\right]\\ &+ (1{-}\eta)\left[(x'{-}z')F^{j'}(x') - (y'{-}z')F^{j'}(x'{+}1)\right].\end{aligned} \tag{7.18}$$

(7.18) holds in Scenario (1) because one of the boundary lines has positive slope, i.e., $(x-z)F^j(x) - (y-z)F^j(x+1) > 0$, while the other has nonpositive slope, and in Scenario (3) because one boundary line has positive slope while the other is the vertical line $\nu = \bar{\nu}$ which corresponds to a particular choice of $j \in \{1,\dots,m\}$ and $(x,y,z) \in \mathcal{I}_{\mathcal{R}}(n)$ such that $(x-z)F^j(x) - (y-z)F^j(x+1) < 0$ by (7.16).

Next, for the parameters $j,j' \in \{1,\dots,m\}$, $(x,y,z),(x',y',z') \in \mathcal{I}_{\mathcal{R}}(n)$, and $\eta \in [0,1]$ obtained above, we construct a game instance $G \in \mathcal{G}_{\Delta(\mathcal{P})}$ such that $\mathrm{PoA}(G) = 1/\rho^{\rm opt}$. Let $\mathcal{R}_1 = \{r_1,\dots,r_n\}$ and $\mathcal{R}_2 = \{r_{n+1},\dots,r_{2n}\}$ denote two disjoint cycles of resources. Every resource $r \in \mathcal{R}_1$ has cost function $C_r(k) = \eta C^j(k)$, and agent-cost function $F_r(k) = \eta F^j(k)$ for all k. Meanwhile, every $r \in \mathcal{R}_2$ has cost function $C_r(k) = (1-\eta)C^{j'}(k)$, and cost generating function $F_r(k) = (1-\eta)F^{j'}(k)$ for all k. We define the agent set $N = \{1,\dots,n\}$, where each agent $i \in N$ has action set $\mathcal{A}_i = \{a_i^{\rm ne}, a_i^{\rm opt}\}$. In action $a_i^{\rm ne}$, the agent i selects x consecutive resources in $\mathcal{R}_1$ starting with r_i, i.e. $\{r_i, r_{(i \bmod n)+1},\dots,r_{((i+x-2) \bmod n)+1}\}$, and x' consecutive resources in $\mathcal{R}_2$ starting with resource r_{n+i}. In $a_i^{\rm opt}$, agent i selects y consecutive resources in $\mathcal{R}_1$ ending with resource $r_{((i+z-2) \bmod n)+1}$, i.e. $\{r_{((i+z-y-1) \bmod n)+1},\dots,r_{((i+z-2) \bmod n)+1}\}$, and y' consecutive resources in $\mathcal{R}_2$ ending with resource $r_{n+((i+z'-2) \bmod n)+1}$. We provide an illustration of this game construction in Figure 7.2. Observe that $a^{\rm ne} = (a_1^{\rm ne},\dots,a_n^{\rm ne})$ satisfies the conditions for a Nash equilibrium,

$$\begin{aligned}
J_i(a^{\mathrm{ne}}) &= \eta x F^j(x) + (1-\eta)x' F^{j'}(x') \\
&= \eta[zF^j(x) + (y-z)F^j(x+1)] \\
&\qquad + (1-\eta)[z'F^{j'}(x') + (y'-z')F^{j'}(x'+1)] \\
&= J_i(a_i^{\mathrm{opt}}, a_{-i}^{\mathrm{ne}}),
\end{aligned}$$

which holds by (7.18). Then, by the above equality and (7.17),

$$\begin{aligned}
0 &= \sum_{i=1}^{n} J_i(a_i^{\mathrm{opt}}, a_{-i}^{\mathrm{ne}}) - \sum_{i=1}^{n} J_i(a^{\mathrm{ne}}) \\
&= \frac{1}{\nu^{\mathrm{opt}}}\Big[n \cdot \eta \left[\rho^{\mathrm{opt}} C^j(x) - C^j(y)\right] \\
&\qquad + n \cdot (1-\eta)\left[\rho^{\mathrm{opt}} C^{j'}(x') + C^{j'}(y')\right]\Big] \\
&= \frac{1}{\nu^{\mathrm{opt}}}\left[\rho^{\mathrm{opt}} C(a^{\mathrm{ne}}) - C(a^{\mathrm{opt}})\right],
\end{aligned}$$

where $a^{\mathrm{opt}} = (a^{\mathrm{opt}})_{i=1}^{n}$. Thus, $\mathrm{PoA}(G) = 1/\rho^{\mathrm{opt}}$. For Scenario (3), observe that $\rho^{\mathrm{opt}} = 0$, and so $1/\rho^{\mathrm{opt}}$ is unbounded. Recall that, in this scenario, there exist $j \in \{1, \ldots, m\}$ and $x \in \{1, \ldots, n\}$ such that $F^j(x) \leq 0$. We use the basis function pair $\{C^j, F^j\}$ to construct a game G with unbounded PoA. Consider a game instance with x agents and resource set $\mathcal{R} = \{r_1, r_2\}$, where $x \in \{1, \ldots, n\}$ is the value that minimizes the function $F(x)$, i.e., $F^j(x) = \min_{k \in \{1,\ldots,n\}} F^j(k) \leq 0$. Every agent $i \in \{1, \ldots, x\}$ has action set $\mathcal{A}_i = \{\{r_1\}, \{r_2\}\}$. The resource r_1 has cost function $C_r(k) = \eta C^j(k)$ and agent-cost function $F_r(k) = \eta F^j(k)$ for all k. Similarly, the resource r_2 has cost function $C_r(k) = (1-\eta)C^j(k)$ and agent-cost function $F_r(k) = (1-\eta)F(k)$. It is straightforward to verify that, for η approaching 0 from above, the allocation in which all agents select r_1 is an equilibrium and the PoA is unbounded. □

Chapter 8

Optimizing under the worst-case perspective

Considering system performance from the worst-case perspective is natural, as we would like to take a robust perspective on efficiency guarantees.

8.1 Optimizing the Price of Anarchy

8.1.1 Optimal taxation rule

First, we develop a methodology to compute optimal local tolling mechanisms through the solution of tractable linear programs. To ease the notation, we introduce the set of integer triplets $\mathcal{I} = \{(x, y, z) \in \mathbb{Z}^3_{\geq 0} \text{ s.t. } 1 \leq x + y + z \leq n \text{ and either } xyz = 0 \text{ or } x + y + z = n\}$, for given $n \in \mathbb{N}$.

Theorem 8.1.1. *A local mechanism minimizing the price of anarchy over congestion games with n agents, resource costs $\ell(x) = \sum_{j=1}^{m} \alpha_j b_j(x)$, $\alpha_j \geq 0$, and basis functions $\{b_1, \dots, b_m\}$ is given by*

$$T^{\mathrm{opt}}(\ell) = \sum_{j=1}^{m} \alpha_j \cdot \tau_j^{\mathrm{opt}}, \quad \textit{where} \;\; \tau_j^{\mathrm{opt}} : \{1, \dots, n\} \to \mathbb{R}, \quad \tau_j^{\mathrm{opt}}(x) = f_j^{\mathrm{opt}}(x) - b_j(x) \tag{8.1}$$

and $\rho_j^{\mathrm{opt}} \in \mathbb{R}$, $f_j^{\mathrm{opt}} : \{1, \dots, n\} \to \mathbb{R}$ *solve the following linear programs (one per each* b_j*)*

$$
\begin{aligned}
&\underset{f\in\mathbb{R}^n,\,\rho\in\mathbb{R}}{\text{maximize}}\ \rho \text{ subject to:}\\
&b_j(x+z)(x+z) - \rho b_j(x+y)(x+y) + f(x+y)y - f(x+y+1)z \geq 0 \quad \forall\,(x,y,z) \in \mathcal{I},
\end{aligned}
\tag{8.2}
$$

where we define $b_j(0) = f(0) = f(n+1) = 0$. *Correspondingly,* $\mathrm{PoA}(T^{\mathrm{opt}}) = \max_j\{1/\rho_j^{\mathrm{opt}}\}$.[1] *These results are tight for pure Nash equilibria, and extend to coarse correlated equilibria.*

The above statement contains two fundamental results. The first part of the statement shows that an optimal tolling mechanism applied to the function $\ell(x) = \sum_{j=1}^{m} \alpha_j b_j(x)$ can be obtained as the linear combination of $\tau_j^{\mathrm{opt}}(x)$, with the *same* coefficients α_j used to define ℓ. Complementary to this, the second part of the statement provides a practical technique to compute $\tau_j^{\mathrm{opt}}(x)$ for each of the basis $b_j(x)$ as the solution of a tractable linear program. Python/Matlab code to design optimal tolls can be found in [105].

We solved the latter linear programs for $n = 100$ and polynomials of maximum degree $1 \leq d \leq 6$. The corresponding results are displayed in Table 8.1 (one can show that these results hold identically for arbitrarily large n using similar techniques as in Section 9.3). In the case of $d = 1$, the optimal price of anarchy is approximately 2.012, matching that of untolled load balancing games on identical machines [94, 106]. We observe that, in this restricted setting, the price of anarchy cannot be improved *at all* through local tolling mechanisms. In fact, no matter what non-negative tolling mechanism we are given, we can always construct a load balancing game on identical machines with a price of anarchy no lower than 2.012.[2]

We conclude observing that the decomposition of resource costs as linear combination of basis functions is, strictly speaking, not required for Theorem 8.1.1 to hold. Nevertheless, pursuing this approach would require to solve a linear program for each function in $\mathcal{L}$, a task that becomes daunting when $\mathcal{L}$ contains infinitely many functions, e.g., in the case of polynomial

[1] If we require tolls to be non-negative, an optimal mechanism is as in (8.1), where we set $\tau_j^{\mathrm{opt}}(x) = f_j^{\mathrm{opt}}(x) \cdot \mathrm{PoA}^{\mathrm{opt}} - b_j(x)$.

[2] To do so, it is sufficient to utilize the instance in [106, Thm 3.4], where the the resource cost x used therein is replaced with $x + \tau(x)$. The Nash equilibrium and the optimal allocation will remain unchanged, yielding the same price of anarchy value.

d	No toll [86]	Global toll from [88, 72]	Optimal local toll (this work)	Optimal constant local toll (this work)	Marginal cost toll (this work)
1	2.50	2	2.012	2.15	3.00
2	9.58	5	5.101	5.33	13.00
3	41.54	15	15.551	18.36	57.36
4	267.64	52	55.452	89.41	391.00
5	1513.57	203	220.401	469.74	2124.21
6	12 345.20	877	967.533	3325.58	21 337.00

Table 8.1: Price of anarchy values for congestion games with resource costs of degree at most d. All results are tight for pure Nash and also hold for coarse correlated equilibria. The columns feature the price of anarchy with no tolls, with global tolls from [88, 72], with optimal local tolls, with optimal constant (i.e. congestion-independent) local tolls, and with marginal cost tolls, respectively. Columns four, five, and six, are composed of entirely novel results, except for the case of constant tolls with $d = 1$, which recovers [88]. Note that i) optimal tolls relying only on local information perform closely to optimal tolls designed using global information, with a difference in performance below 1% for $d = 1$; ii) congestion-independent tolls result in a price of anarchy that is comparable to that obtained using congestion-aware local tolls for polynomials of low degree. The code used to generate this table can be downloaded from [105].

congestion games. In this case, Theorem 8.1.1 allows to compute optimal tolls by solving *finitely* many linear programs.

8.1.2 Explicit solution and simplified linear program.

Next, we derive a simplified linear program as well as an analytical solution to the problem of designing optimal taxation rules. We do so under the assumption that all basis functions are positive, increasing, and convex in the discrete sense.[3]

Theorem 8.1.2. *Consider congestion games with n agents, where resource costs take the form $\ell(x) = \sum_{j=1}^{m} \alpha_j b_j(x)$, $\alpha_j \geq 0$, and the basis functions $b_j : \{1, \ldots, n\} \to \mathbb{R}$ are positive, convex, strictly increasing.*[4]

i) A tolling mechanism minimizing the price of anarchy is as in (8.1)*, where each* f_j^{opt} :

[3]We say that a function $f : \{1, \ldots, n\} \to \mathbb{R}$ is convex if $f(x+1) - f(x)$ is non-decreasing in its domain.

[4]The result also holds if convexity and strict increasingness of $b_j(x)$ are weakened to strict convexity of $b_j(x)x$ and $b_j(n) > b_j(n-1)$. One such example is that of $b_j(x) = \sqrt{x}$.

$\{1,\dots,n\} \to \mathbb{R}$ solves the following simplified linear program

$$
\begin{aligned}
(f_j^{\text{opt}}, \rho_j^{\text{opt}}) \in & \underset{f\in\mathbb{R}^n,\,\rho\in\mathbb{R}}{\mathit{maximize}}\ \rho \ \textit{subject to:} \\
& b_j(v)v - \rho b_j(u)u + f(u)u - f(u+1)v \geq 0, \\
& \qquad\qquad \forall u,v \in \{0,\dots,n\}\ \textit{s.t.}\ u+v \leq n, \\
& b_j(v)v - \rho b_j(u)u + f(u)(n-u) - f(u+1)(n-u) \geq 0, \\
& \qquad\qquad \forall u,v \in \{0,\dots,n\}\ \textit{s.t.}\ u+v > n,
\end{aligned}
\tag{8.3}
$$

with $f(0) = f(n+1) = 0$. The corresponding optimal price of anarchy is $\max_j\{1/\rho_j^{\text{opt}}\}$.

ii) An explicit expression for each f_j^{opt} is given by the following recursion, where $f_j^{\text{opt}}(1) = b_j(1)$,

$$
\begin{aligned}
f_j^{\text{opt}}(u+1) &= \min_{v\in\{1,\dots,n\}} \beta(u,v) f_j^{\text{opt}}(u) + \gamma(u,v) - \delta(u,v)\rho_j^{\text{opt}}, \\
\beta(u,v) = \frac{\min\{u,n-v\}}{\min\{v,n-u\}}, &\quad \gamma(u,v) = \frac{b(v)v}{\min\{v,n-u\}}, \quad \delta(u,v) = \frac{b(u)u}{\min\{v,n-u\}},
\end{aligned}
\tag{8.4}
$$

$$
\rho_j^{\text{opt}} = \min_{\substack{v_1,\dots,v_{n-1}\in\{1,\dots,n\},\\ v_n\in\{0,\dots,n\}}} \frac{(n-v_n)\left(\prod_{u=1}^{n-1}\beta_u b_j(1) + \sum_{u=1}^{n-2}\left(\prod_{i=u+1}^{n-1}\beta_i\right)\gamma_u + \gamma_{n-1}\right) + b(v_n)v_n}{(n-v_n)\left(\sum_{u=1}^{n-2}\left(\prod_{i=u+1}^{n-1}\beta_i\right)\delta_u + \delta_{n-1}\right) + b(n)n},
\tag{8.5}
$$

where we use the short-hand notation β_u instead of $\beta(u,v_u)$, and similarly for γ_u and δ_u.

Before delving into the proof, we observe that the key difficulty in designing optimal tolls resides in the expressions of ρ_j^{opt} arising from (8.5). Nevertheless, for any possible choice of $\bar{\rho}_j$ that approximates ρ_j^{opt} from below, i.e., $\bar{\rho}_j \leq \rho_j^{\text{opt}}$, one can directly utilize the recursion in (8.4) to design a valid tolling mechanism. The resulting price of anarchy would then amount to $\max_j\{1/\bar{\rho}_j\} > \max_j\{1/\rho_j^{\text{opt}}\}$. This follows from the ensuing proof.

8.2 Universal guarantees on the optimal Price of Anarchy

Recall that a guarantee on the Price of Anarchy for a taxation/utility rule translates directly to an approximation ratio of the underlying set of equilibria. Though several works provide tight bounds on the approximation ratio of polynomial-time centralized algorithms for the class of problems we consider (see, e.g., [107, 108, 109]), there is currently no result in the literature that establishes comparable bounds on the best achievable Price of Anarchy, aside from the general bound put forward in [110] that is provably inexact. Our main result in this section is an efficient technique for computing a utility mechanism that guarantees a Price of Anarchy of $1 - c/e$ in all resource allocation games with nonnegative, nondecreasing concave welfare functions with maximum curvature c.

Definition 8.2.1 (Curvature [111]). *The curvature of a nondecreasing concave function* $W : \mathbb{N} \to \mathbb{R}$ *is*

$$c = 1 - \frac{W(n) - W(n-1)}{W(1)}. \tag{8.6}$$

In the literature on submodular maximization, the curvature is commonly used to compactly parameterize broad classes of functions. The notion of curvature we consider was originally defined by Conforti et al. [111] in the context of general nondecreasing submodular set functions. In our specific setup, this reduces to the expression in Definition 8.2.1. Observe that all nondecreasing concave functions have curvature $c \in [0, 1]$. Thus, $c = 1$ can be considered in scenarios where the maximum curvature among functions in the set $\mathcal{W}$ is not known.

Theorem 8.2.1. *Let $\mathcal{G}$ denote the set of all resource allocation games with nonnegative, nondecreasing concave welfare functions with maximum curvature c. An optimal utility mechanism achieves* $\mathrm{PoA}(\mathcal{G}) = 1 - c/e$ *and can be computed efficiently.*

A significant consequence of this result is a universal guarantee that the best achievable Price of Anarchy is always greater than $1 - 1/e \approx 63.2\%$ for resource allocation games with nonnegative, nondecreasing concave welfare functions. Note that since $1 - 1/e$ is the optimal Price of Anarchy in general covering games [112], it cannot be further improved without more information about the underlying set of welfare functions. Our guarantee improves to $1 - c/e$ if the curvature c of the underlying set of welfare functions is known.

Observe that the result in Theorem 8.2.1 also implies that one can efficiently compute a "universal" utility mechanism, in that it would guarantee a Price of Anarchy greater than or equal to $1 - 1/e$ with respect to any game with nonnegative, nondecreasing concave welfare functions. This follows from the observation that $c \leq 1$ always holds for nondecreasing concave welfare functions. Of course, if more information is available about the underlying set of welfare functions (e.g., the maximum curvature), then this lower bound can be improved. In the case where the entire set of welfare functions $\mathcal{W}$ is known *a priori* and $|\mathcal{W}|$ is "small enough", then the optimal utility mechanism can be computed using existing methodologies (see, e.g., [113]).[5]

Consider the sets represented in Figure 8.1. From our reasoning, it holds that as the size of the set of welfare functions considered is reduced, the prices of anarchy of the corresponding optimal utility mechanisms increase. The set of games induced by welfares in the green ellipse, for example, coincides with the vehicle-target assignment problem, as described in [114], where $p_t = p \in [0, 1]$ for all $t \in \mathcal{T}$. Note that the welfare function W_t of each target $t \in \mathcal{T}$ in this problem is nonnegative, nondecreasing concave (i.e., the green ellipse is a subset of the dotted red box). Thus, we can immediately observe that the best achievable Price of Anarchy in the corresponding resource allocation game G satisfies

$$\mathrm{PoA}(G) \geq 1 - \frac{1}{e},$$

which is achieved by the universal utility mechanism from Theorem 8.2.1. Since there is only a single welfare function in this setting (ignoring uniform scalings) the optimal utility mechanism can be computed for a modest number of agents, as aforementioned.

In Figure 8.2, we plot the price of anarchy corresponding to the optimal utility mechanism within this setting (labelled "Optimal"), the price of anarchy achieved by the universal utility

[5] In this case, the optimal utility mechanism can be found as the solution of $|\mathcal{W}|$ linear programs with number of constraints that is quadratic in the maximum number of agents n, and $n + 1$ decision variables. For this reason, we say that the optimal utility mechanism can only be computed for modest values of $|\mathcal{W}|$ and n.

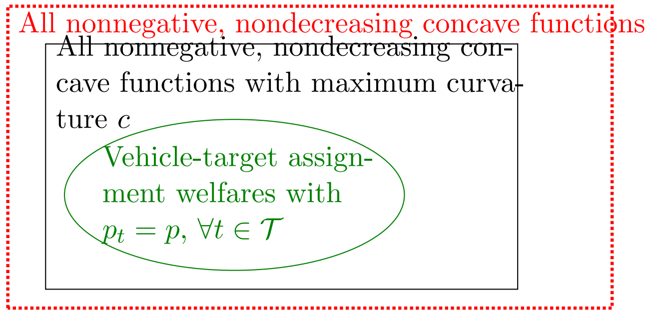

Figure 8.1: The set of games induced by the set of all nonnegative, nondecreasing concave functions contains the set of all nonnegative, nondecreasing concave functions with maximum curvature c, which in turn contains the set of all vehicle-target assignment problems with $p_t = p$.

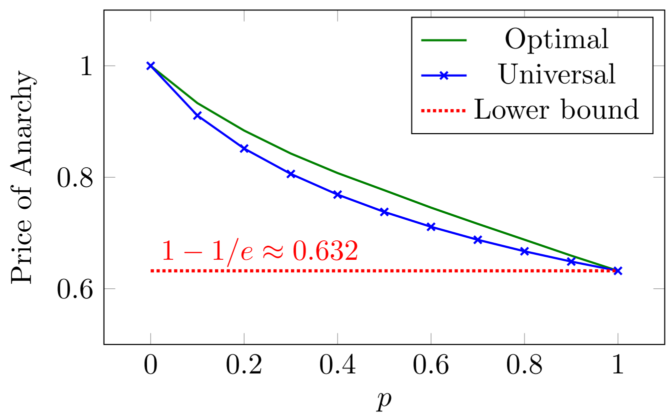

Figure 8.2: The price of anarchy of the universal utility mechanism obtained in this work and the optimal utility mechanism in the vehicle-target assignment problems with $p_t \in [0, p]$ for all $t \in \mathcal{T}$. Note that this utility mechanism is designed for the set of all nonnegative, nondecreasing concave welfare functions but its price of anarchy is close to the best achievable within this particular setting.

mechanism (labelled "Universal") and the $1 - 1/e$ lower bound from Theorem 8.2.1 (labelled "Lower bound"). As expected, the optimal utility mechanism corresponds with the best price of anarchy as it was designed specifically for the underlying welfare function. However, knowledge of the set of welfare functions corresponds with only a small increase in the price of anarchy; the price of anarchy achieved by the universal utility mechanism is surpisingly close to the best achievable by any mechanism for all values of $p \in [0, 1]$. Note that the universal utility mechanism is only guaranteed to achieve a price of anarchy of $1 - 1/e$.

8.3 Chapter proofs

8.3.1 Proofs from Section 8.1

Proof of Theorem 8.1.1

Proof. We divide the proof in two parts for ease of exposition.

Part 1. We show that any local mechanism minimizing the price of anarchy over all linear local mechanisms, does so also over all linear and non-linear local mechanisms. We let T^{opt} be a mechanism that minimizes the price of anarchy over all linear local mechanisms, i.e., over all T satisfying

$$T\left(\sum_{j=1}^{m} \alpha_j b_j\right) = \sum_{j=1}^{m} \alpha_j T(b_j),$$

for all $\alpha_j \geq 0$. We intend to show that $\text{PoA}(T^{\text{opt}}) \leq \text{PoA}(T)$ for any possible T (linear or non-linear). Towards this goal, assume, for a contradiction, that there exists a tolling mechanism $\hat{T}$ such that

$$\text{PoA}(T^{\text{opt}}) > \text{PoA}(\hat{T}). \tag{8.7}$$

Let $\mathcal{G}_b$ be the class of games in which any resource e can only utilize a resource cost $\ell_e \in \{b_1, \dots, b_m\}$. Since $\mathcal{G}_b \subset \mathcal{G}$, we have

$$\text{PoA}(\hat{T}) \geq \sup_{G \in \mathcal{G}_b} \sup_{a \in \text{NE}(G)} \frac{\text{SC}(a)}{\text{MinCost}(G)}. \tag{8.8}$$

Additionally, let $\mathcal{G}(\mathbb{Z}_{\geq 0}) \subset \mathcal{G}$ be the class of games with $\alpha_j \in \mathbb{Z}_{\geq 0}$ for all $j \in \{1, \dots, m\}$, for all resources in $\mathcal{R}$. Construct the mechanism $\bar{T}$ by "linearizing" the mechanism $\hat{T}$, i.e., as

$$\bar{T}(\ell) = \bar{T}\left(\sum_{j=1}^{m} \alpha_j b_j\right) = \sum_{j=1}^{m} \alpha_j \hat{T}(b_j).$$

We observe that the efficiency of any instance $G \in \mathcal{G}_b$ to which the tolling mechanism $\hat{T}$ is applied, coincides with that of an instance $G \in \mathcal{G}(\mathbb{Z}_{\geq 0})$ to which $\bar{T}$ is applied, and vice-versa.

Thus,

$$\sup_{G\in\mathcal{G}_b}\sup_{a\in\mathrm{NE}(G)}\frac{\mathrm{SC}(a)}{\mathrm{MinCost}(G)}=\sup_{G\in\mathcal{G}(\mathbb{Z}_{\geq 0})}\sup_{a\in\mathrm{NE}(G)}\frac{\mathrm{SC}(a)}{\mathrm{MinCost}(G)}=\mathrm{PoA}(\bar{T}),\tag{8.9}$$

where the last equality holds due to Lemma 8.3.1. Putting together (8.7), (8.8), and (8.9) gives

$$\mathrm{PoA}(T^{\mathrm{opt}})>\mathrm{PoA}(\bar{T}).\tag{8.10}$$

Since T^{opt} minimizes the price of anarchy over all linear mechanisms, and since $\bar{T}$ is linear by construction, it must be $\mathrm{PoA}(T^{\mathrm{opt}})\leq\mathrm{PoA}(\bar{T})$, a contradiction of (8.10). Thus, T^{opt} minimizes the price of anarchy over any mechanism.

Part 2. We will derive a linear program to design optimal linear mechanisms. Putting this together with the claim in Part 1 will conclude the proof. Towards this goal, we will prove that any mechanism of the form

$$T(\ell)=\sum_{j=1}^{m}\alpha_j\tau_j^{\mathrm{opt}}\quad\text{with}\quad\tau_j^{\mathrm{opt}}(x)=\lambda\cdot f_j^{\mathrm{opt}}(x)-b_j(x)\tag{8.11}$$

is optimal, regardless of the value of $\lambda\in\mathbb{R}_{>0}$. While this is slightly more general than needed, setting $\lambda=1$ will give the first claim. Additionally, setting $\lambda=\mathrm{PoA}^{\mathrm{opt}}$ will give the second claim as this choice will ensure non-negativity of the tolls.

Before turning to the proof, we recall a result from Chapter 7 that allows us to compute the price of anarchy for *given* linear tolling mechanism $T(\ell)=\sum_{j=1}^{m}\alpha_j\tau_j$. Upon defining $f_j(x)=b_j(x)+\tau_j(x)$ for all $1\leq x\leq n$ and $j\in\{1,\ldots,m\}$, the authors show that the price of anarchy of T computed over congestion games $\mathcal{G}$ is identical for pure Nash and coarse correlated equilibria and is given by $\mathrm{PoA}(T)=1/\rho^{\mathrm{opt}}$, where ρ^{opt} is the value of the following program

$$\begin{aligned}&\underset{\rho\in\mathbb{R},\nu\in\mathbb{R}_{\geq 0}}{\text{maximize}}\ \rho\ \text{subject to:}\\&b_j(x+z)(x+z)-\rho b_j(x+y)(x+y)+\nu[f_j(x+y)y-f_j(x+y+1)z]\geq 0,\\&\qquad\qquad\forall\,(x,y,z)\in\mathcal{I},\forall j\in\{1,\ldots,m\}.\end{aligned}\tag{8.12}$$

We also remark that, when all functions $\{f_j\}_{j=1}^m$ are non-decreasing, it is sufficient to only consider a reduced set of constraints, following a similar argument to that in [115, Cor. 1]. In this case, the linear program simplifies to

$$\begin{aligned} &\underset{\rho\in\mathbb{R},\nu\in\mathbb{R}_{\geq 0}}{\text{maximize}}\ \rho \text{ subject to:}\\ &b_j(v)v - \rho b_j(u)u + \nu[f_j(u)u - f_j(u+1)v] \geq 0,\\ &\qquad\qquad\forall u, v \in \{0,\dots,n\} \text{ s.t. } u+v \leq n, \quad \forall j \in \{1,\dots,m\},\\ &b_j(v)v - \rho b_j(u)u + \nu[f_j(u)(n-v) - f_j(u+1)(n-u)] \geq 0,\\ &\qquad\qquad\forall u, v \in \{0,\dots,n\} \text{ s.t. } u+v > n, \quad \forall j \in \{1,\dots,m\}. \end{aligned} \tag{8.13}$$

We now leverage (8.12) to prove that any mechanism in (8.11) is optimal, as required. Towards this goal, we begin by observing that the optimal price of anarchy obtained when the resource costs are generated using all the basis functions $\{b_1,\dots,b_m\}$ is no smaller than the optimal price of anarchy obtained when the resource costs are generated using a single basis function $\{b_j\}$ at a time (and therefore is no smaller than the highest of these optimal price of anarchy values). This follows readily since the former class of games is a superset of the latter. Additionally, observe that a set of tolls minimizing the price of anarchy over the games generated using a single basis function $\{b_j\}$ is precisely that in (8.11). This is because minimizing the price of anarchy amounts to designing f_j to maximize ρ in (8.12), i.e., to solving the following program

$$\begin{aligned} &\underset{f\in\mathbb{R}^n}{\text{maximize}}\ \max_{\rho\in\mathbb{R},\nu\in\mathbb{R}_{\geq 0}} \rho \text{ subject to:}\\ &b_j(x+z)(x+z) - \rho b_j(x+y)(x+y) + \nu[f(x+y)y - f(x+y+1)z] \geq 0, \quad \forall\,(x,y,z)\in\mathcal{I}, \end{aligned} \tag{8.14}$$

which can be equivalently written as

$$
\begin{aligned}
&\underset{\tilde{f}\in\mathbb{R}^n}{\text{maximize}}\ \max_{\rho\in\mathbb{R}}\ \rho \text{ subject to:}\\
&b_j(x+z)(x+z) - \rho b_j(x+y)(x+y) + \tilde{f}(x+y)y - \tilde{f}(x+y+1)z \geq 0, \quad \forall\ (x,y,z)\in\mathcal{I},
\end{aligned}
\tag{8.15}
$$

where we defined $\tilde{f} = \nu \cdot f$. While f_j^{opt} is defined in (8.2) precisely as the solution of this last program, resulting in a price of anarchy of $1/\rho_j^{\text{opt}}$, note that $\lambda \cdot f_j^{\text{opt}}$ is also a solution since its price of anarchy matches $1/\rho_j^{\text{opt}}$ (in fact, it can be computed using (8.12) for which $(\rho,\nu) = (\rho_j^{\text{opt}}, 1/\lambda)$ are feasible).

The above reasoning shows that the optimal price of anarchy for a game with resource costs generated by $\{b_1,\dots,b_m\}$ must be no smaller than $\max_j\{1/\rho_j^{\text{opt}}\}$. We now show that this holds with equality. Towards this goal, we note, thanks to (8.12), that utilizing tolls as in (8.7) for a game generated by $\{b_1,\dots,b_m\}$ results in a price of anarchy of precisely $\max_j\{1/\rho_j^{\text{opt}}\}$. This follows as $(\min_j\{\rho_j^{\text{opt}}\}, 1/\lambda)$ is feasible for this program for any choice of $\lambda > 0$. This proves, as requested, that any tolling mechanism defined in (8.11) is optimal.

We now verify that the choice $\lambda = \text{PoA}^{\text{opt}} = \max_j\{1/\rho_j^{\text{opt}}\}$ ensures positivity of the tolls, which is equivalent to $f_j^{\text{opt}}(x) - b_j(x)/\lambda \geq 0$ for all $x \in \{1,\dots,n\}$. This follows readily, as setting $x = z = 0$ in (8.2) results in the constraint $f(y) - \rho b_j(y) \geq 0$ for all $y \in \{1,\dots,n\}$. Since f_j^{opt} and ρ_j^{opt} must be feasible for this constraint, we have $f_j^{\text{opt}}(y) - \rho_j^{\text{opt}} b_j(y) \geq 0$. One concludes observing that $f_j^{\text{opt}}(y) - b_j(y)/\lambda \geq f_j^{\text{opt}}(y) - \rho_j^{\text{opt}} b_j(y) \geq 0$, since $\lambda \geq 1/\rho_j^{\text{opt}}$. We conclude remarking that all results hold for both Nash and coarse correlated equilibria, as they were derived from (8.12). □

Lemma 8.3.1. *Consider the class of congestion games $\mathcal{G}$. For any linear tolling mechanism T, it is*

$$
\text{PoA}(T) = \sup_{G\in\mathcal{G}(\mathbb{Z}_{\geq 0})}\ \sup_{a\in\text{NE}(G)} \frac{\text{SC}(a)}{\text{MinCost}(G)},
$$

where $\mathcal{G}(\mathbb{Z}_{\geq 0}) \subset \mathcal{G}$ is the subclass of games with $\alpha_j \in \mathbb{Z}_{\geq 0}$ for all $j \in \{1,\dots,m\}$, for all resources in $\mathcal{R}$.

Proof. We divide the proof in two steps. First, we show that

$$\mathrm{PoA}(T) = \sup_{G\in\mathcal{G}(\mathbb{Q}_{\geq 0})} \sup_{a\in\mathrm{NE}(G)} \frac{\mathrm{SC}(a)}{\mathrm{MinCost}(G)}, \tag{8.16}$$

where $\mathcal{G}(\mathbb{Q}_{\geq 0}) \subset \mathcal{G}$ is the subclass of games with $\alpha_j \in \mathbb{Q}_{\geq 0}$ for all $j \in \{1,\dots,m\}$, for all resources in $\mathcal{R}$. Towards this goal, observe that (8.16) holds trivially with $\geq$ in place of the equality sign, as $\mathbb{R}_{\geq 0} \supset \mathbb{Q}_{\geq 0}$. To show that the converse inequality also holds, observe that the price of anarchy of a given linear mechanisms T (computed over all meaningful instances where $\mathrm{SC}(a^{\mathrm{ne}}) > 0$) can be computed utilizing the linear program reported in (8.12). By strong duality, we have $\mathrm{PoA}(T) = 1/C^\star$, where $C^\star$ is the value of the dual program of (8.12), i.e.,

$$C^\star = \underset{\theta(x,y,z,j)}{\text{minimize}} \sum_{x,y,z,j} b_j(x+z)(x+z)\theta(x,y,z,j) \tag{8.17}$$

$$\sum_{x,y,z,j} \left[f_j(x+y)y - f_j(x+y+1)z\right]\theta(x,y,z,j) \leq 0, \tag{8.18}$$

$$\sum_{x,y,z,j} b_j(x+y)(x+y)\theta(x,y,z,j) = 1, \tag{8.19}$$

$$\theta(x,y,z,j) \geq 0, \forall (x,y,z,j) \in \mathcal{I}, \tag{8.20}$$

where we define $b_j(0) = f_j(0) = f_j(n+1) = 0$ for convenience, $\mathcal{I} = \{(x,y,z,j) \in \mathbb{Z}^4_{\geq 0} \text{ s.t. } 1 \leq x+y+z \leq n,\ 1 \leq j \leq m\}$, and the minimum is intended over the entire tuple $\{\theta(x,y,z,j)\}_{(x,y,z,j)\in\mathcal{I}}$. Let $\{\theta^\star(x,y,z,j)\}_{(x,y,z,j)\in\mathcal{I}}$ denote an optimal solution (which exists, due to the non-emptiness and boundedness of the constraint set, which can be proven using the same argument in [115, Thm. 2]). If all $\theta^\star(x,y,z,j)$ are rational, then consider the game G defined as follows. For every $i \in \{1,\dots,n\}$ and for every $(x,y,z,j) \in \mathcal{I}$, we create a resource identified with $e(x,y,z,j,i)$, and assign to it the resource cost $\alpha_j b_j$, where $\alpha_j = \theta^\star(x,y,z,j)/n$. The game G features n players, where player $p \in \{1,\dots,n\}$ can either select the resources in the allocation a_p^{opt} or in

a_p^{ne}, defined by

$$a_p^{\mathrm{opt}} = \cup_{i=1}^{n} \cup_{j=1}^{m} \{e(x,y,z,j,i) \,:\, x+y \geq 1 + ((i-p) \bmod n)\},$$
$$a_p^{\mathrm{ne}} = \cup_{i=1}^{n} \cup_{j=1}^{m} \{e(x,y,z,j,i) \,:\, x+z \geq 1 + ((i-p+z) \bmod n)\}.$$

Note that the above construction is an extension of that appearing in [115] to the case of multiple basis functions. Since G has

$$\sup_{a \in \mathrm{NE}(G)} \mathrm{SC}(a) = \sum_{x,y,z,j} b_j(x+y)(x+y)\theta^\star(x,y,z,j) = 1,$$
$$\mathrm{MinCost}(G) \leq \sum_{x,y,z,j} b_j(x+z)(x+z)\theta^\star(x,y,z,j) = C^\star,$$

(see [115, Thm. 2] for this), its price of anarchy is no smaller than $1/C^\star$. Observe that G features only non-negative rational resource costs' coefficients (i.e., $G \in \mathcal{G}(\mathbb{Q}_{\geq 0})$), therefore (8.16) follows readily.

If at least one entry in the tuple $\{\theta^\star(x,y,z,j)\}_{(x,y,z,j)\in\mathcal{I}}$ is not rational, we will prove the existence of a sequence of games $G^k \in \mathcal{G}(\mathbb{Q}_{\geq 0})$ whose worst-case efficiency converges to $\mathrm{PoA}(T)$ as $k \to \infty$. This would imply that (8.16) holds with $\leq$ in place of the equality sign, concluding the proof. To do so, let us consider the set

$$S = \{\{\theta(x,y,z,j)\}_{(x,y,z,j)\in\mathcal{I}} \ \text{ s.t. } (8.18), \text{ and } (8.20) \text{ hold}\}.$$

Observe that S is non-empty, and that for any tuple belonging to S, we can find a sequence of non-negative rational tuples $\{\{\theta^k(x,y,z,j)\}_{(x,y,z,j)\in\mathcal{I}}\}_{k=1}^{\infty}$ (i.e., $\theta^k(x,y,z,j) \in \mathbb{Q}_{\geq 0}$ for all x,y,z,j and k), that converges to it.

Let $\{\{\theta^k(x,y,z,j)\}_{(x,y,z,j)\in\mathcal{I}}\}_{k=1}^{\infty}$ be a sequence of tuples converging to $\theta^\star(x,y,z,j)$ such that $(x,y,z,j) \in \mathcal{I}$, which belongs to S. For each tuple $\{\theta^k(x,y,z,j)\}_{(x,y,z,j)\in\mathcal{I}}$ in the sequence, define the game G^k following the same construction introduced above with $\theta^k(x,y,z,j)$ in place of $\theta^\star(x,y,z,j)$. Following the same reasoning as above, it is $\sup_{a\in\mathrm{NE}(G^k)} \mathrm{SC}(a) = \sum_{x,y,z,j} b_j(x+$

$y)(x+y)\theta^k(x,y,z,j)$, and $\mathrm{MinCost}(G^k) \leq \sum_{x,y,z,j} b_j(x+z)(x+z)\theta^k(x,y,z,j)$. Therefore

$$\mathrm{PoA}^k = \sup_{a\in\mathrm{NE}(G^k)} \frac{\mathrm{SC}(a)}{\mathrm{MinCost}(G^k)} \geq \frac{\sum_{x,y,z,j} b_j(x+y)(x+y)\theta^k(x,y,z,j)}{\sum_{x,y,z,j} b_j(x+z)(x+z)\theta^k(x,y,z,j)},$$

from which we conclude that

$$\begin{aligned}\lim_{k\to\infty} \mathrm{PoA}^k &\geq \lim_{k\to\infty} \frac{\sum_{x,y,z,j} b_j(x+y)(x+y)\theta^k(x,y,z,j)}{\sum_{x,y,z,j} b_j(x+z)(x+z)\theta^k(x,y,z,j)} \\ &= \frac{\sum_{x,y,z,j} b_j(x+y)(x+y)\theta^\star(x,y,z,j)}{\sum_{x,y,z,j} b_j(x+z)(x+z)\theta^\star(x,y,z,j)} = \frac{1}{C^\star},\end{aligned}$$

as $\theta^k(x,y,z,j) \to \theta^\star(x,y,z,j)$ for $k\to\infty$. This completes the first step.

The second and final step consist in showing that

$$\sup_{G\in\mathcal{G}(\mathbb{Q}_{\geq 0})} \sup_{a\in\mathrm{NE}(G)} \frac{\mathrm{SC}(a)}{\mathrm{MinCost}(G)} = \sup_{G\in\mathcal{G}(\mathbb{Z}_{\geq 0})} \sup_{a\in\mathrm{NE}(G)} \frac{\mathrm{SC}(a)}{\mathrm{MinCost}(G)}.$$

Towards this goal, for any given game from the above-defined sequence $G^k \in \mathcal{G}(\mathbb{Q}_{\geq 0})$, let d_{G^k} denote the lowest common denominator among the resource cost coefficients α_j, across all the resources of the game. Define $\hat{\alpha}_j = \alpha_j \cdot d_{G^k} \in \mathbb{Z}_{\geq 0}$ for all $j \in \{1,\dots,m\}$, for all resources in $\mathcal{R}$. Since the tolling mechanisms T is linear by assumption, the equilibrium conditions are independent to uniform scaling of the resource costs and tolls by the coefficient d_{G^k}. Therefore any game in the sequence G^k with tolling mechanism T and resource cost coefficients $\{\alpha_j\}_{j=1}^m$ has the same worst-case equilibrium efficiency as a game $\hat{G}^k$ which is identical to G^k except that it has resource cost coefficients $\{\hat{\alpha}_j\}_{j=1}^m$. Observing that $\hat{G}^k$ belongs to $\mathcal{G}(\mathbb{Z}_{\geq 0})$ concludes the proof. □

Proof of Theorem 8.1.2. Before delving into the proof, we observe that the key difficulty in designing optimal tolls resides in the expressions of ρ_j^{opt} arising from (8.5). Nevertheless, for any possible choice of $\bar{\rho}_j$ that approximates ρ_j^{opt} from below, i.e., $\bar{\rho}_j \leq \rho_j^{\mathrm{opt}}$, one can directly utilize the recursion in (8.4) to design a valid tolling mechanism. The resulting price of anarchy would then amount to $\max_j\{1/\bar{\rho}_j\} > \max_j\{1/\rho_j^{\mathrm{opt}}\}$. This follows from the ensuing proof.

Proof. As shown in Theorem 8.1.1, computing an optimal tolling mechanism amounts to utilizing (8.1), where each τ_j^{opt} has been designed through the solution of the program in (8.2). In light of this, we prove the theorem as follows: first, we consider a simplified linear program, where only a subset of the constraints enforced in (8.2) are considered. Second, we show that a solution of this simplified program is given by $(\rho_j^{\mathrm{opt}}, f_j^{\mathrm{opt}})$ as defined above. Third, we show that f_j^{opt} is non-decreasing, thus ensuring that $(\rho_j^{\mathrm{opt}}, f_j^{\mathrm{opt}})$ is also feasible for the original over constrained program in (8.2). From this we conclude that $(\rho_j^{\mathrm{opt}}, f_j^{\mathrm{opt}})$ must also be a solution of (8.2), i.e., the second claim in the Theorem. We conclude with some cosmetics, and transform the simplified linear program whose solution is given by $(\rho_j^{\mathrm{opt}}, f_j^{\mathrm{opt}})$ in (8.3), thus obtaining the first claim. Throughout the proof, we drop the index j from b_j as the proof can be repeated for each basis separately.

Simplified linear program. We begin by rewriting the program in (8.2), where instead of the indices (x, y, z), we use the corresponding indices (u, v, x) defined as $u = x+y$, $v = x+z$. The constraint indexed by (u, v, x) reads as $b(v)v-\rho b(u)u+f(u)(u-x)-f(u+1)(v-x) \geq 0$. We now consider only the constraints where x is set to $x = \min\{0, u+v-n\}$, and $u, v \in \{0, \dots, n\}$, Such constraints read as $b(v)v - \rho b(u)u + \min\{u, n-v\}f(u) - \min\{v, n-u\}f(u+1) \geq 0$.[6] Finally, we exclude the constraints with $v = 0$, $u \in \{1, \dots, n-1\}$ and obtain the following *simplified linear program*

$$\begin{aligned}
&\underset{f\in\mathbb{R}^n,\,\rho\in\mathbb{R}}{\text{maximize}}\ \rho \text{ subject to:}\\
&b(v)v - \rho b(u)u + \min\{u, n-v\}f(u) - \min\{v, n-u\}f(u+1) \geq 0, \\
&\qquad\qquad\qquad\qquad \forall (u,v) \in \{0,\dots,n\} \times \{1,\dots,n\} \cup (n,0).
\end{aligned} \tag{8.21}$$

Proof that $(\rho^{\mathrm{opt}}, f^{\mathrm{opt}})$ solve (8.21). Towards the stated goal, we begin by observing that $(\rho^{\mathrm{opt}}, f^{\mathrm{opt}})$ is feasible by construction. For $u = 0$ this follows as the tightest constraints

[6]Note that considering all these constraints with $u, v \in \{0, \dots, n\}$ results precisely in (8.3). To see this, simply distinguish the cases based on whether $u + v \leq n$ or $u + v > n$.

in (8.21) read as $f^{\mathrm{opt}}(1) \geq b(1)$ and we selected $f^{\mathrm{opt}}(1) = b(1)$. Feasibility is immediate to verify for $u \in \{1, \dots, n-1\}$, $v \in \{1, \dots, n\}$ as applying its definition gives $f^{\mathrm{opt}}(u+1) \leq \beta(u,v) f^{\mathrm{opt}}(u) + \gamma(u,v) - \delta(u,v)\rho^{\mathrm{opt}}$. Using the expressions of β, γ, δ, and rearranging gives exactly the constraint (u,v) in (8.21). The only element of difficulty consists in showing that also the constraints with $u = n$, $v \in \{0, \dots, n\}$ are satisfied. Towards this goal, we observe that utilizing the recursive definition of f^{opt} we obtain an expression for $f^{\mathrm{opt}}(n)$ as a function of ρ^{opt} with a nested succession of minimizations, which can be jointly extracted as follows

$$f^{\mathrm{opt}}(n) = \min_{v_{n-1}} \left\{ \dots + \min_{v_{n-2}} \left\{ \dots + \min_{v_1} \{ \dots \} \right\} \right\} = \min_{v_{n-1}} \min_{v_{n-2}} \dots \min_{v_1} \{ \dots \} .$$

This holds as $f^{\mathrm{opt}}(u+1) = \min_{v_u} [\beta_u \min_{v_{u-1}} (\beta_{u-1} f^{\mathrm{opt}}(u-1) - \delta_{u-1}\rho^{\mathrm{opt}} + \gamma_{u-1}) - \delta_u \rho^{\mathrm{opt}} + \gamma_u]$, and since $\beta_u \geq 0$, the latter simplifies to $f^{\mathrm{opt}}(u+1) = \min_{v_u} \min_{v_{u-1}} \beta_u \beta_{u-1} f^{\mathrm{opt}}(u-1) - (\beta_u \delta_{u-1} + \delta_u)\rho^{\mathrm{opt}} + \beta_u \gamma_{u-1} + \gamma_u$. Repeating the argument recursively gives the desired expression. Hence,

$$\begin{aligned} f^{\mathrm{opt}}(n) &= \min_{(v_1,\dots,v_{n-1}) \in \{1,\dots,n\}^{n-1}} \prod_{u=1}^{n-1} \beta_u b_j(1) + \sum_{u=1}^{n-2} \left(\prod_{i=u+1}^{n-1} \beta_i \right) (\gamma_u - \delta_u \rho^{\mathrm{opt}}) + (\gamma_{n-1} - \delta_{n-1}\rho^{\mathrm{opt}}) \\ &\doteq \min_{(v_1,\dots,v_{n-1}) \in \{1,\dots,n\}^{n-1}} q(v_1, \dots, v_{n-1}; \rho^{\mathrm{opt}}), \end{aligned}$$

where we implicitly defined $q(v_1, \dots, v_{n-1}; \rho^{\mathrm{opt}})$. The constraints we intend to verify read as $b(v)v - \rho b(n)n + (n-v) f^{\mathrm{opt}}(n) \geq 0$ for all $v \in \{0, \dots, n\}$, and can be equivalently written as $\min_{v_n \in \{0,\dots,n\}} [b(v_n)v_n - \rho b(n)n + (n-v_n) f^{\mathrm{opt}}(n)] \geq 0$. We substitute the resulting expression of $f^{\mathrm{opt}}(n)$, extract the minimization over v_n as in the above, and are therefore left with $\min_{(v_1,\dots,v_{n-1},v_n) \in \{1,\dots,n\}^{n-1} \times \{0,\dots,n\}} [b(v_n)v_n - \rho b(n)n + (n-v) q(v_1, \dots, v_{n-1}; \rho^{\mathrm{opt}})] \geq 0$, which holds if and only if $b(v_n)v_n - \rho b(n)n + (n-v_n) q(v_1, \dots, v_{n-1}; \rho^{\mathrm{opt}}) \geq 0$ for all possible tuples $(v_1, \dots, v_n)$. Rearranging these constraints and solving for ρ^{opt} will result in a set of inequalities on ρ^{opt} (one inequality for each tuple). Our choice of ρ^{opt} in (8.5) is precisely obtained by turning the most binding of these into an equality. This ensures that $(\rho^{\mathrm{opt}}, f^{\mathrm{opt}})$ are feasible also when $u = n$.

We now prove, by contradiction, that $(\rho^{\mathrm{opt}}, f^{\mathrm{opt}})$ is optimal. To do so, we assume that there exists $\hat{f}$, that is feasible and achieves a higher value $\hat{\rho} > \rho^{\text{opt}}$. Since $(\hat{f}, \hat{\rho})$ is feasible, using the constraint with $u = 0$, $v = 1$, we have $\hat{f}(1) \leq b(1) = f^{\text{opt}}(1)$. Observing that $\min\{v, n-u\} > 0$ due to $v > 0$, $u < n$ and leveraging the constraints with $u = 1$ as well as the corresponding specific choice of $v = v_1^*$ (for given $u \in \{1, \ldots, n-1\}$, we let v_u^* be an index $v \in \{1, \ldots, n\}$ where the minimum in (8.4) is attained), it must be that $\hat{f}(2)$ satisfies

$$\hat{f}(2) \leq \frac{b(v_1^*)v_1^* - \hat{\rho}b(1) + \min\{1, n-v_1^*\}\hat{f}(1)}{\min\{v_1^*, n-1\}} < \frac{b(v_1^*)v_1^* - \rho^{\text{opt}}b(1) + \min\{1, n-v_1^*\}f^{\text{opt}}(1)}{\min\{v_1^*, n-1\}}$$

which is equal to $f^{\text{opt}}(2)$. Here the first inequality follows by feasibility of $\hat{f}$, the second is due to $\hat{\rho} > \rho^{\text{opt}}$ and $\hat{f}(1) \leq f^{\text{opt}}(1)$. The final equality follows due to the definition of $f^{\text{opt}}(2)$. Hence we have shown that $\hat{f}(2) < f^{\text{opt}}(2)$. Noting that the only information we used to move from level u to $u+1$ is that $\hat{\rho} > \rho^{\text{opt}}$ and $\hat{f}(u) \leq f^{\text{opt}}(u)$, one can apply this argument recursively up until $u = n-1$, and thus obtain $\hat{f}(n) < f^{\text{opt}}(n)$. Nevertheless, leveraging the constraints with $u = n$ and $v = v_n^*$ gives $b(v_n^*)v_n^* - \hat{\rho}b(n)n + (n - v_n^*)\hat{f}(n) \geq 0$, or equivalently $\hat{\rho} \leq (b(v_n^*)v_n^* + (n - v_n^*)\hat{f}(n))/(b(n)n)$. Thus

$$\hat{\rho} \leq \frac{b(v_n^*)v_n^* + (n - v_n^*)\hat{f}(n)}{b(n)n} \leq \frac{b(v_n^*)v_n^* + (n - v_n^*)f^{\text{opt}}(n)}{b(n)n} = \rho^{\text{opt}},$$

where we used the fact that $n - v_n^* \geq 0$ and $\hat{f}(n) < f^{\text{opt}}(n)$. Note that $\hat{\rho} \leq \rho^{\text{opt}}$ contradicts the assumption $\hat{\rho} > \rho^{\text{opt}}$, thus concluding this part of the proof.

Proof that f^{opt} is non-decreasing. By contradiction, let us assume f^{opt} is decreasing at some index. Lemma 8.3.2 in the Appendix shows that, if this is the case, then f^{opt} continues to decrease, so that $f^{\text{opt}}(n) \leq f^{\text{opt}}(n-1)$. Note that it must be $f^{\text{opt}}(n) > 0$, as if it were $f^{\text{opt}}(n) \leq 0$, then by definition of ρ^{opt} we would have

$$\rho^{\text{opt}} = \min_{v \in \{0, \ldots, n\}} \frac{b(v)v + (n-v)f^{\text{opt}}(n)}{nb(n)} = \frac{0 + f^{\text{opt}}(n)}{b(n)} \leq 0,$$

since the minimum is attained at the lowest feasible v due to $b(v)v$ and $-vf^{\text{opt}}(n)$ non-decreasing and increasing, respectively. This is a contradiction as the price of anarchy is bounded already in the un-tolled setup.[7] It must therefore be that the price of anarchy is bounded also when optimal tolls are used. Additionally, as we have removed a number of constraints from the linear program, the corresponding price of anarchy will be even lower. Therefore it must be that $1/\rho^{\rm opt}$ is non-negative and bounded, so that $\rho^{\rm opt} > 0$ contradicting the last equation.

Thus, in the following we proceed with the case of $f^{\rm opt}(n) > 0$. It must be that

$$\begin{aligned}\rho^{\rm opt} &= \min_{v\in\{0,\dots,n\}} \frac{b(v)v + (n-v)f^{\rm opt}(n)}{nb(n)} \le \min_{v\in\{1,\dots,n\}} \frac{b(v)v + (n-v)f^{\rm opt}(n)}{nb(n)}\\ &\le \min_{v\in\{1,\dots,n\}} \frac{b(v)v + (n-v)f^{\rm opt}(n-1)}{nb(n)},\end{aligned}$$

where the first inequality holds as we are restricting the domain of minimization, the second because $f^{\rm opt}(n) \le f^{\rm opt}(n-1)$ and $n - v \ge 0$. Let us observe that $f^{\rm opt}(n)$ is defined as $f(n) = \min_{v\in\{1,\dots,n\}}[b(v)v + (n-v)f^{\rm opt}(n-1)] - \rho^{\rm opt}(n-1)b(n-1)$. Substituting $\min_{v\in\{1,\dots,n\}}[b(v)v + (n-v)f^{\rm opt}(n-1)] = f^{\rm opt}(n) + \rho^{\rm opt}(n-1)b(n-1)$ in the former bound on $\rho^{\rm opt}$, we get

$$\rho^{\rm opt} \le \frac{f^{\rm opt}(n) + \rho^{\rm opt}(n-1)b(n-1)}{nb(n)} \qquad\implies\qquad \rho^{\rm opt} \le \frac{f^{\rm opt}(n)}{nb(n) - (n-1)b(n-1)}.$$

We want to prove that this gives rise to a contradiction. To do so, we will show that

$$\frac{f^{\rm opt}(n)}{nb(n) - (n-1)b(n-1)} < \min_{v\in\{0,\dots,n\}} \frac{b(v)v + (n-v)f^{\rm opt}(n)}{nb(n)}. \tag{8.22}$$

As a matter of fact, if the latter inequality holds true, the proof is immediately concluded as

$$\begin{aligned}&\rho^{\rm opt} \le \frac{f^{\rm opt}(n)}{nb(n) - (n-1)b(n-1)} < \min_{v\in\{0,\dots,n\}} \frac{b(v)v + (n-v)f^{\rm opt}(n)}{nb(n)} = \rho^{\rm opt}\\ &\implies \rho^{\rm opt} < \rho^{\rm opt},\end{aligned}$$

where the first inequality has been shown above, the second is what remains to be proved, and

[7] To see this, consider the linear program used to determine the price of anarchy in the un-tolled case, i.e., (8.13) where we set $f_j(x) = b_j(x)$. When $\nu = 1$, it is always possible to find $\rho > 0$, so that the corresponding price of anarchy is bounded.

the latter equality is by definition. Therefore, we are left to show (8.22), which holds if we can show that $\forall v \in \{0, \ldots, n\}$ it is

$$g(v) \doteq \frac{h(v) + (n-v)f^{\mathrm{opt}}(n)}{h(n)} - \frac{f^{\mathrm{opt}}(n)}{h(n) - h(n-1)} > 0,$$

where $h : \mathbb{R} \to \mathbb{R}_{\geq 0}$ is a function such that $h(v) = b(v)v$ for $v \in \{0, \ldots, n\}$. We choose h to be continuously differentiable, strictly increasing, and strictly convex; one such function always exists.[8] We first consider the point $v = 0$. Observe that $g(0) > 0$ when $n > 1$ as

$$g(0) = \frac{f^{\mathrm{opt}}(n)}{b(n)} - \frac{f^{\mathrm{opt}}(n)}{nb(n) - (n-1)b(n-1)} > 0 \quad \iff \quad f^{\mathrm{opt}}(n)[(n-1)b(n) - (n-1)b(n-1)] > 0,$$

which holds as $f^{\mathrm{opt}}(n) > 0$, $n > 1$, and $b(n) > b(n-1)$ strictly.

If $g'(v) \geq 0$ at $v = 0$, the proof is complete as g is convex and due to $g'(0) \geq 0$ it is non-decreasing for any $v \geq 0$ so that the constraint will be satisfied for all $v \geq 0$.

If this is not the case, then $g'(0) < 0$, which we consider now. Note that, at the point $v = n-1$, the derivative $g'(n-1) = [h'(n-1) - f^{\mathrm{opt}}(n)]/h(n)$ satisfies

$$h(n)g'(n-1) = h'(n-1) - f^{\mathrm{opt}}(n) \geq h'(n-1) - (h(n-1) - h(n-2)) \geq 0$$

where the last inequality is due to convexity, while the first inequality holds as $f^{\mathrm{opt}}(n) \leq h(n-1) - h(n-2)$ thanks to Lemma 8.3.2 and $n \geq 2$.[9] Therefore since $g'(0) < 0$, $g'(n-1) \geq 0$ and g convex, there must exist an unconstrained minimizer $v^\star \in (0, n-1]$. We will guarantee that $g(v^\star) > 0$ so that for any (real and thus integer) $v \in [0, n]$ it is $g(v) > 0$. The unconstrained minimizer satisfies $f^{\mathrm{opt}}(n) = h'(v^\star)$, which we substitute, and are thus left with proving the

[8] Observe that the function $b(v)v$ is positive, strictly increasing, and strictly convex in the discrete sense in its domain due to the assumptions.

[9] In fact, either n is the first index starting from which f^{opt} decreases (i.e. $f^{\mathrm{opt}}(n) < f^{\mathrm{opt}}(n-1)$) in which case $f^{\mathrm{opt}}(n) \leq \rho^{\mathrm{opt}}[b(n-1)(n-1) - b(n-2)(n-2)] \leq b(n-1)(n-1) - b(n-2)(n-2)$ due to $\rho^{\mathrm{opt}} \leq 1$, or the function starts decreasing at a $u+1 < n$ in which case Lemma 8.3.2 also shows that

$$f^{\mathrm{opt}}(n) \leq \cdots \leq f^{\mathrm{opt}}(u+1) \leq \rho^{\mathrm{opt}}[b(u)u - b(u-1)(u-1)] \leq b(u)u - b(u-1)(u-1) \leq b(n-1)(n-1) - b(n-2)(n-2),$$

where the inequalities hold due to $\rho^{\mathrm{opt}} \leq 1$ and the convexity of $b(u)u$.

final inequality
$$\frac{h(v^\star)+(n-v)h'(v^\star)}{h(n)}-\frac{h'(v^\star)}{h(n)-h(n-1)}>0,$$

which is equivalent to

$$[h(n)-h(n-1)]h(v^\star)>h'(v^\star)[(n-v^\star)(h(n-1)-h(n))+h(n)],$$

where we recall $0<v^\star\leq n-1$. As the left hand side is positive due to h increasing and $v^\star>0$, the inequality holds trivially if the right hand side is less or equal to zero, i.e., if $h(n)\leq(n-v^\star)(h(n)-h(n-1))$. In the other case, when $(n-v^\star)(h(n-1)-h(n))+h(n)>0$, we leverage the fact that $h'(v^\star)<(h(n)-h(v^\star))/(n-v^\star)$ by strict convexity of $h(x)$ in $x=v^\star>0$, so that

$$\begin{aligned}h'(v^\star)[(n-v^\star)(h(n-1)-h(n))+h(n)]&<\frac{h(n)-h(v^\star)}{n-v^\star}[(n-v^\star)(h(n-1)-h(n))+h(n)]\\&=\frac{h(n)}{n-v^\star}[h(n)-h(v^\star)]+[h(n)-h(n-1)][h(v^\star)-h(n)]\leq[h(n)-h(n-1)]h(v^\star),\end{aligned}$$

where the last inequality follows since $[h(n)-h(n-1)][h(v^\star)-h(n)]\leq 0$ and from $\frac{h(n)-h(v^\star)}{n-v^\star}\leq h(n)-h(n-1)$, which holds for $0<v^\star\leq n-1$ by convexity. This concludes this part of the proof.

Proof that $(\rho^{\text{opt}},f^{\text{opt}})$ **is feasible also for** (8.2) **and final cosmetics.** Recall from the first part of the proof that the constraints in (8.2) can be equivalently written as $b(v)v-\rho b(u)u+f(u)(u-x)-f(u+1)(v-x)\geq 0$. Since f^{opt} is non-decreasing, following the argument in [115, Cor. 1] one verifies that the tightest constraints are obtained when $x=\min\{0,u+v-n\}$. These constraints are already included in our simplified program of (8.21), with the exception of those with $v=0$ and $u\in\{0,\dots,n-1\}$ which we have removed. To show that also these hold, we note that the constraint with $v=0$ reads as $uf^{\text{opt}}(u)\geq\rho^{\text{opt}}ub(u)$, and is trivially satisfied for $u=0$. We now show that also the constraints with $v=0$, $u>0$ hold. To do so, we consider the constraint corresponding to $v = 1$

$$b(1)1 - \rho b(u)u + uf^{\text{opt}}(u) - f^{\text{opt}}(u+1) \geq 0.$$

Since f^{opt} is non-decreasing as shown in previous point then $f^{\text{opt}}(u+1) \geq f^{\text{opt}}(1) = b(1)$. Hence,

$$0 \leq b(1)1 - \rho b(u)u + uf^{\text{opt}}(u) - f^{\text{opt}}(u+1) \leq b(1)1 - \rho^{\text{opt}} b(u)u + uf^{\text{opt}}(u) - b(1).$$

Thus, from the left and right hand sides we obtain the desired result $uf^{\text{opt}}(u) \geq \rho^{\text{opt}} ub(u)$.

We conclude with some cosmetics: the simplified linear program in (8.21) is almost identical to that in (8.3), except for the constraints with $v = 0$ and $u \in \{0, \ldots, n-1\}$, which we have removed in (8.21). Nevertheless, we have just verified that an optimal solution does satisfy these constraints too. Hence, we simply add them back to obtain (8.3). □

Lemma 8.3.2. *Let $b : \mathbb{N} \to \mathbb{R}_{\geq 0}$ be a nondecreasing, convex function, and let $0 < \rho \leq 1$ be a given parameter. Further, define the function $f : \{1, \ldots, n\} \to \mathbb{R}$ such that $f(1) = b(1)$ and*

$$f(u+1) \doteq \min_{v_u \in \{1,\ldots,n\}} \frac{\min\{u, n - v_u\} \cdot f(u) - b(u)u \cdot \rho + b(v_u)v_u}{\min\{v_u, n-u\}}, \tag{8.23}$$

for all $u \in \{1, \ldots, n-1\}$. Then, for the lowest value $1 \leq \hat{u} \leq n-1$ such that $f(\hat{u}+1) < f(\hat{u})$, it must hold that $f(u+1) < f(u)$ for all $u \in \{\hat{u}, \ldots, n-1\}$.

Proof. The proof is presented in two parts as follows: in Part 1, we identify inequalities given that $f(\hat{u}+1) < f(\hat{u})$, for $1 \leq \hat{u} \leq n-1$ as defined in the claim; and, in Part 2, we use a recursive argument to prove that $f(u+1) < f(u)$ holds for all $\hat{u}+1 \leq u \leq n-1$, using the inequalities derived in Part 1.

Part 1. We define v_u^* as one of the arguments that minimize the right-hand side of (8.23) for each $u \in \{1, \dots, n-1\}$. By assumption, it must hold that $f(\hat{u}+1) < f(\hat{u})$, which implies that

$$\begin{aligned}
f(\hat{u}) &> \min_{v \in \{1,\dots,n\}} \frac{\min\{\hat{u}, n-v\}}{\min\{v, n-\hat{u}\}} f(\hat{u}) - \frac{b(\hat{u})\hat{u}}{\min\{v, n-\hat{u}\}}\rho + \frac{b(v)v}{\min\{v, n-\hat{u}\}} \\
&= \min_{v \in \{1,\dots,n\}} \frac{\min\{\hat{u}-1, n-v\}}{\min\{v, n-\hat{u}+1\}} f(\hat{u}-1) - \frac{b(\hat{u}-1)(\hat{u}-1)}{\min\{v, n-\hat{u}+1\}}\rho + \frac{b(v)v}{\min\{v, n-\hat{u}+1\}} \\
&\qquad + \frac{\min\{\hat{u}, n-v\}}{\min\{v, n-\hat{u}\}} f(\hat{u}) - \frac{\min\{\hat{u}-1, n-v\}}{\min\{v, n-\hat{u}+1\}} f(\hat{u}-1) - \frac{b(\hat{u})\hat{u}}{\min\{v, n-\hat{u}\}}\rho \\
&\qquad + \frac{b(\hat{u}-1)(\hat{u}-1)}{\min\{v, n-\hat{u}+1\}}\rho + \frac{b(v)v}{\min\{v, n-\hat{u}\}} - \frac{b(v)v}{\min\{v, n-\hat{u}+1\}},
\end{aligned}$$

where the strict inequality holds by the definition of $f(\hat{u}+1)$. Recall that

$$f(\hat{u}) \doteq \min_{v \in \{1,\dots,n\}} \frac{\min\{\hat{u}-1, n-v\}}{\min\{v, n-\hat{u}+1\}} f(\hat{u}-1) - \frac{b(\hat{u}-1)(\hat{u}-1)}{\min\{v, n-\hat{u}+1\}}\rho + \frac{b(v)v}{\min\{v, n-\hat{u}+1\}}.$$

Thus, if $v_{\hat{u}}^* \le n-\hat{u}$, the above strict inequality with $f(\hat{u})$ can only be satisfied if

$$f(\hat{u}+1) < f(\hat{u}) \le \hat{u} f(\hat{u}) - (\hat{u}-1) f(\hat{u}-1) < [b(\hat{u})\hat{u} - b(\hat{u}-1)(\hat{u}-1)] \cdot \rho.$$

Similarly, if $v_{\hat{u}}^* \ge n-\hat{u}+1$, then it must hold that

$$\begin{aligned}
&(n - v_{\hat{u}}^*)\left[\frac{f(\hat{u})}{n-\hat{u}} - \frac{f(\hat{u}-1)}{n-\hat{u}+1}\right] + \left[\frac{1}{n-\hat{u}} - \frac{1}{n-\hat{u}+1}\right] b(v_{\hat{u}}^*) v_{\hat{u}}^* \\
&< \left[\frac{b(\hat{u})\hat{u}}{n-\hat{u}} - \frac{b(\hat{u}-1)(\hat{u}-1)}{n-\hat{u}+1}\right] \cdot \rho \\
\implies &\left[\frac{1}{n-\hat{u}} - \frac{1}{n-\hat{u}+1}\right] [(n - v_{\hat{u}}^*) f(\hat{u}) + b(v_{\hat{u}}^*) v_{\hat{u}}^*] < \left[\frac{b(\hat{u})\hat{u}}{n-\hat{u}} - \frac{b(\hat{u}-1)(\hat{u}-1)}{n-\hat{u}+1}\right] \cdot \rho \\
\iff &f(\hat{u}+1) < [b(\hat{u})\hat{u} - b(\hat{u}-1)(\hat{u}-1)]\rho,
\end{aligned}$$

where the first line implies the second line because $f(\hat{u}) \ge f(\hat{u}-1)$, by the definition of $\hat{u}$ in the claim, and the second line is equivalent to the third by the definitions of $f(\hat{u}+1)$ and $v_{\hat{u}}^*$. This concludes Part 1 of the proof.

Part 2. In this part of the proof, we show by recursion that if $f(\hat{u}+1) < f(\hat{u})$, then $f(u+1) < f(u)$ for all $u \in \{\hat{u}+1, ..., n-1\}$. We do so by showing that, if $f(u) < f(u-1) < \cdots < f(\hat{u}+1)$ for any $u \in \{\hat{u}+1, \ldots, n-1\}$, then it must hold that $f(u+1) < f(u)$. Thus, in the following reasoning, we assume that $u \in \{\hat{u}+1, \ldots, n-1\}$, and that $f(u) < f(u-1) < \cdots < f(\hat{u}+1)$.

We begin with the case of $v^*_{u-1} < n-u+1$, which gives us that $v^*_{u-1} \leq n-u$. Observe that

$$
\begin{aligned}
f(u+1) \doteq & \min_{v_u \in \{1,\ldots,n\}} \frac{\min\{u, n-v_u\}}{\min\{v_u, n-u\}} f(u+1) - \frac{b(u)u}{\min\{v_u, n-u\}}\rho + \frac{b(v_u)v_u}{\min\{v_u, n-u\}} \\
= & \min_{v_u \in \{1,\ldots,n\}} \frac{\min\{u-1, n-v_u\}}{\min\{v_u, n-u+1\}} f(u-1) - \frac{b(u-1)(u-1)}{\min\{v_u, n-u+1\}}\rho + \frac{b(v_u)v_u}{\min\{v_u, n-u+1\}} \\
& + \frac{\min\{u, n-v_u\}}{\min\{v_u, n-u\}} f(u+1) - \frac{\min\{u-1, n-v_u\}}{\min\{v_u, n-u+1\}} f(u-1) - \frac{b(u)u}{\min\{v_u, n-u\}}\rho \\
& + \frac{b(u-1)(u-1)}{\min\{v_u, n-u+1\}}\rho + \frac{b(v_u)v_u}{\min\{v_u, n-u\}} - \frac{b(v_u)v_u}{\min\{v_u, n-u+1\}} \\
\leq & f(u) + \frac{u}{v^*_{u-1}} f(u) - \frac{u-1}{v^*_{u-1}} f(u-1) - \frac{b(u)u - b(u-1)(u-1)}{v^*_{u-1}}\rho \\
< & f(u) + \frac{1}{v^*_{u-1}} f(\hat{u}+1) - \frac{1}{v^*_{u-1}}[b(u)u - b(u-1)(u-1)]\rho < f(u),
\end{aligned}
$$

where the first inequality holds by evaluating the minimization at $v_u = v^*_{u-1}$, the second inequality holds because $f(u) < f(u-1)$ and $f(u) \leq f(\hat{u}+1)$, by assumption, and the final inequality holds by the result showed in Part 1 and because $b(\cdot)$ is nondecreasing and convex.

Next, we consider the scenario in which $v^*_{u-1} > n-u+1$. Observe that

$$\begin{aligned}
f(u+1) &\le f(u) + (n - v^*_{u-1})\left[\frac{f(u)}{n-u} - \frac{f(u-1)}{n-u+1}\right] + \left[\frac{1}{n-u} - \frac{1}{n-u+1}\right] b(v^*_{u-1})v^*_{u-1} \\
&\quad - \frac{b(u)u}{n-u}\rho + \frac{b(u-1)(u-1)}{n-u+1}\rho \\
&< f(u) + \left[\frac{1}{n-u} - \frac{1}{n-u+1}\right][(n-v^*_{u-1})f(u-1) + b(v^*_{u-1})v^*_{u-1}] \\
&\quad - \frac{b(u)u}{n-u}\rho + \frac{b(u-1)(u-1)}{n-u+1}\rho \\
&= f(u) + \left[\frac{1}{n-u} - \frac{1}{n-u+1}\right][(n-u+1)f(u) + b(u-1)(u-1)\rho] \\
&\quad - \frac{b(u)u}{n-u}\rho + \frac{b(u-1)(u-1)}{n-u+1}\rho \\
&\le f(u) + \frac{1}{n-u}f(\hat{u}+1) - \frac{1}{n-u}[b(u)u - b(u-1)(u-1)]\rho \\
&< f(u),
\end{aligned}$$

where the first inequality holds by evaluating the minimization at $v_u = v^*_{u-1}$, the second inequality holds because $f(u) < f(u-1)$, by assumption, the equality holds by the definitions of $f(u)$ and v^*_{u-1}, the third inequality holds because $f(u) \le f(\hat{u}+1)$, by assumption, and the final inequality holds by the identity we showed in Part 1 and because b is nondecreasing and convex.

Finally, we consider the scenario in which $v^*_{u-1} = n-u+1$. Observe that

$$\begin{aligned}
f(u+1) &\le f(u) + \frac{u-1}{n-u}f(u) - \frac{u-1}{n-u+1}f(u-1) - \frac{b(u)u}{n-u}\rho + \frac{b(u-1)(u-1)}{n-u+1}\rho \\
&\quad + \left[\frac{1}{n-u} - \frac{1}{n-u+1}\right] b(n-u+1)(n-u+1) \\
&< f(u) + \left[\frac{1}{n-u} - \frac{1}{n-u+1}\right][(n-v^*_{u-1})f(u-1) + b(v^*_{u-1})v^*_{u-1}] \\
&\quad - \frac{b(u)u}{n-u}\rho + \frac{b(u-1)(u-1)}{n-u+1}\rho \\
&< f(u),
\end{aligned}$$

where the first inequality holds by evaluating the minimization at $v_u = v^*_{u-1}$, the second inequality holds because $f(u) < f(u-1)$, by assumption, and the final inequality holds by the same reasoning as for $v^*_{u-1} > n - u + 1$. □

8.3.2 Proofs from Section 8.2

In this section, we prove the claim in Theorem 8.2.1 by constructing a utility mechanism that achieves the best achievable price of anarchy of $1 - c/e$ with respect to the set of all nonnegative, nondecreasing concave welfare functions with maximum curvature $c \in [0, 1]$. In scenarios where a more specific set of welfare functions is considered, we outline how the techniques used to prove Theorem 8.2.1 can be generalized to derive tighter *a priori* bounds on the best achievable price of anarchy.

Before presenting the proof of Theorem 8.2.1, we provide an informal outline of the three steps underpinning the result. These steps correspond with the three parts of the formal proof, but are listed in a different order for sake of clarity. For the reader's convenience, we include the part of the proof that corresponds with each of the steps in our informal outline. The proof is summarized as follows:

–Step #1: We demonstrate that any concave welfare function can be decomposed as a linear combination with nonnegative coefficients of a specialized set of basis functions. [Section 8.3.3, Part ii)]

–Step #2: We derive optimal basis utility functions for each of the basis functions in the specialized set. [Section 8.3.3, Part i)]

–Step #3: We construct local utility functions as linear combinations over the optimal basis utility functions from Step 2 with the nonnegative coefficients derived in Step 1. Finally, we demonstrate that this tractable approach for constructing resource utility functions provides near optimal efficiency guarantees. [Section 8.3.3, Part iii)]

8.3.3 Proof of Theorem 8.2.1

Here we consider the class of games induced by the set of all concave welfare functions with maximum curvature $c \in [0, 1]$. The proof of Theorem 8.2.1 proceeds in the following three parts:

i) Given a value $c \in [0, 1]$, we derive explicit expressions for the local utility functions that maximize the price of anarchy relative to a restricted class of nonnegative, nondecreasing concave welfare functions with curvature c. Among the optimal price of anarchy values obtained for the functions in this restricted class, the lowest is equal to $1 - c/e$;

ii) We show that any nonnegative, nondecreasing concave welfare function W with curvature less than or equal to c can be represented as a linear combination with explicitly defined nonnegative coefficients over this restricted class; and,

iii) We demonstrate that using the local utility functions computed as a linear combination over the optimal local utility functions from i) with the nonnegative coefficients from ii) guarantees that $\mathrm{PoA}(\mathcal{G}) = 1 - c/e$ within the set of resource allocation games $\mathcal{G}$ induced by all nonnegative, nondecreasing concave welfare functions with maximum curvature c.

The above parts successfully prove Theorem 8.2.1 as we argue here. Note that, by part i), the lowest optimal price of anarchy among welfare functions in the restricted class considered is equal to $1 - c/e$, for given curvature $c \in [0, 1]$. By part iii), this implies that all resource allocation games induced by nonnegative, nondecreasing concave welfare functions with maximum curvature c have optimal price of anarchy equal to $1 - c/e$. This is because, by part ii), any such welfare function can be represented as a nonnegative linear combination over the restricted class of welfare functions we consider. Since the best achievable price of anarchy for at least one of the functions in the restricted class is also $1-c/e$, one cannot further improve the price of anarchy within the set of games considered. In addition, parts i)–iii) combine to prove that a corresponding utility mechanism that maximizes the price of anarchy entails computing nonnegative linear combinations over a class of functions with explicit expressions. Thus, the computation of optimal local utility functions is polynomial in the number of players.

Part i). In this part of the proof, we provide explicit expressions for local utility functions that maximize the price of anarchy with respect to a restricted set of welfare functions, as well as the corresponding optimal price of anarchy. To that end, given parameters $\alpha \in [0,1]$ and $\beta \in \mathbb{N}_{\geq 1}$, we define the (α,β)-*coverage* function as

$$V_\beta^\alpha(x) := (1-\alpha)\cdot x + \alpha \cdot \min\{x,\beta\}. \tag{8.24}$$

It is straightforward to verify that every (α,β)-coverage function is nonnegative, nondecreasing concave. In the lemma below, we derive a local utility function that maximizes the price of anarchy of the set of resource allocation games induced by any given (α,β)-coverage function. We use this result to derive the optimal utility functions for a broad range of local welfare functions in Part iii).

Lemma 8.3.3. *Consider the set of resource allocation games $\mathcal{G}$ induced by the (α,β)-coverage function*

$$V_\beta^\alpha(x) = (1-\alpha)\cdot x + \alpha \cdot \min\{x,\beta\},$$

where $\alpha \in [0,1]$ and $\beta \in \mathbb{N}_{\geq 1}$. Let $\rho = (1-\alpha\cdot\beta^\beta e^{-\beta}/(\beta!))^{-1}$, and define F_β^α as in the following recursion: $F_\beta^\alpha(1) := W(1)$,

$$F_\beta^\alpha(x+1) := \max\left\{\frac{1}{\beta}[xF_\beta^\alpha(x) - V_\beta^\alpha(x)\rho] + 1, 1-\alpha\right\},\ \forall x = 1,\ldots,n-1. \tag{8.25}$$

Then, the local utility function F_β^α maximizes the price of anarchy and the corresponding price of anarchy is $\mathrm{PoA}(\mathcal{G}) = 1/\rho$.

According to the result in Lemma 8.3.3, the maximum achievable price of anarchy in resource allocation games induced by a (α,β)-coverage function with $\alpha = 1$ and $\beta \geq 1$ is $1-\beta^\beta e^{-\beta}/(\beta!)$. Surprisingly, Barman et al. [109] show that that the optimal approximation ratio of any polynomial-time algorithm for the same class of resource allocation problems is also $1-\beta^\beta e^{-\beta}/(\beta!)$. Similarly, the optimal price of anarchy for the (α,β)-coverage function with $\alpha\in[0,1]$ and $\beta=1$ is $1-\alpha/e$, which matches the best achievable approximation ratio of any polynomial-time algorithm for this problem setting [107].

Part ii). In the next result, we show that any nonnegative, nondecreasing concave welfare function with maximum curvature $c\in[0,1]$ can be represented as a nonnegative linear combination over the set of (c,k)-coverage functions with $k=1,\dots,n$.

Lemma 8.3.4. *Let $W:\mathbb{N}\to\mathbb{R}$ denote a nonnegative, nondecreasing concave function with curvature less than or equal to $c\in[0,1]$. Then, the nonnegative coefficients $\eta_1,\dots,\eta_n$ satisfy*

$$W(x)=\sum_{k=1}^{n}\eta_k\cdot V_k^c(x),\quad \forall x=0,1,\dots,n, \tag{8.26}$$

where $\eta_1:=[2W(1)-W(2)]/c$, $\eta_k:=[2W(k)-W(k-1)-W(k+1)]/c$, for $k=2,\dots,n-1$, and $\eta_n:=W(1)-\sum_{k=1}^{n-1}\eta_k$.

Proof. The proof is by construction. Define coefficients $\eta_1:=[2W(1)-W(2)]/c$, $\eta_j:=[2W(j)-W(j-1)-W(j+1)]/c$, $j=2,\dots,n-1$, and $\eta_n:=W(1)-\sum_{j=1}^{n-1}\eta_j=W(1)-[W(1)+W(n-1)-W(n)]/c$. It is straightforward to verify that $\eta_j\geq 0$ for all $k=1,\dots,n$ recalling that $W(0)=0$ and $W(x)$ is nonnegative, nondecreasing concave for $x\geq 0$. We defer the proof that $W(x)=\sum_{k=1}^{n}\eta_k\cdot V_k^c(x)$ for all $x=1,\dots,n$ to the proof of Corollary 8.4.1, where one need only substitute $W^{ub}(x)=x$ and $W^{lb}(x)=V_1^c(x)$, for $x\geq 0$. □

8.3.4 Proof of Lemma 8.3.3

Proof. We first dispense with the situation where $n\leq\beta$. In this case, the local welfare function is identical to $W(x)=x$, and thus the price of anarchy is 1 for choice of $F(x)=W(x)/x$. For the remainder of the proof, we only consider $n>\beta$.

The remainder of the proof is structured as follows: (i) we introduce a relaxation of the linear program in Equation (8.2); (ii) in this relaxed linear program, we determine what are the most restrictive constraints for each $x \in \{1, \ldots, n-1\}$; (iii) we show that a feasible solution to the relaxed linear program is nonincreasing, i.e., $F(x+1) \leq F(x)$, for every $x \in \{1, \ldots, n-1\}$ such that $F(x) > 1-\alpha$, and $F(x+1) = 1-\alpha$ otherwise; (iv) we show that (F, ρ) as defined in the claim is a solution to the relaxed linear program for $n \to \infty$; and (v) we observe that (F, ρ) as defined in the claim is feasible in the linear program in Equation (8.2) and thus a solution to this linear program as well.

Relaxed linear program. First we consider a relaxation of the linear program in Equation (8.2). In this relaxed linear program, only the constraints where $z = \min\{0, x+y-n\}$ and $x, y \in \{0, \ldots, n\}$ are retained. Finally, we exclude the constraint with $y = 0$, for all $x \in \{1, \ldots, n-1\}$, resulting in the following *relaxed linear program*:

$$\begin{aligned} &\max_{F \in \mathbb{R}^n, \rho \in \mathbb{R}} \rho \quad \text{subject to:} \\ &W(y) - \rho W(x) + \min\{x, n-y\}F(x) - \min\{y, n-x\}F(x+1) \leq 0, \\ &\qquad \forall (x, y) \in \{0, \ldots, n\} \times \{1, \ldots, n\} \cup (n, 0). \end{aligned} \tag{8.27}$$

Tightest constraints on ρ. We characterize what value $y \in \{1, \ldots, n\}$ parameterizes the tightest constraint for each $x \in \{1, \ldots, n-1\}$. For any $x = 1, \ldots, n-1$, if $1 \geq F(x), F(x+1) \geq 1-\alpha$, we observe that the constraint with $y = \beta$ is strictest. For $y < \beta$, it holds that

$$\begin{aligned} \rho W(x) &\geq \beta + \min\{x, n-\beta\}F(x) - \min\{\beta, n-x\}F(x+1) \\ &\geq y + \min\{x, n-y\}F(x) - \min\{y, n-x\}F(x+1), \end{aligned}$$

where the final inequality holds when $x \leq n-\beta$ because $\beta - y \geq (\beta - y)F(x+1)$; when $n-\beta < x \leq n-y$ because $\beta - y - (x+\beta-n)F(x) \geq n-x-y \geq (n-x-y)F(x+1)$ since

$x+\beta-n>0$; and when $n-y<x$ because $\beta-y\geq(\beta-y)F(x)$. For constraints with $y>\beta$,

$$\begin{aligned}&\rho W(x)\\&\geq \alpha\beta+(1-\alpha)\beta+\min\{x,n-\beta\}F(x)-\min\{\beta,n-x\}F(x+1)\\&\geq \alpha\beta+(1-\alpha)y+\min\{x,n-y\}F(x)-\min\{y,n-x\}F(x+1),\end{aligned}$$

where the final inequality holds when $x\leq n-y$ because $(y-\beta)F(x+1)\geq(y-\beta)(1-\alpha)$, when $n-y<x\leq n-\beta$ because $(1-\alpha)(\beta-y)+(x+y-n)F(x)\geq(1-\alpha)(\beta+x-n)\geq(\beta+x-n)F(x+1)$ since $x+y-n>0\geq\beta+x-n$, and when $n-\beta<x$ because $(1-\alpha)(\beta-y)\geq(\beta-y)F(x)$.

For any $x=1,\ldots,n-1$, if $F(x)\geq 1-\alpha\geq F(x+1)$ and $n-x\geq\beta$, then the constraint with $y=n-x$ is strictest among all constraints as, for any $y\neq n-x$, it holds that

$$\begin{aligned}&\rho W(x)\\&\geq \alpha\beta+(1-\alpha)(n-x)+xF(x)-(n-x)F(x+1)\\&\geq \alpha\beta+(1-\alpha)y+\min\{x,n-y\}F(x)-\min\{y,n-x\}F(x+1)\\&\geq W(y)+\min\{x,n-y\}F(x)-\min\{y,n-x\}F(x+1),\end{aligned}$$

where the inequality holds because $(1-\alpha)(n-x-y)\geq(n-x-y)F(x+1)$ when $x\leq n-y$ and $(1-\alpha)(n-x-y)\geq(n-x-y)F(x)$ when $x>n-y$. For any $x=1,\ldots,n-1$, if $F(x)\geq 1-\alpha\geq F(x+1)$ and $n-x<\beta$, then $y=\beta$ is strictest as for any $y\neq\beta$, it holds that

$$\begin{aligned}\rho W(x)&\geq \beta+\min\{x,n-\beta\}F(x)-\min\{\beta,n-x\}F(x+1)\\&\geq W(y)+\min\{x,n-y\}F(x)-\min\{y,n-x\}F(x+1),\end{aligned}$$

where $\beta-y+(n-\beta-x)F(x)\geq(\beta-y)[1-F(x)]+(n-x-y)F(x)\geq(n-x-y)F(x+1)$ when $x\leq n-y$ since $y\leq n-x<\beta$, $(\beta-y)[1-F(x)]\geq 0$ when $x>n-y$ and $n-x<y\leq\beta$, and $(1-\alpha)(\beta-y)+(y-\beta)F(x)\geq 0$ when $y>\beta$ since $F(x)\geq 1-\alpha$.

For any $x=1,\ldots,n-1$, if $F(x+1),F(x)\leq 1-\alpha$, then the constraint with $y=n$ is strictest among all constraints as, for any $y < n$, it holds that

$$\begin{aligned}
&\rho W(x)\\
&\geq \alpha\beta + (1-\alpha)n + \min\{x, 0\}F(x) - \min\{n, n-x\}F(x+1)\\
&\geq \alpha\beta + (1-\alpha)y + \min\{x, n-y\}F(x) - \min\{y, n-x\}F(x+1)\\
&\geq W(y) + \min\{x, n-y\}F(x) - \min\{y, n-x\}F(x+1),
\end{aligned}$$

where the second last inequality holds because $(n-y)[1-\alpha-F(x+1)] \geq x[F(x)-F(x+1)]$ when $x \leq n-y$ and $(n-y)(1-\alpha) \geq (n-y)F(x)$ when $x > n-y$.

Thus, if $F(x) \geq F(x+1) \geq 1-\alpha$, it is sufficient to consider only the constraint with $y = \beta$ and $z = \max\{0, x+\beta-n\}$. If $F(x) \geq 1-\alpha \geq F(x+1)$ and $n-x > \beta$, it is sufficient to consider only the constraint with $y = n-x$ and $z = 0$. Otherwise, if $F(x) \geq 1-\alpha \geq F(x+1)$ and $n-x < \beta$, then we consider only the constraint $y = \beta$ and $z = \max\{0, x+y-n\}$. Finally, if $1-\alpha \geq F(x) \geq F(x+1)$, then the constraint with $y = n$ and $z = x$ is the strictest.

Proof that a solution to Equation (8.27) has F 'nonincreasing'. For this portion of the proof, consider a function F defined for any given $\rho > 1$ as follows: $F(1) = W(1)$ and, for all $x \in \{1, \ldots, n-1\}$,

$$F(x+1) = \max_{y \in \{1,\ldots,n\}} \frac{\min\{x, n-y\}F(x) - W(x)\rho + W(y)}{\min\{y, n-x\}}.$$

For conciseness, we will use the shorthand $\kappa_x = \frac{\min\{x,n-y^*\}}{\min\{y^*,n-x\}}$, $\lambda_x = \frac{W(x)}{\min\{y^*,n-x\}}$ and $\mu_x = \frac{W(y^*)}{\min\{y^*,n-x\}}$ where $y^* \in \{1, \ldots, n\}$ maximizes the above expression for each $x \in \{1, \ldots, n\}$. Thus, $F(x+1) = \kappa_x F(x) - \lambda_x \rho + \mu_x$ for each $x \in \{1, \ldots, n\}$.

We assume, by contradiction, that we are given ρ such that F is increasing at some index, i.e., $F(x+1) > F(x)$ for $x \in \{1, \ldots, n\}$. The forthcoming Lemma 8.3.5 shows that, if this is the case, then F must continue to increase, so that $F(n) > F(n-1)$. We wish to show the following: (i) if F first increases at a point $x \in \{1, \ldots, n-1\}$ where $F(x) > 1-\alpha$, this leads to a contradiction for the value of ρ; and, (ii) if F first increases at a point x where $F(x) \leq 1-\alpha$, then either $F(j) \leq 1-\alpha$ for all $j \geq x$ is feasible or (F, ρ) is infeasible. It is important to note that the value $nF(n)/W(n)$ must be bounded, otherwise

$$\rho \geq \max_{y \in \{0,\ldots,n\}} \frac{W(y)+(n-y)F(n)}{W(n)} \geq \frac{nF(n)}{W(n)} \to \infty.$$

This is a contradiction as the price of anarchy will be at least 0.5, even if we use the marginal contribution utility [110]. Since we are optimizing for the price of anarchy (i.e., $1/\rho$), we need only consider values of ρ no greater than 2.

Observe that if F first increases at some point $x \in \{1,\ldots,n\}$ such that $F(x) > 1-\alpha$, then $F(n) > F(n-1) > 1-\alpha$ and

$$\begin{aligned}
\rho &= \max_{y \in \{0,\ldots,n\}} \frac{W(y)+(n-y)F(n)}{W(n)} \\
&\geq \max_{y \in \{1,\ldots,n\}} \frac{W(y)+(n-y)F(n)}{W(n)} \\
&= \max_{y \in \{1,\ldots,n-1\}} \frac{W(y)+(n-y)F(n)}{W(n)} \\
&> \max_{y \in \{1,\ldots,n-1\}} \frac{W(y)+(n-y)F(n-1)}{W(n)},
\end{aligned}$$

where the first inequality holds because we reduce the domain of maximization, the second equality holds because $W(n-1)+F(n) > W(n)$ since $n > \beta$ and $F(n) > 1-\alpha$, and the final inequality holds because $F(n) > F(n-1)$. Since $y \leq \beta \leq n-1$ corresponds with the strictest constraints when $F(n), F(n-1) \geq 1-\alpha$, we can substitute

$$\max_{y \in \{1,\ldots,n-1\}} [W(y)+(n-y)F(n-1)] = F(n) + \rho W(n-1)$$

in the former bound on ρ to get $\rho > [F(n) + \rho W(n-1)]/W(n)$. Since $n > \beta$, this implies that

$$\begin{cases} F(n) < 0 & \text{if } \alpha = 1, \\ \rho > \frac{F(n)}{W(n)-W(n-1)} = \frac{F(n)}{1-\alpha} & \text{if } \alpha \in [0,1) \end{cases}$$

For $\alpha = 1$ this is a contradiction, since we have that $F(n) > 1 - \alpha = 0$. For the remaining $\alpha \in [0,1)$, we want to prove that this also gives rise to a contradiction. To do so, we show that

$$\frac{F(n)}{1-\alpha} \geq \max_{y \in \{0,\ldots,n\}} \frac{W(y) + (n-y)F(n)}{W(n)} = \rho.$$

Observe that $y = \beta$ maximizes the right-hand side if $1 - \alpha < F(n) \leq 1$, and $y = 0$ maximizes the right-hand side if $F(n) > 1$. For $F(n) \in (1-\alpha, 1]$ and $y = \beta$, it holds that

$$\begin{aligned} &\frac{F(n)}{1-\alpha} - \frac{W(\beta) + (n-\beta)F(n)}{W(n)} \geq 0 \\ \iff &F(n)W(n) - (1-\alpha)[W(\beta) + (n-\beta)F(n)] \geq 0 \\ \impliedby &F(n)\frac{W(n) - W(\beta)}{n-\beta} = (1-\alpha)F(n) \end{aligned}$$

where the first and second line are equivalent because $\alpha \in [0,1)$ and the third line implies the second because $F(n) > 1-\alpha$. The final inequality holds because $[W(n)-W(\beta)] = (n-\beta)(1-\alpha)$ and $n > \beta$. For $F(n) > 1$ and $y = 0$, it holds that

$$\frac{F(n)}{1-\alpha} - \frac{nF(n)}{W(n)} \geq 0$$

since $W(n)/n \geq 1 - \alpha$ by definition. Thus, in the above reasoning, we have shown that, if F first increases at a point $x \in \{1, \ldots, n\}$ and $F(x) > 1 - \alpha$, then, if $\alpha = 1$, it holds that

$1 - \alpha < F(x) < \cdots < F(n) < 1 - \alpha$; and, if $\alpha \in [0, 1)$, it holds that

$$\rho > \frac{F(n)}{1-\alpha} \geq \max_{y \in \{0,\ldots,n\}} \frac{W(y) + (n-y)F(n)}{W(n)} = \rho,$$

which is a contradiction.

Now we consider the scenario where we are at a point x such that $F(x) \leq 1 - \alpha$ and F is monotonically nonincreasing before x. We show that either selecting $F(x) = \cdots = F(n) = 1-\alpha$ is feasible for ρ or that the value ρ is infeasible. We first consider the case where the strictest constraint on the value of $F(x+1)$ has $y \leq n-x$ and show that $F(x+1)$ cannot be greater than $F(x) \leq 1 - \alpha$. In the proof of Lemma 8.3.5, we showed that if $y \leq n - x$, $F(x+1) > F(x)$ and $F(x) \leq F(x-1)$, then it must hold that $F(x) > [W(x) - W(x-1)]\rho \geq 1-\alpha$. As we have assumed $F(x) \leq 1-\alpha$, it must be that $F(x+1) \leq F(x) \leq 1-\alpha$ if $y \leq n-x$. We complete our reasoning for the case when $y > n-x$ corresponds to the strictest constraint on the value of $F(x+1)$. We showed above that if $F(x) \leq 1-\alpha$ and $x > n-y$, then the strictest constraint is parameterized by $y = n$. For any $x \geq \beta$, it must hold that $F(x+1) \geq -W(x)\rho/(n-x) + W(n)/(n-x)$. Since $-W(x)\rho/(n-x) + W(n)/(n-x) \leq 1-\alpha$ the constraint is satisfied for choice of $F(x+1) \leq 1-\alpha$. Else, if $x < \beta$, since $\beta < n$, $F(x+1) > 1-\alpha$ implies that $F(n) > F(n-1) > 1-\alpha$, since $n - 1 \geq \beta$. We already proved above that this scenario leads to a contradiction on the value of ρ. Repeating this reasoning for all $j > x$ such that $F(x) \leq 1-\alpha$, we argue that $F(j) = 1-\alpha$ is feasible. Since the strictest constraint for each $F(j), F(j+1) \leq 1-\alpha$ has $y = n$, there is no recursion and the optimal value ρ has no dependence on the values of $F(x), \ldots, F(n)$, even if it begins increasing. We have also shown that the lower bound on F is lower than or equal to $1-\alpha$ for any feasible ρ, and so, F with $F(j) = 1-\alpha$, for all $j \in \{x, \ldots, n\}$, must be feasible.

For any feasible ρ, we have successfully shown that $F(x)$ must be nonincreasing when it is greater than $1-\alpha$, and that $F(x) = \cdots = F(n) = 1-\alpha$ is feasible otherwise. This concludes this part of the proof.

Proof that (F, ρ) *solves Equation (8.27).* We begin by showing that (F, ρ) as defined in the claim are feasible. For $x=0$, the constraints in Equation (8.27) read as $F(1) \geq W(y)/\min\{y, n-x\}$, for all $y=1,\ldots,n$, which is satisfied for $F(1)=W(1)$. Now consider $(x,y) \in \{1,\ldots,n-1\} \times \{1,\ldots,n\}$. In the above reasoning, we showed that a feasible (F,ρ) within Equation (8.27) will have F nonincreasing while $F(x) > 1-\alpha$ and $F(x) = 1-\alpha$ otherwise. Furthermore, we showed that when $F(x) \geq F(x+1) > 1-\alpha$ or when $F(x) \geq 1-\alpha \geq F(x+1)$ and $n-x < \beta$, then the strictest constraint has $y=\beta$. Observe that $\kappa_x = \min\{x, n-\beta\}/\min\{\beta, n-x\}$, $\lambda_x = W(x)/\min\{\beta, n-x\}$ and $\mu_x = \beta/\min\{\beta, n-x\}$ correspond with the recursive definition of $F(x+1)$.

We showed above that $F(x+1) = 1-\alpha$ when $\kappa_x F(x) - \lambda_x \rho + \mu_x \leq 1-\alpha$, is feasible as long as ρ is feasible, since the values of F less than or equal to $1-c$ have no impact on the optimal value of the relaxed linear program. Consider the expression for ρ that can be obtained by completing the recursion as follows,

$$1-\alpha = F(\hat{x}+1) \geq \Pi_{u=1}^{\hat{x}} \kappa_u F(1) + \sum_{u=1}^{\hat{x}-1} (\Pi_{v=u+1}^{\hat{x}} \kappa_v)(\mu_u - \lambda_u \rho) + \mu_{\hat{x}} - \lambda_{\hat{x}} \rho.$$

Rearranging this expression, we obtain,

$$\rho \geq \frac{\Pi_{u=1}^{\hat{x}} \kappa_u F(1) + \sum_{u=1}^{\hat{x}-1} (\Pi_{v=u+1}^{\hat{x}} \kappa_v)\mu_u + \mu_{\hat{x}} + \alpha - 1}{\sum_{u=1}^{\hat{x}-1} (\Pi_{v=u+1}^{\hat{x}} \kappa_v)\lambda_u + \lambda_{\hat{x}}}.$$

Observe that for $n \to \infty$, $\min\{x, n-\beta\} = x$ and $\min\{\beta, n-x\} = \beta$. Thus, the above expression for ρ simplifies to

$$\begin{aligned}
\rho &\geq \frac{\frac{\hat{x}!}{\beta^{\hat{x}}} + \sum_{u=1}^{\hat{x}-1} \frac{\hat{x}!}{j!} \frac{1}{\beta^{\hat{x}-j}} + 1 + \alpha - 1}{\sum_{j=1}^{\beta} \frac{\hat{x}!}{j!\beta^{\hat{x}-j}} \frac{j}{\beta} + \sum_{j=\beta+1}^{\hat{x}-1} \frac{\hat{x}!}{j!\beta^{\hat{x}-j}} \frac{\alpha\beta + (1-\alpha)j}{\beta}} \\
&= \frac{1 + \sum_{j=1}^{\hat{x}-1} \frac{\beta^j}{j!} + \alpha \frac{\beta^{\hat{x}}}{\hat{x}!}}{\sum_{j=0}^{\beta-1} \frac{\beta^j}{j!} + \sum_{j=\beta+1}^{\hat{x}-1} \frac{\beta^j}{j!} \frac{\alpha\beta + (1-\alpha)j}{\beta}} = \frac{e^\beta}{e^\beta - \alpha \frac{\beta^\beta}{\beta!}}.
\end{aligned}$$

Noting that PoA $= 1/\rho$ concludes this part of the proof.

Feasibility of (F,ρ) in Equation (8.2). To conclude the proof, we simply observe that since

$F(x)$ is nonincreasing for all x, the strictest constraints in the linear program in Equation (8.2) correspond with the choice of $z = \min\{0, x + y - n\}$. Thus, since (F, ρ) is a solution to the relaxed linear program and feasible in the original linear program, it must also be a solution to the original. $\square$

Lemma 8.3.5. *Let W be a nonnegative, nondecreasing concave function, and let $\rho \geq 1$ be a given parameter. Further, define the function F such that $F(1) = W(1)$ and*

$$F(j+1) := \max_{\ell \in \{1,\ldots,n\}} \frac{\min\{j, n-\ell\}F(j) - W(j)\rho + W(\ell)}{\min\{\ell, n-j\}}, \tag{8.28}$$

for all $j = 1, \ldots, n-1$. Then, for the lowest value $\hat{j} = 1, \ldots, n-1$ such that $F(\hat{j}+1) > F(\hat{j})$, it must hold that $F(j+1) > F(j)$ for all $j = \hat{j}, \ldots, n-1$.

Proof. The proof is presented in two parts as follows: in part (i), we identify an inequality that must hold given that $F(\hat{j}+1) > F(\hat{j})$ for $1 \leq \hat{j} \leq n-1$ as defined in the claim; and, in part (ii), we use a recursive argument to prove that $F(j+1) > F(j)$ holds for all $\hat{j}+1 \leq j \leq n-1$, using the inequality we derived in part (i).

Part (i). We define ℓ_j^* as one of the arguments that minimizes the right-hand side of Equation (8.28) for each $x = 1, \ldots, n-1$. By assumption, it must hold that $F(\hat{j}+1) > F(\hat{j})$, which implies that

$$\begin{aligned}
F(\hat{j}) &< \max_{1\le \ell\le n} \frac{\min\{\hat{j}, n-\ell\}}{\min\{\ell, n-\hat{j}\}} F(\hat{j}) - \frac{W(\hat{j})}{\min\{\ell, n-\hat{j}\}}\rho + \frac{W(\ell)}{\min\{\ell, n-\hat{j}\}} \\
&= \max_{1\le \ell\le n} \frac{\min\{\hat{j}-1, n-\ell\}}{\min\{\ell, n-\hat{j}+1\}} F(\hat{j}-1) - \frac{W(\hat{j}-1)}{\min\{\ell, n-\hat{j}+1\}}\rho \\
&\qquad + \frac{W(\ell)}{\min\{\ell, n-\hat{j}+1\}} + \frac{\min\{\hat{j}, n-\ell\}}{\min\{\ell, n-\hat{j}\}} F(\hat{j}) \\
&\qquad - \frac{\min\{\hat{j}-1, n-\ell\}}{\min\{\ell, n-\hat{j}+1\}} F(\hat{j}-1) - \frac{W(\hat{j})}{\min\{\ell, n-\hat{j}\}}\rho \\
&\qquad + \frac{W(\hat{j}-1)}{\min\{\ell, n-\hat{j}+1\}}\rho + \frac{W(\ell)}{\min\{\ell, n-\hat{j}\}} \\
&\qquad - \frac{W(\ell)}{\min\{\ell, n-\hat{j}+1\}},
\end{aligned}$$

where the strict inequality holds by definition of $F(\hat{j}+1)$. Recall that

$$\begin{aligned}
F(\hat{j}) := \max_{1\le \ell\le n} &\frac{\min\{\hat{j}-1, n-\ell\}}{\min\{\ell, n-\hat{j}+1\}} F(\hat{j}-1) - \frac{W(\hat{j}-1)}{\min\{\ell, n-\hat{j}+1\}}\rho \\
&+ \frac{W(\ell)}{\min\{\ell, n-\hat{j}+1\}}.
\end{aligned}$$

Thus, if $\ell^*_{\hat{j}} \le n-\hat{j}$, the above strict inequality with $F(\hat{j})$ can only be satisfied if

$$F(\hat{j}+1) > F(\hat{j}) \ge \hat{j}F(\hat{j}) - (\hat{j}-1)F(\hat{j}-1) > [W(\hat{j}) - W(\hat{j}-1)]\cdot\rho.$$

Similarly, if $\ell^*_{\hat{j}} \geq n - \hat{j} + 1$, then it must hold that

$$\begin{aligned}
&(n - \ell^*_{\hat{j}})\left[\frac{F(\hat{j})}{n-\hat{j}} - \frac{F(\hat{j}-1)}{n-\hat{j}+1}\right] + \left[\frac{1}{n-\hat{j}} - \frac{1}{n-\hat{j}+1}\right]W(\ell_{\hat{j}}) \\
&\quad > \left[\frac{W(\hat{j})}{n-\hat{j}} - \frac{W(\hat{j}-1)}{n-\hat{j}+1}\right]\cdot\rho \\
\Longrightarrow &\left[\frac{1}{n-\hat{j}} - \frac{1}{n-\hat{j}+1}\right][(n-\ell^*_{\hat{j}})F(\hat{j}) + W(\ell_{\hat{j}})] \\
&\quad > \left[\frac{W(\hat{j})}{n-\hat{j}} - \frac{W(\hat{j}-1)}{n-\hat{j}+1}\right]\cdot\rho \\
\Longleftrightarrow &F(\hat{j}+1) > [W(\hat{j}) - W(\hat{j}-1)]\rho,
\end{aligned}$$

where the first line implies the second line because $F(\hat{j}) \leq F(\hat{j}-1)$, by the definition of $\hat{j}$ in the claim, and the second line is equivalent to the third by the definitions of $F(\hat{j}+1)$ and $\ell^*_{\hat{j}}$. This concludes part (i) of the proof.

Part (ii). In this part of the proof, we show by recursion that if $F(\hat{j}+1) > F(\hat{j})$, then $F(j+1) > F(j)$ for all $j = \hat{j}+1, ..., n-1$. We do so by showing that, if $F(j) > F(j-1) > \cdots > F(\hat{j}+1)$ for any $\hat{j}+1 \leq j \leq n-1$, then it must hold that $F(j+1) > F(j)$. Thus, in the following reasoning, we assume that $\hat{j}+1 \leq j \leq n-1$, and that $F(j) > F(j-1) > \cdots > F(\hat{j}+1)$.

We begin with the scenario in which $\ell^*_{j-1} < n - j + 1$, which gives us that $\ell^*_{j-1} \leq n - j$. Recall that

$$\begin{aligned}
F(j+1) := \max_{1\leq \ell_j \leq n} &\frac{\min\{j, n-\ell_j\}}{\min\{\ell_j, n-j\}}F(j+1) - \frac{W(j)}{\min\{\ell_j, n-j\}}\rho \\
&+ \frac{W(\ell_j)}{\min\{\ell_j, n-j\}}.
\end{aligned}$$

Thus, it must hold that

$$
\begin{aligned}
&F(j+1)\\
&= \max_{1\le \ell_j\le n} \frac{\min\{j-1, n-\ell_j\}}{\min\{\ell_j, n-j+1\}} F(j-1) - \frac{W(j-1)}{\min\{\ell_j, n-j+1\}}\rho\\
&\qquad + \frac{W(\ell_j)}{\min\{\ell_j, n-j+1\}} + \frac{\min\{j, n-\ell_j\}}{\min\{\ell_j, n-j\}} F(j+1)\\
&\qquad - \frac{\min\{j-1, n-\ell_j\}}{\min\{\ell_j, n-j+1\}} F(j-1) - \frac{W(j)}{\min\{\ell_j, n-j\}}\rho\\
&\qquad + \frac{W(j-1)}{\min\{\ell_j, n-j+1\}}\rho + \frac{W(\ell_j)}{\min\{\ell_j, n-j\}}\\
&\qquad - \frac{W(\ell_j)}{\min\{\ell_j, n-j+1\}}\\
&\ge F(j) + \frac{j}{\ell^*_{j-1}} F(j) - \frac{j-1}{\ell^*_{j-1}} F(j-1) - \frac{W(j)-W(j-1)}{\ell^*_{j-1}}\rho\\
&> F(j) + \frac{1}{\ell^*_{j-1}} F(\hat{j}+1) - \frac{1}{\ell^*_{j-1}}[W(j)-W(j-1)]\rho\\
&> F(j),
\end{aligned}
$$

where the first inequality holds by evaluating the maximization at $\ell_j = \ell^*_{j-1}$, the second inequality holds because $F(j) > F(j-1)$ and $F(j) \ge F(\hat{j}+1)$, by assumption, and the final inequality holds by the identity we showed in part (i) and because $W(\cdot)$ is concave.

Next, consider the scenario in which $\ell^*_{j-1} > n - j + 1$. Observe that

$$\begin{aligned}
&F(j+1) \\
&\geq F(j) + (n - \ell^*_{j-1})\left[\frac{F(j)}{n-j} - \frac{F(j-1)}{n-j+1}\right] \\
&\quad + \left[\frac{1}{n-j} - \frac{1}{n-j+1}\right] W(\ell^*_{j-1}) - \frac{W(j)}{n-j}\rho + \frac{W(j-1)}{n-j+1}\rho \\
&> F(j) + \left[\frac{1}{n-j} - \frac{1}{n-j+1}\right][(n - \ell^*_{j-1})F(j-1) + W(\ell^*_{j-1})] \\
&\quad - \frac{W(j)}{n-j}\rho + \frac{W(j-1)}{n-j+1}\rho \\
&= F(j) + \left[\frac{1}{n-j} - \frac{1}{n-j+1}\right][(n-j+1)F(j) + W(j-1)\rho] \\
&\quad - \frac{W(j)}{n-j}\rho + \frac{W(j-1)}{n-j+1}\rho \\
&\geq F(j) + \frac{1}{n-j}F(\hat{j}+1) - \frac{1}{n-j}[W(j) - W(j-1)]\rho \\
&> F(j),
\end{aligned}$$

where the first inequality holds by evaluating the maximization at $\ell_j = \ell^*_{j-1}$, the second inequality holds because $F(j) > F(j-1)$, by assumption, the equality holds by the definitions of $F(j)$ and ℓ^*_{j-1}, the third inequality holds because $F(j) \geq F(\hat{j}+1)$, by assumption, and the final inequality holds by the identity we showed in part (i) and because $W(\cdot)$ is concave.

Finally, we consider the scenario in which $\ell^*_{j-1} = n - j + 1$. Observe that

$$\begin{aligned}
&F(j+1) \\
&\geq F(j) + \frac{j-1}{n-j}F(j) - \frac{j-1}{n-j+1}F(j-1) - \frac{W(j)}{n-j}\rho + \frac{W(j-1)}{n-j+1}\rho \\
&\quad + \left[\frac{1}{n-j} - \frac{1}{n-j+1}\right] W(n-j+1) \\
&> F(j) + \left[\frac{1}{n-j} - \frac{1}{n-j+1}\right][(n - \ell^*_{j-1})F(j-1) + W(\ell^*_{j-1})] \\
&\quad - \frac{W(j)}{n-j}\rho + \frac{W(j-1)}{n-j+1}\rho \\
&> F(j),
\end{aligned}$$

where the first inequality holds by evaluating the maximization at $\ell_j = \ell^*_{j-1}$, the second inequality holds because $F(j) > F(j-1)$, by assumption, and the final inequality holds by the same reasoning as for $\ell^*_{j-1} > n - j + 1$. $\square$

Part iii). We begin by describing a utility mechanism parameterized by the maximum curvature and maximum number of players. Let $\mathcal{G}$ denote the set of resource allocation games induced by all nonnegative, nondecreasing concave functions with maximum curvature $c \in [0,1]$ with a maximum of n players . Consider any resource allocation game $G \in \mathcal{G}$ and assign the following local utility function to each $r \in \mathcal{R}$:

$$F_r(x) = \sum_{k=1}^{n} \eta_k \cdot F^c_k(x), \quad \forall x = 1, \ldots, n,$$

where $\eta_1 := [2W_r(1) - W_r(2)]/c$, $\eta_k := [2W_r(k) - W_r(k-1) - W_r(k+1)]/c$, for $k = 2, \ldots, n-1$, and $\eta_n := W_r(1) - \sum_{k=1}^{n-1} \eta_k$, $W_r : \mathbb{N} \to \mathbb{R}$ is the welfare function on the resource r and each function $F^c_k : \mathbb{N} \to \mathbb{R}$, $k = 1, \ldots, n$, is the optimal local utility function for $V^c_k(x)$ defined recursively in Lemma 8.3.3. In this part, we show that $\mathrm{PoA}(G) \geq 1 - c/e$ holds for this utility mechanism.

Given maximum curvature $c \in [0,1]$, Lemma 8.3.3 proves that among the (c,k)-coverage functions with $k = 1, \ldots, n$, the $(c,1)$-coverage function has best achievable price of anarchy $1 - c/e$ which is strictly lower than the best achievable price of anarchy for any (c,k)-coverage function with $k > 1$. This implies that the best achievable price of anarchy must satisfy $\mathrm{PoA}(\mathcal{G}) \leq 1 - c/e$, since any game G in the set of resource allocation games induced by the $(c,1)$-coverage function must also be in the set $\mathcal{G}$, i.e., $G \in \mathcal{G}$, and there is at least one such game with $\mathrm{PoA}(G) = 1 - c/e$. We now show that $\mathrm{PoA}(\mathcal{G}) \geq 1 - c/e$ also holds. Recall from Lemma 8.3.4 that the nonnegative coefficients $\eta_1, \ldots, \eta_n$ defined above satisfy

$$W_r(x) = \sum_{k=1}^{n} \eta_k \cdot V^c_k(x) \quad \forall x = 0, 1, \ldots, n.$$

It must then hold that, for any $r \in \mathcal{R}$, $(F_r, (1-c/e)^{-1})$ is a feasible point in the linear program in Equation (8.2) for any n and the corresponding W_r. Observe that each constraint in the linear program must be satisfied since, by Lemma 8.3.4, it can be represented as a nonnegative linear combination of the constraints in the n linear programs for V_k^c and $(F_k^c, (1-c/e)^{-1})$, $k = 1, \ldots, n$, i.e., for all $r \in \mathcal{R}$ and all $(x, y, z) \in \mathcal{I}(n)$ it must hold that

$$\begin{aligned}
(1-c/e)^{-1} W_r(x) &\geq \sum_{k=1}^{n} \eta_k \cdot \left[1 - c \cdot \frac{k^k e^{-k}}{k!}\right] V_k^c(x) \\
&\geq \sum_{k=1}^{n} \eta_k \cdot \left[V_k^c(y) + (x-z) F_k^c(x) - (y-z) F_k^c(x+1)\right] \\
&= W_r(y) + (x-z) F_r(x) - (y-z) F_r(x+1),
\end{aligned}$$

where the first inequality holds because $1 - c/e \leq 1 - c \cdot k^k e^{-k}/(k!)$ for all $k \geq 1$ and since W_r, $V_k^c(x)$, $k = 1, \ldots, n$, and the coefficients $\eta_1, \ldots, \eta_n$ are nonnegative, and the second inequality holds because $(F^j, \rho^j) = (F_k^c, 1 - c \cdot k^k e^{-k}/(k!))$ is a feasible point in the linear program in Equation (8.2) for $W^j = V_k^c$, by the result in Lemma 8.3.3.

8.4 Proof of Corollary 8.4.1.

Corollary 8.4.1. *Let $\mathcal{W}$ denote a set of nonnegative, nondecreasing concave welfare functions and n be the maximum number of agents. Let W^{ub} and W^{lb} be two nonnegative, nondecreasing concave functions that satisfy the following for all $W \in \mathcal{W}$: (i) $W^{lb}(x+1) - W^{lb}(x) \leq [W(x+1) - W(x)]/W(1) \leq W^{ub}(x+1) - W^{ub}(x)$, for all $x = 1, \ldots, n-1$; and, (ii) $[W(x+1) - 2W(x) + W(x-1)]/W(1) \leq W^{ub}(x+1) - 2W^{ub}(x) + W^{ub}(x-1) \leq W^{lb}(x+1) - 2W^{lb}(x) + W^{lb}(x-1)$, for all $x = 2, \ldots, n-1$. Finally, define the candidate functions $W^{(k)}$, $k = 1, \ldots, n$, as follows:*

$$W^{(k)}(x) = \begin{cases} W^{ub}(x) & \text{if } 1 \leq x \leq k, \\ W^{ub}(k) + W^{lb}(x) - W^{lb}(k) & \text{if } x > k. \end{cases} \tag{8.29}$$

Then, for any welfare function $W \in \mathcal{W}$, there exist nonnegative coefficients $\eta_1, \ldots, \eta_n$ that *satisfy*

$$W(x) = \sum_{k=1}^{n} \eta_k \cdot W^{(k)}(x), \quad \forall x = 0, 1, \ldots, n.$$

Proof. First, observe that there must exist functions W^{ub} and W^{lb}. Simply observe that $W^{ub}(x) = x$ and $W^{lb}(x) = V_1^1(x) = \min\{x, 1\}$ are valid for any set of nonnegative, nondecreasing concave functions.

The rest of the proof follows by construction. Define the coefficients η_j, $j = 0, \ldots, n$, as follows:

$$\eta_1 = \frac{W^{ub}(2) - W^{ub}(1) - W(2) + W(1)}{W^{ub}(2) - W^{ub}(1) - W^{lb}(2) + W^{lb}(1)},$$

$$\eta_j = \frac{W^{ub}(j+1) - W^{ub}(j) - W(j+1) + W(j)}{W^{ub}(j+1) - W^{ub}(j) - W^{lb}(j+1) + W^{lb}(j)} - \sum_{k=1}^{j-1} \eta_k,$$

for $j = 2, \ldots, n-1$ and $\eta_n = 1 - \sum_{k=1}^{n-1} \eta_k$.

First, we prove that the coefficients $\eta_1, \ldots, \eta_n$ are nonnegative. It is simple to see that $\eta_1 \geq 0$ since $W^{ub}(2) - W^{ub}(1) \geq W(2) - W(1) \geq W^{lb}(2) - W^{lb}(1)$. Similarly, $\eta_n \geq 0$ since $\eta_n = 1 - [W^{ub}(n) - W^{ub}(n-1) - W(n) + W(n-1)]/[W^{ub}(n) - W^{ub}(n-1) - W^{lb}(n) + W^{lb}(n-1)]$. Finally, for any $j \in \{2, \ldots, n-1\}$,

$$\begin{aligned}
\eta_j &= \frac{W^{ub}(j+1) - W^{ub}(j) - W(j+1) + W(j)}{W^{ub}(j+1) - W^{ub}(j) - W^{lb}(j+1) + W^{lb}(j)} \\
&\quad - \frac{W^{ub}(j) - W^{ub}(j-1) - W(j) + W(j-1)}{W^{ub}(j) - W^{ub}(j-1) - W^{lb}(j) + W^{lb}(j-1)} \\
&\geq \frac{W^{ub}(j+1) - 2W^{ub}(j) + W^{ub}(j-1)}{W^{ub}(j) - W^{ub}(j-1) - W^{lb}(j) + W^{lb}(j-1)} \\
&\quad - \frac{W(j+1) - 2W(j)W(j) + W(j-1)}{W^{ub}(j) - W^{ub}(j-1) - W^{lb}(j) + W^{lb}(j-1)} \geq 0,
\end{aligned}$$

where the equality holds by definition, the first inequality holds because $W^{lb}(j+1) - 2W^{lb}(j) + W^{lb}(j-1) \geq W^{ub}(j+1) - 2W^{ub}(j) + W^{ub}(j-1)$ and the final inequality holds because $W^{ub}(j+1) - 2W^{ub}(j) + W^{ub}(j-1) \geq W(j+1) - 2W(j) + W(j-1)$.

We conclude the proof by observing that, for all $x = 1, \ldots, n$,

$$
\begin{aligned}
&\sum_{k=1}^{x-1} \eta_k W^{lb}(x) + \sum_{k=1}^{x-1} \eta_k [W^{ub}(k) - W^{lb}(k)] + \sum_{k=x}^{n} \eta_k W^{ub}(x) \\
&= \frac{W^{ub}(x) - W^{ub}(x-1) - W(x) + W(x-1)}{W^{ub}(x) - W^{ub}(x-1) - W^{lb}(x) + W^{lb}(x-1)} W^{lb}(x) \\
&\quad + \sum_{k=1}^{x-1} \eta_k [W^{ub}(k) - W^{lb}(k)] \\
&\quad + \left[1 - \frac{W^{ub}(x) - W^{ub}(x-1) - W(x) + W(x-1)}{W^{ub}(x) - W^{ub}(x-1) - W^{lb}(x) + W^{lb}(x-1)}\right] W^{ub}(x) \\
&= W^{ub}(x) - W^{ub}(x) - W^{ub}(x-1) - W(x) + W(x-1) \\
&\quad + \sum_{k=1}^{x-2} \eta_k [W^{ub}(k) - W^{lb}(k) - W^{ub}(x-1) + W^{lb}(x-1)] \\
&= W(x) + \sum_{k=1}^{x-2} \eta_k [W^{ub}(k) - W^{lb}(k) - W^{ub}(x-1) + W^{lb}(x-1)] \\
&\quad + W^{ub}(x-1) - W(x-1) \\
&= W(x),
\end{aligned}
$$

where the final equality holds once the expression is simplified for the remaining η_k values. $\square$

Chapter 9

Unintended consequences of the worst-case perspective

A number of recent results – including those in the previous chapter – focus on deriving taxes that optimize the Price of Anarchy as a surrogate for optimizing the system performance [72, 88, 90]. Naturally, this raises concerns about the consequences of optimizing for the worst-case equilibrium efficiency on other performance metrics. In this chapter, we seek to understand the impact of optimizing the Price of Anarchy on other performance metrics. In particular, we consider its consequences on the Price of Stability, and on the system performance within transient states which we term *Price of Urgency*.

9.1 Trade-offs with the Price of Stability

In this section, we demonstrate that there exists an inherent trade-off between the Price of Anarchy and the Price of Stability in congestion games, and put forward techniques to study this trade-off. We study this trade-off from two distinct perspectives: an outer trade-off corresponding with the joint optimization of the Price of Anarchy and Price of Stability of a family of congestion games using taxes; and, an inner trade-off corresponding with the jointly achievable Price of Anarchy and Price of Stability pairs for a family of congestion games under a given taxation rule.

Figure 9.1: *The Pareto frontier between the Price of Anarchy and Price of Stability in congestion games with affine and quadratic resource costs.* The Pareto frontier lies within the region below the upper bound curves (solid black) and above the lower bound curves (dotted black), which were derived with the techniques put forward in Theorems 9.1.2 and 9.1.3. The joint Price of Anarchy and Price of Stability values for the mechanism that minimizes the Price of Anarchy (in red), no taxes (in green), and the mechanism that minimizes the Price of Stability (in blue) are reported in the table on the right. Note that all joint performance guarantees in the region above the upper bound curves are suboptimal, in the pink region below the lower bound curves are unachievable by any local mechanism and in the grey region are inadmissible as $\text{PoA}(\mathcal{G}_T) \geq \text{PoS}(\mathcal{G}_T)$ must hold. Although the upper and lower bound curves do not always match, we show that they are tight at the endpoints for any family of convex, nondecreasing resource costs. Similar bounds on the Pareto frontier can be derived for any family of resource cost functions using the techniques outlined in Theorems 9.1.2 and 9.1.3.

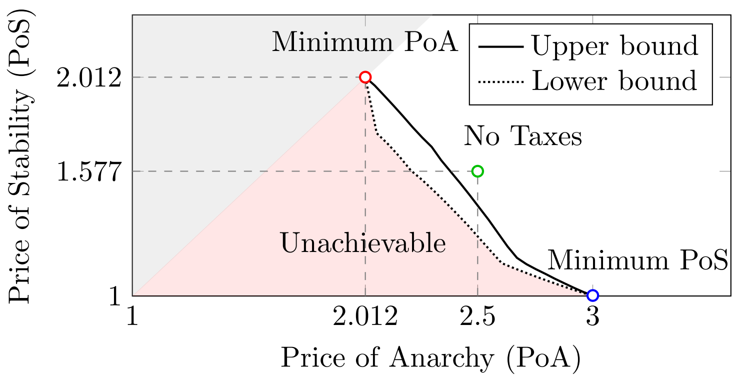

Minimum PoA	
$\text{PoA}(\mathcal{G}_T)$	$\text{PoS}(\mathcal{G}_T)$
2.012	2.012

No Taxes	
$\text{PoA}(\mathcal{G}_T)$	$\text{PoS}(\mathcal{G}_T)$
2.500	1.577

Minimum PoS	
$\text{PoA}(\mathcal{G}_T)$	$\text{PoS}(\mathcal{G}_T)$
3.000	1.000

(a) Congestion games with affine resource costs.

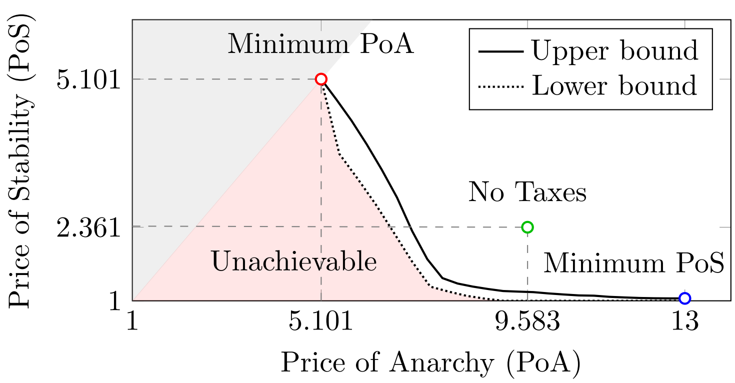

Minimum PoA	
$\text{PoA}(\mathcal{G}_T)$	$\text{PoS}(\mathcal{G}_T)$
5.101	5.101

No Taxes	
$\text{PoA}(\mathcal{G}_T)$	$\text{PoS}(\mathcal{G}_T)$
9.583	2.361

Minimum PoS	
$\text{PoA}(\mathcal{G}_T)$	$\text{PoS}(\mathcal{G}_T)$
13.000	1.000

(b) Congestion games with quadratic resource costs.

9.1.1 The existence of a trade-off

We first seek to investigate how the efficiency of the best-case equilibria is affected when we optimize for the worst-case equilibrium efficiency. In the next result, we prove that any local taxation rule T that minimizes the Price of Anarchy has corresponding Price of Stability equal to the Price of Anarchy in congestion games with convex, nondecreasing resource cost functions. Additionally, we show that linear taxation rules (i.e., local rules T that satisfy $T(\sum_{j=1}^{m} \alpha_j b_j) = \sum_{j=1}^{m} \alpha_j T(b_j)$) are Pareto optimal over all possible local taxation rules, i.e., for any (possibly nonlinear) local rule T, there exists a linear rule T^{lin} such that $\text{PoA}(\mathcal{G}^n_{T^{\text{lin}}}) \leq \text{PoA}(\mathcal{G}^n_T)$ and $\text{PoS}(\mathcal{G}^n_{T^{\text{lin}}}) \leq \text{PoS}(\mathcal{G}^n_T)$.

Theorem 9.1.1. *Consider the family of resource cost functions $\mathcal{L} = span(b_1, \ldots, b_m)$ corresponding to convex, nondecreasing basis functions $b_1, \ldots, b_m$, and maximum number of users n. The following statements hold:*

i) Let T^{lin} denote a Pareto optimal rule in the set of all linear taxation rules. Then, T^{lin} is Pareto optimal over all (possibly nonlinear) local taxation rules.

ii) Let T^{PoA} denote the rule that minimizes the Price of Anarchy as defined in Theorem 8.1.1 where $(F_j^{\text{PoA}}, \rho_j^{\text{PoA}})$, $j = 1, \ldots, m$, are solutions to the m linear programs in (8.3) (one for each b_j). It holds that $\text{PoS}(\mathcal{G}^n_{T^{\text{PoA}}}) = \text{PoA}(\mathcal{G}^n_{T^{\text{PoA}}})$. Furthermore, the functions F_j^{PoA}, $j = 1, \ldots, m$, are nondecreasing and unique up to rescaling.

The proof is presented in Section 9.3. We highlight that the performance guarantee $\text{PoS}(\mathcal{G}^n_{T^{\text{PoA}}}) = \text{PoA}(\mathcal{G}^n_{T^{\text{PoA}}})$ in Theorem 9.1.1ii) is achieved by the same game instance $G \in \mathcal{G}^n_{T^{\text{PoA}}}$. Moreover, the instance G is a simple, n-user game in which each user has 2 single-selection actions and there is a unique pure Nash equilibrium. We depict the structure of the worst-case game instance as a graph in Figure 9.2 where the users are represented by the edges, and the resources are the nodes.

Alternatively, one might wish to understand how unilaterally minimizing the Price of Stability impacts the Price of Anarchy. Here, we show that the marginal cost rule is the unique

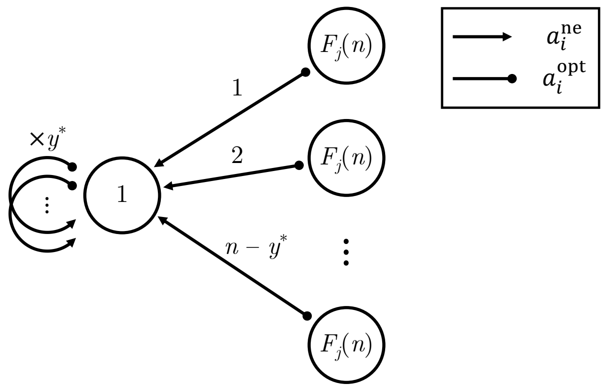

Figure 9.2: *The worst-case game structure satisfying* $\text{PoA}(G) = \text{PoS}(G) = \text{PoA}(\mathcal{G}^n_{T^{\text{PoA}}}) = \text{PoS}(\mathcal{G}^n_{T^{\text{PoA}}})$ *in Theorem 9.1.1.* Consider the local taxation rule T^{PoA} that minimizes the Price of Anarchy as described in Proposition 8.1.1. As shown in Theorem 9.1.1, the resulting Price of Stability is equal to the minimum Price of Anarchy, i.e., $\text{PoS}(\mathcal{G}^n_{T^{\text{PoA}}}) = \text{PoA}(\mathcal{G}^n_{T^{\text{PoA}}})$. We depict the worst-case game structure with the graph above where the n edges are the users and the $1 \leq n - y^* + 1 \leq n + 1$ nodes are the resources. Each resource $e \in \mathcal{E}$ has resource cost function $\ell_e(x) = \alpha_e \cdot b_j(x)$, where $\alpha_e \geq 0$ is the value of the node (either 1 or $F_j(n)$) and b_j is one of the basis functions. Each user $i \in N$ has two single-selection actions, i.e., $\mathcal{A}_i = \{a_i^{\text{ne}}, a_i^{\text{opt}}\}$. In the above depiction, the arrow (resp. round) tip of each edge $i \in N$ indicates the resource i selects in a_i^{ne} (resp. a_i^{opt}). Observe that all n users select the left resource in the joint action $a^{\text{ne}} = (a_1^{\text{ne}}, \ldots, a_n^{\text{ne}})$ which is the unique equilibrium action since the functions F_j^{opt}'s are nondecreasing (by Theorem 9.1.1ii)). In contrast, y^* users select the left resource and the remaining $n - y^*$ users select individual resources in the joint action $a^{\text{opt}} = (a_1^{\text{opt}}, \ldots, a_n^{\text{opt}})$ which is the optimal action. To obtain the result, we select $y^* \in \{0, 1, \ldots, n\}$ and $b_j \in \{b_1, \ldots, b_m\}$ that maximize $\text{SC}(a^{\text{ne}})/\text{MinCost}(G)$.

local rule that achieves the minimum Price of Stability of 1 in congestion games.

Proposition 9.1.1. *Consider the family of resource cost functions* $\mathcal{L} = \textit{span}(b_1, \ldots, b_m)$ *corresponding to positive, nondecreasing basis functions* $b_1, \ldots, b_m$*, and maximum number of users* n*. Then, the marginal cost rule*

$$T^{\text{mc}}(\ell) = \sum_{j=1}^{m} \alpha_j \tau_j^{\text{mc}}, \quad \textit{where } \tau_j^{\text{mc}} : \{1, \ldots, n\} \to \mathbb{R}, \quad \tau_j^{\text{mc}}(x) = (x-1)[b_j(x) - b_j(x-1)]$$

is the unique local taxation rule (up to rescaling) with $\text{PoS}(\mathcal{G}^n_{T^{\text{mc}}}) = 1$.

Proof. We prove the claim in two parts: (i) show that any local taxation rule $T \neq T^{\text{mc}}$ has $\text{PoS}(\mathcal{G}^n_T) > 1$; and (ii) prove that an optimal assignment is an equilibrium under T^{mc}.

Part (i): Assume, by contradiction, that there exists a local taxation rule T with $\text{PoS}(\mathcal{G}^n_T) =$

1 with $T(\ell)(k) > T^{\mathrm{mc}}(\ell)(k)$ for some integer $1 \leq k \leq n$ and resource cost function $\ell \in \mathcal{L}$. Consider the game G with user set $N = \{1, \ldots, k\}$ and two resources $\mathcal{E} = \{e_0, e_1\}$. The resource e_0 has resource cost $\ell(x)$ and resource e_1 has resource cost $[\ell(k)k - \ell(k-1)(k-1) + \epsilon] \cdot \ell(x)$ for $x = 1, \ldots, k$ where $0 < \epsilon < T(\ell)(k) - T^{\mathrm{mc}}(\ell)(k)$. Every user $i \in \{1, \ldots, k-1\}$ has only one action, $a_i = \{e_0\}$, while user k has action set $\mathcal{A}_k = \{a_k, a_k'\}$, where $a_k = \{e_0\}$ and $a_k' = \{e_1\}$. Observe that if $T(\ell)(k) > T^{\mathrm{mc}}(\ell)(k)$, then the unique pure Nash equilibrium corresponds with when user k selects a_k' resulting in social cost $\ell(k-1)(k-1) + \ell(k)k - \ell(k-1)(k-1) + \epsilon$. Thus, the Price of Stability in this game is $[\ell(k)k + \epsilon]/[\ell(k)k] > 1$, which contradicts $\mathrm{PoS}(\mathcal{G}_T^n) = 1$. We conclude this part by observing that a similar argument holds for $T(\ell)(k) < T^{\mathrm{mc}}(\ell)(k)$, when the resource cost of e_1 is $[\ell(k)k - \ell(k-1)(k-1) - \epsilon] \cdot \ell(x)$ for $x = 1, \ldots, k$ and $0 < \epsilon < T^{\mathrm{mc}}(\ell)(k) - T(\ell)(k)$. In this case, user k's Nash action is a_k and the Price of Stability is $\ell(k)k/[\ell(k)k - \epsilon] > 1$.

Part (ii): Consider an optimal assignment a^{opt} in a given game G. It is straightforward to show that this assignment must be an equilibrium under T^{mc}:

$$
\begin{aligned}
C_i(a^{\mathrm{opt}}) - C_i(a_i, a_{-i}^{\mathrm{opt}}) &= \sum_{e \in a_i^{\mathrm{opt}}} \Big[\ell_e(|a^{\mathrm{opt}}|_e)|a^{\mathrm{opt}}|_e - \ell_e(|a^{\mathrm{opt}}|_e - 1)(|a^{\mathrm{opt}}|_e - 1)\Big] \\
&\quad - \sum_{e \in a_i} \Big[\ell_e(|a_i, a_{-i}^{\mathrm{opt}}|_e)|a_i, a_{-i}^{\mathrm{opt}}|_e - \ell_e(|a_i, a_{-i}^{\mathrm{opt}}|_e - 1)(|a_i, a_{-i}^{\mathrm{opt}}|_e - 1)\Big] \\
= \sum_{e \in a_i^{\mathrm{opt}} \setminus a_i} \ell_e(|a^{\mathrm{opt}}|_e)|a^{\mathrm{opt}}|_e &+ \sum_{e \in a_i \setminus a_i^{\mathrm{opt}}} \ell_e(|a_i, a_{-i}^{\mathrm{opt}}|_e - 1)(|a_i, a_{-i}^{\mathrm{opt}}|_e - 1) \\
&\quad - \sum_{e \in a_i^{\mathrm{opt}} \setminus a_i} \ell_e(|a^{\mathrm{opt}}|_e - 1)(|a^{\mathrm{opt}}|_e - 1) - \sum_{e \in a_i \setminus a_i^{\mathrm{opt}}} \ell_e(|a_i, a_{-i}^{\mathrm{opt}}|_e)|a_i, a_{-i}^{\mathrm{opt}}|_e] \\
= \sum_{e \in \mathcal{E}} \ell_e(|a^{\mathrm{opt}}|_e)|a^{\mathrm{opt}}|_e &- \sum_{e \in \mathcal{E}} \ell_e(|a_i, a_{-i}^{\mathrm{opt}}|_e)|a_i, a_{-i}^{\mathrm{opt}}|_e \\
= \mathrm{MinCost}(G) &- \mathrm{SC}(a_i, a_{-i}^{\mathrm{opt}}),
\end{aligned}
$$

where the third equality holds because we add and subtract $\ell_e(|a^{\mathrm{opt}}|_e)|a^{\mathrm{opt}}|_e$ for all $e \in \mathcal{E} \setminus (a_i^{\mathrm{opt}} \cup a_i)$ and all $e \in a_i^{\mathrm{opt}} \cap a_i$. The final line must be nonpositive for all actions $a_i \in \mathcal{A}_i$ and all users $i \in N$ by the definition of $\mathrm{MinCost}(G)$, concluding the proof. □

In Corollary 7.4.1, we provide a tractable linear program that can be used to compute the Price of Anarchy in congestion games under the marginal cost rule. To do so, one simply computes the agent-cost function corresponding to each basis resource-cost function and the marginal cost taxation rule. Thus, we have established that the local taxation rules that unilaterally minimize the Price of Anarchy and Price of Stability are linear and unique (Proposition 9.1.1 and Theorem 9.1.1), and have provided characterizations of the corresponding Price of Stability and Price of Anarchy, respectively.

9.1.2 The outer trade-off between anarchy and stability

In the previous sections, we showed that the unique local taxation rule that minimizes the Price of Anarchy has corresponding Price of Stability equal to the Price of Anarchy (Theorem 9.1.1). Furthermore, the well-known marginal cost rule is the unique local taxation rule that minimizes the Price of Stability, always achieving a Price of Stability of 1 (Proposition 9.1.1). Since the minimum achievable Price of Anarchy is strictly greater than 1 for all nondecreasing, convex resource costs (except constant) [89], it immediately follows that the taxation rule that minimizes the Price of Anarchy is distinct from the taxation rule that minimizes the Price of Stability and, thus, there must exist a trade-off between these two metrics. In this section, we develop analytical techniques for deriving upper and lower bounds on the Pareto frontier between the Price of Anarchy and Price of Stability in congestion games, which permit us to better understand the trade-off between these two metrics. Though the results in this section depend on an upper-bound, n, on the number of users, we discuss how these techniques can be extended to remove the dependence on n in Section 9.3.

An upper bound. Before presenting our upper bound, we introduce a modified version of the smoothness argument in [96] that provides upper bounds on the Price of Stability of congestion games. Recall that all congestion games are potential games and admit the potential function

$\Phi : \mathcal{A} \to \mathbb{R}$ in (6.6).

Proposition 9.1.2. *Let $\mathcal{G}$ denote any family of congestion games, and suppose that there exist $\zeta > 0$, $\lambda > 0$ and $\mu < 1$ such that, for every game $G \in \mathcal{G}$ and any two assignments $a, a' \in \mathcal{A}$, it holds that*

$$\mathrm{SC}(a) + \sum_{i \in N} [C_i(a_i', a_{-i}) - C_i(a)] + \zeta[\Phi(a') - \Phi(a)] \leq \lambda \mathrm{SC}(a') + \mu \mathrm{SC}(a). \tag{9.1}$$

Then, the Price of Stability satisfies $\mathrm{PoS}(\mathcal{G}) \leq \lambda/(1-\mu)$.

Proof. Consider any game $G \in \mathcal{G}$ and let $a^{\mathrm{opt}} \in \mathcal{A}$ denote an optimal assignment, i.e., $\mathrm{SC}(a^{\mathrm{opt}}) = \mathrm{MinCost}(G)$. Thus, let $a^{\mathrm{ne}} \in \mathrm{NE}(G)$ denote a pure Nash equilibrium that satisfies $\Phi(a^{\mathrm{ne}}) \leq \Phi(a^{\mathrm{opt}})$.[1] Since $C_i(a^{\mathrm{ne}}) \leq C_i(a_i^{\mathrm{opt}}, a_{-i}^{\mathrm{ne}})$ for all $i \in N$ and $\Phi(a^{\mathrm{ne}}) \leq \Phi(a^{\mathrm{opt}})$, it follows from (9.1) that

$$\mathrm{SC}(a^{\mathrm{ne}}) \leq \lambda \mathrm{SC}(a^{\mathrm{opt}}) + \mu \mathrm{SC}(a^{\mathrm{ne}}).$$

Rearranging the above inequality gives us that $\mathrm{SC}(a^{\mathrm{ne}})/\mathrm{SC}(a^{\mathrm{opt}}) \leq \lambda/(1-\mu)$. Since a^{ne} is not necessarily the pure Nash equilibrium in $\mathrm{NE}(G)$ with minimum social cost, it holds that $\mathrm{PoS}(G) \leq \lambda/(1-\mu)$, and it could hold that $\mathrm{PoS}(G) < \lambda/(1-\mu)$ in general. □

The smoothness argument in Proposition 9.1.2 provides an upper bound on the Price of Stability by bounding the efficiency of all pure Nash equilibria with potential lower than the potential at the optimal assignment. Our next result shows how one can leverage this smoothness argument to optimize an upper bound on the Price of Stability under a maximum allowable Price of Anarchy constraint.

Theorem 9.1.2. *Consider the family of resource cost functions $\mathcal{L} = span(b_1, \ldots, b_m)$ corresponding to basis functions $b_1, \ldots, b_m$, and maximum number of users n. Further, consider a maximum allowable Price of Anarchy $\bar{\Pi} \geq \mathrm{MinPoA}(n, \mathcal{L})$. Let $\{F_1^{\mathrm{opt}}, \ldots, F_m^{\mathrm{opt}}\}$, ν^{opt}, ρ^{opt},*

[1] Observe that such a pure Nash equilibrium must always exist since any potential minimizer is a pure Nash equilibrium.

$\gamma^{\rm opt}$, $\kappa^{\rm opt}$ *be solutions to the following:*[2]

$$\begin{aligned}
&\underset{\{F_j\},\nu^{-1}\rho,\gamma,\nu^{-1},\kappa}{\text{maximize}} \quad \gamma \quad \text{subject to:}\\
&\nu^{-1}\rho \geq \bar{\Pi}^{-1}\nu^{-1}, \quad \nu^{-1} \geq 0, \quad \kappa \geq 0\\
&\nu^{-1}b_j(y)y - \nu^{-1}\rho b_j(x)x + (x-z)F_j(x) - (y-z)F_j(x+1) \geq 0,\\
&\qquad\qquad\qquad\qquad \forall (x,y,z) \in \mathcal{I}(n), \forall j \in \{1,\dots,m\},\\
&b_j(y)y - \gamma b_j(x)x + (x-z)F_j(x) - (y-z)F_j(x+1) + \kappa\left[\sum_{k=1}^{x} F_j(k) - \sum_{k=1}^{y} F_j(k)\right] \geq 0,\\
&\qquad\qquad\qquad\qquad \forall (x,y,z) \in \mathcal{I}(n), \forall j \in \{1,\dots,m\},
\end{aligned} \tag{9.2}$$

where we define $b_j(0) = F_j(0) = F_j(n+1) = 0$*. Then, the local taxation rule* $T^{\rm opt}$ *defined as* $T^{\rm opt}(b_j)(x) = F_j^{\rm opt}(x) - b_j(x)$, $j = 1,\dots,m$*, achieves Price of Anarchy* $\mathrm{PoA}(\mathcal{G}^n_{T^{\rm opt}}) = 1/\rho^{\rm opt} \leq \bar{\Pi}$ *and Price of Stability* $\mathrm{PoS}(\mathcal{G}^n_{T^{\rm opt}}) \leq 1/\gamma^{\rm opt}$.

The proof is presented in Section 9.3, and amounts to reformulating the problem of computing the local taxation rule that optimizes the smoothness bound in Proposition 9.1.2 as a tractable optimization problem. The optimization problem in (9.2) is a bilinear program with a single bilinearity, since κ is multiplied with F in the final set of constraints. Such programs can be solved efficiently using, e.g., the method of bisections, which involves solving a finite number of linear programs for appropriate guesses of the value $\kappa^{\rm opt}$.

A possible interpretation of the above result is that the local rule $T^{\rm opt}$ guarantees that every game $G \in \mathcal{G}^n_{T^{\rm opt}}$ has at least one pure Nash equilibrium with social cost at most $1/\gamma^{\rm opt}$ times greater than $\mathrm{MinCost}(G)$. Recall from the proof of Proposition 9.1.2 that this equilibrium may not represent the best performing equilibrium of G, so this represents an upper bound on the Price of Stability, in general.

[2] In (9.2) and forthcoming optimization problems, we reformulate the decision variables and constraints to ensure they follow a (bi)linear program structure for the reader's convenience. For example, $\nu^{-1}\rho$ is a decision variable in (9.2). Note that the corresponding optimal values $\{F_1^{\rm opt},\dots,F_m^{\rm opt}\}$, $\nu^{\rm opt}$, $\rho^{\rm opt}$, etc. are uniquely determined by solutions to these problems.

A lower bound The following theorem states our corresponding lower bound on the best achievable Price of Stability for a maximum allowable Price of Anarchy $\bar{\Pi}$:

Theorem 9.1.3. *Consider the family of resource cost functions $\mathcal{L} = span(b_1, \ldots, b_m)$ corresponding to basis functions $b_1, \ldots, b_m$, and maximum number of users n. Further, consider a maximum allowable Price of Anarchy $\bar{\Pi} \geq \mathrm{MinPoA}(n, \mathcal{L})$. Let $F_j^{\mathrm{opt}}, \nu_j, \rho_j$ be optimal values that solve the following m linear programs (one for each j):*

$$\underset{F,\nu^{-1},\rho\nu^{-1}}{\mathit{maximize}} \sum_{x=1}^{n} F(x) \quad \textit{subject to:}$$

$$\begin{aligned}
&\rho\nu^{-1} \geq \bar{\Pi}^{-1}\nu^{-1}, \quad \nu^{-1} \geq 0, \quad F(1) = 1,\\
&\nu^{-1} b_j(y)y - \rho\nu^{-1} b_j(x)x + (x-z)F(x) - (y-z)F(x+1) \geq 0, \forall (x,y,z) \in \mathcal{I}(n),
\end{aligned} \tag{9.3}$$

where we define $b_j(0) = F_j(0) = F_j(n+1) = 0$. Then, the Price of Stability of any local taxation rule T with $\mathrm{PoA}(\mathcal{G}_T^n) \leq \bar{\Pi}$ must satisfy $\mathrm{PoS}(\mathcal{G}_T^n) \geq \max_j\{1/\gamma_j^{\mathrm{opt}}\}$, where

$$\gamma_j^{\mathrm{opt}} = \min_{0 \leq v < u \leq n} \frac{b(v)v + \sum_{k=1}^{u-v} F_j^{(u,v)}(k)}{b(u)u},$$

where $F_j^{(u,v)}(k) = \max_{v+k \leq x \leq u} F_j^{\mathrm{opt}}(x)$ for $k = 1, \ldots, u - v$.

We highlight some important observations regarding the above result in the discussion below:

Note that the linear program in (9.3) must be feasible for all values $\bar{\Pi} \geq \mathrm{MinPoA}(n, \mathcal{L})$ as there exists at least one set of feasible values F, ν, ρ by Theorem 8.1.1 and Theorem 9.1.1. Furthermore, the linear program must provide a (tight) lower bound of $\mathrm{PoS}(\mathcal{G}_T^n) \geq 1$ for any $\bar{\Pi}$ greater than the Price of Anarchy of the marginal cost rule, $\mathrm{PoA}(\mathcal{G}_{T^{\mathrm{mc}}}^n)$, since the Price of Stability of the marginal cost rule is 1. When the basis functions are convex and nondecreasing, the linear program must also provide a (tight) lower bound $\mathrm{PoS}(T) \geq \bar{\Pi}$ when $\bar{\Pi} = \mathrm{MinPoA}(n, \mathcal{L})$, since we showed in Part (ii) of the proof of Theorem 9.1.1 that a worst case game in this setting has the same structure as the construction we use to obtain this lower bound.[3] Additionally, the game construction from which we obtain this lower bound on the Price of Stability has a unique pure Nash equilibrium a^{ne} where each user $i \in N$ strictly prefers to play a_i^{ne} when users $1, \ldots, i-1$ play their respective actions in a^{ne}. It is straightforward to verify that a^{ne} is also the unique coarse-correlated equilibrium of the game and, thus, our lower bound extends to the best case coarse correlated equilibrium efficiency.

9.1.3 The inner trade-off between anarchy and stability

In the previous sections, we study the Price of Anarchy and Price of Stability of a given family of instances $\mathcal{G}_T$ as independent, worst-case measures of the equilibrium efficiency, i.e., we summarize the equilibrium efficiency of all game instances under a given taxation rule T with only two numbers, $\text{PoA}(\mathcal{G}_T)$ and $\text{PoS}(\mathcal{G}_T)$. Note, however, that the values $\text{PoA}(\mathcal{G}_T), \text{PoS}(\mathcal{G}_T)$ may not be achieved within the same game instance. Specifically, there need not exist a game instance $G \in \mathcal{G}_T$ such that $\text{PoA}(G) = \text{PoA}(\mathcal{G}_T)$ *and* $\text{PoS}(G) = \text{PoS}(\mathcal{G}_T)$. Rather, it could be that there exist two distinct games $G, G' \in \mathcal{G}_T$ satisfying $\text{PoA}(G) = \text{PoA}(\mathcal{G}_T) > \text{PoA}(G')$ and $\text{PoS}(G') = \text{PoS}(\mathcal{G}_T) > \text{PoS}(G)$. This motivates our investigation – in this section – of those Price of Anarchy and Price of Stability pairs that can be achieved within the same game instance, where we wish to understand if considering such *attainable joint performance measure* offers more refined insights on the joint optimization of the worst and best equilibrium efficiency.

More specifically, for a given family $\mathcal{G}$, we aim to capture the dependence of the Price of Stability of an invidual instance on its Price of Anarchy. To that end, for given $\tau \in [1, \text{PoA}(\mathcal{G})]$, we define $\mathcal{G}^\tau$ as

$$\mathcal{G}^\tau := \{G \in \mathcal{G} \text{ s.t. } \text{PoA}(G) = \tau\}. \tag{9.4}$$

Our goal is to characterize how the value $\text{PoS}(\mathcal{G}^\tau)$ evolves with $\tau \in [1, \text{PoA}(\mathcal{G})]$.

Our next result establishes that the tension between the Price of Anarchy and Price of Sta-

[3] Though the game construction we consider to obtain the result in Theorem 9.1.3 is of the same structure as in Figure 9.2, the selection of the resources' coefficients is more nuanced in general since we have no guarantee on the monotonicity of the resource cost function F in this setting. See Section 9.3 for more details.

bility persists under the attainable joint performance measure. This is based on the observation that the independently measured Price of Anarchy and Price of Stability corresponding with taxes that minimize either the Price of Anarchy or the Price of Stability are in fact attained within the same game instance.

Corollary 9.1.1. *For any family of nondecreasing, convex latency functions $\mathcal{L}$, and maximum number of users n, the following statements hold:*

- *Let T^{PoA} denote a taxation rule that minimizes the Price of Anarchy of the corresponding family of instances, i.e.,*
$$T^{\mathrm{PoA}} \in \arg\min_{T} \mathrm{PoA}(\mathcal{G}_T^n). \tag{9.5}$$
There exists an instance $G \in \mathcal{G}^n_{T^{\mathrm{PoA}}}$ such that $\mathrm{PoA}(G) = \mathrm{PoS}(G) = \mathrm{PoA}(\mathcal{G}^n_{T^{\mathrm{PoA}}})$.

- *Let T^{PoS} denote a taxation rule that minimizes the Price of Stability of the corresponding family of instances, i.e.,*
$$T^{\mathrm{PoS}} \in \arg\min_{T} \mathrm{PoS}(\mathcal{G}_T^n). \tag{9.6}$$
There exists an instance $G \in \mathcal{G}^n_{T^{\mathrm{PoS}}}$ such that $\mathrm{PoA}(G) = \mathrm{PoA}(\mathcal{G}^n_{T^{\mathrm{PoS}}})$ and $\mathrm{PoS}(G) = \mathrm{PoS}(\mathcal{G}^n_{T^{\mathrm{PoS}}}) = 1$.

Corollary 9.1.1 establishes that the extreme points of the Price of Anarchy, Price of Stability trade-off curve coincide whether we consider the independent, worst-case performance measure, or the attainable joint performance measure. It remains to be seen whether these two coincide in general, i.e., that the independent, worst-case performance guarantee is always attainable by the same game instance for any family of instances under any taxation rule. To that end, we put forward a modified smoothness condition and game construction to characterize the relation between the Price of Anarchy and Price of Stability under the attainable joint performance measure. We show that there can be a significant separation between the independent, worst-case performance guarantee and the attainable joint performance guarantees.

An upper bound. We obtain an upper bound on the attainable joint performance guarantees using the following smoothness condition, which applies to any family of potential games, where each game has corresponding potential function $\Phi : \mathcal{A} \to \mathbb{R}$:

Proposition 9.1.3. *Given any family of congestion games $\mathcal{G}$ and parameter $\tau \in [1, \mathrm{PoA}(\mathcal{G})]$, suppose that there exist parameters $\kappa, \lambda_1, \lambda_2 \geq 0$, and $\mu, \nu \in \mathbb{R}$ such that, for every game $G \in \mathcal{G}^\tau$ and actions $a, a', a'' \in \mathcal{A}$, it holds that*

$$\sum_{i=1}^{n} \left[\lambda_1 [C_i(a'_i, a_{-i}) - C_i(a)] + \lambda_2 [C_i(a''_i, a_{-i}) - C_i(a)] \right] + \kappa [\Phi(a'') - \Phi(a')] \\ \leq \nu \cdot \mathrm{SC}(a) - \mathrm{SC}(a') + \mu \cdot \mathrm{SC}(a''). \tag{9.7}$$

Then, the Price of Stability satisfies $\mathrm{PoS}(\mathcal{G}^\tau) \leq \mu + \tau\nu$.

Proof. Consider any game $G \in \mathcal{G}^\tau$ and let $a^{\mathrm{opt}} \in \mathcal{A}$ denote an optimal assignment, i.e., $\mathrm{SC}(a^{\mathrm{opt}}) = \mathrm{MinCost}(G)$. Let each $a^{\mathrm{ne},1}, a^{\mathrm{ne},2} \in \mathrm{NE}(G)$ denote a pure Nash equilibrium of G, not necessarily distinct. We let $a^{\mathrm{ne},2}$ be a pure Nash equilibrium that satisfies $\Phi(a^{\mathrm{ne},2}) \leq \Phi(a^{\mathrm{opt}})$. Since $C_i(a^{\mathrm{ne},1}) \leq C_i(a_i^{\mathrm{opt}}, a_{-i}^{\mathrm{ne},1})$ and $C_i(a^{\mathrm{ne},1}) \leq C_i(a_i^{\mathrm{ne},2}, a_{-i}^{\mathrm{ne},1})$ for all $i \in N$, and $\Phi(a^{\mathrm{ne},2}) \leq \Phi(a^{\mathrm{opt}})$, it follows from (9.7) that

$$\mathrm{SC}(a^{\mathrm{ne},2}) \leq \nu \cdot \mathrm{SC}(a^{\mathrm{ne},1}) + \mu \cdot \mathrm{SC}(a^{\mathrm{opt}}).$$

Dividing both sides of the above inequality by $\mathrm{SC}(a^{\mathrm{opt}})$, we obtain

$$\frac{\mathrm{SC}(a^{\mathrm{ne},2})}{\mathrm{SC}(a^{\mathrm{opt}})} \leq \nu \cdot \frac{\mathrm{SC}(a^{\mathrm{ne},1})}{\mathrm{SC}(a^{\mathrm{opt}})} + \mu \leq \mu + \tau\nu,$$

where the final inequality holds since $\mathrm{SC}(a^{\mathrm{ne},1})/\mathrm{SC}(a^{\mathrm{opt}}) \leq \mathrm{PoA}(G) = \tau$ by the definition of $\mathcal{G}^\tau$ from (9.4). Following the same reasoning as in the proof of Proposition 9.1.2, it follows that $\mathrm{PoS}(G) \leq \mu + \tau\nu$. □

Observe that by using this smoothness condition, an upper bound on the attainable joint performance guarantees can be obtained for the family of instances corresponding to any class of latency functions and any taxation rule. In our next result, we use this smoothness argument to derive an upper bound on the attainable joint performance guarantees in affine and quadratic congestion games without taxes:

Corollary 9.1.2. *Consider the family of affine resource cost functions, i.e., $\mathcal{L} = span(b_1, b_2)$ where $b_1(x) = 1$ and $b_2(x) = x$, and let $\mathcal{G}_0$ represent the family of affine congestion games without taxes. It holds that*

$$\mathrm{PoS}(\mathcal{G}_0^\tau) \leq \min\left\{\tau, \mathrm{PoS}(\mathcal{G}_0), \frac{-4}{3}\tau + \frac{13}{3}\right\}, \tag{9.8}$$

for all $\tau \in [1, 5/2]$, where $\mathrm{PoS}(\mathcal{G}_0) == 1 + \sqrt{3}/3 \approx 1.577$. Next, consider the family of quadratic resource cost functions, i.e., $\mathcal{L} = span(b_1, b_2, b_3)$ where $b_1(x) = 1$, $b_2(x) = x$ and $b_3(x) = x^2$, and let $\mathcal{G}_0$ represent the family of quadratic congestion games without taxes. It holds that

$$\mathrm{PoS}(\mathcal{G}_0^\tau) \leq \min\left\{\tau, \mathrm{PoS}(\mathcal{G}_0), \frac{-1}{3}\tau + \frac{151}{36}\right\}, \tag{9.9}$$

for all $\tau \in [1, 115/12]$, where $\mathrm{PoS}(\mathcal{G}_0) \approx 2.361$.

Proof. The proof follows from the smoothness condition in (9.7) by showing that the smoothness parameters $\lambda_1 = \lambda_2 = 1$, $\kappa = 3$, $\mu = 13/3$ and $\nu = -4/3$ are feasible for all affine congestion games without taxes, and that the smoothness parameters $\lambda_1 = 1/8$, $\lambda_2 = 29/72$, $\kappa = 35/9$, $\mu = 151/36$ and $\nu = -1/3$ are feasible for all quadratic congestion games without taxes. □

In Figure 9.3, we plot the upper bound on the attainable joint performance guarantees for affine congestion games without taxes provided in (9.8) (solid black line). Observe that the upper bound demonstrates that the independent, worst-case performance guarantee, ($\mathrm{PoA}(\mathcal{G}_T)$, $\mathrm{PoS}(\mathcal{G}_T)$) – which is $(2.500, 1.577)$ for the family of affine congestion games without taxes – cannot be achieved by any instance. Additionally, and perhaps surprisingly, the upper bound guarantees that any affine congestion game without taxes with worst case Price of Anarchy has Price of Stability equal to 1. Note that these two observations do not necessarily hold for every rule T, as we show in Corollary 9.1.1.

A lower bound. Next, we wish to characterize a lower bound that complements the upper bound that we obtained using the smoothness condition in (9.7). We provide such a lower bound by means of a game construction: Let $a^{\mathrm{w-ne}}, a^{\mathrm{b-ne}}, a^{\mathrm{opt}} \in \mathcal{A}$ respectively denote the worst-case pure Nash equilibrium, best-case pure Nash equilibrium, and optimal joint allocation of a game. Observe that, for $a^{\mathrm{w-ne}}$ to be a pure Nash equilibrium, the following constraints must hold:

$$C_i(a^{\mathrm{w-ne}}) \leq C_i(a_i, a_{-i}^{\mathrm{w-ne}}), \forall a_i \in \mathcal{A}_i, \forall i \in N$$

Furthermore, to ensure that $a^{\mathrm{b-ne}}$ is a pure Nash equilibrium, it is sufficient to impose the constraints:

$$C_i(a_{1:i}^{\mathrm{b-ne}}, a_{i+1:n}) < C_i(a_{1:i-i}^{\mathrm{b-ne}}, a_{i:n}), \forall a \neq a^{\mathrm{w-ne}} \in \mathcal{A}, i \in N,$$

Note that $a^{\mathrm{w-ne}}$ and $a^{\mathrm{b-ne}}$ are the *only* pure Nash equilibria of the game under the imposed user cost structure.

The game construction belongs to a subset of the family of instances $\mathcal{G}_T$ that only contains instances with at most two pure Nash equilibria ($a^{\mathrm{w-ne}}$ and $a^{\mathrm{b-ne}}$), of which $a^{\mathrm{b-ne}}$ is the game's potential minimizer. Observe that by maximizing $\mathrm{SC}(a^{\mathrm{b-ne}})$ while requiring that $\mathrm{SC}(a^{\mathrm{w-ne}}) = \tau \cdot \mathrm{SC}(a^{\mathrm{opt}})$, we can obtain a lower bound on $\mathrm{PoS}(\mathcal{G}_T^\tau)$. Since the constraints we consider impose a particular user cost structure on the games we consider, this lower bound may not necessarily be a tight characterization. Nonetheless, the advantage of this lower bound is that – under an appropriate parameterization – it can be computed via linear programming methods for a given maximum number of users n. We provide the details on such a parameterization and corresponding linear program in Section 9.3 for ease of presentation. In Figure 9.3, we plot the lower bound on $\mathrm{PoS}(\mathcal{G}_T^\tau)$ for affine congestion games without taxes computed for a maximum of $n = 4$ users (solid orange line).

Simulation results We provide a simulation example to compare the independent, worst-case performance guarantee and the attainable joint performance guarantees for the Price of Anarchy and Price of Stability. In our simulation example, we consider the family of affine congestion games without taxes, for which the independent, worst-case guarantee is $(2.500, 1.577)$.

Consider an affine congestion game with $n = 4$ users and $|\mathcal{E}| = 10$ resources. To each of the edges $e \in \mathcal{E}$, assign the resource cost function $\ell_e(x) = \alpha_e \cdot x$, where α_e is sampled independently from the uniform distribution between 0 and 1. Further assign to each user $i \in N$ three actions, each action consisting of the unique resources among two resources drawn (with replacement) uniformly from the set of resources.

We generate 10^5 such random instances of affine congestion games without taxes. For each of these instances, we compute the system cost at all the pure Nash equilibria, as well as the minimum achievable system cost. From these values, we obtain the Price of Anarchy and Price of Stability of each instance. In Figure 9.3, we plot the Price of Anarchy, Price of Stability pair for each of the 10^5 random instances as navy blue '+' marks. Observe that the joint performance of each of the generated instances falls within our bounds on the attainable joint performance guarantees. Furthermore, though the attainable joint performance guarantees of the instances are well below the theoretical worst-case, the distribution of the instances mimicks the shape of our bounds, i.e., games with high Price of Anarchy have low Price of Stability, and vice versa.

9.2 Trade-offs with the Price of Urgency

So far in our discussion of congestion and resource-allocation games, we have considered the efficiency of equilibria either from the worst-case or the best-case perspective as measured by the Price of Anarchy and Price of Stability, respectively. Note, however, that these efficiency guarantees correspond with asymptotic solutions in the sense that convergence guarantees to Nash equilibrium are fairly pessimistic. In fact, in the class of games that we consider, arriving at equilibrium can take an exponential amount of time [99], rendering the resulting approxima-

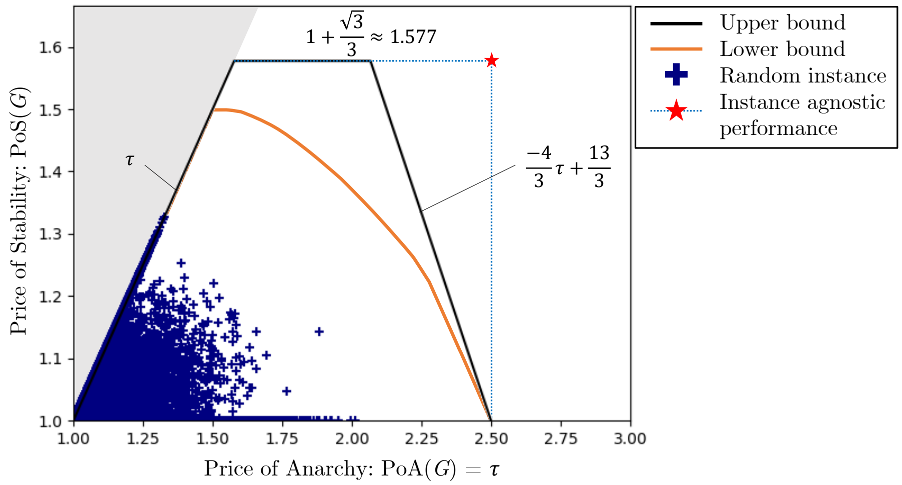

Figure 9.3: *The attainable joint performance guarantees in affine congestion games without taxes.* We plot our upper bound (solid, black line) and lower bound (solid, orange line) on the set of feasible (PoA, PoS) pairs in the family of affine congestion games without taxes. We also plot the (PoA, PoS) of 10^5 randomly generated instances from this family (navy '+' marks), which all fall within the bounds (details on how these instances were generated are provided in the main text). Observe that the upper bound rules out any instances with joint performance equal to, or close to, the independent, worst-case performance guarantee (2.500, 1.577) (red star). Furthermore, our upper and lower bounds coincide at the point (2.50, 1.00), which implies that any worst-case affine congestion game without taxes G from a PoA perspective must satisfy $\text{PoS}(G) = 1$. Finally, although examples of worst-case instances do not arise in the randomly generated instances, their distribution mimicks the shape of our bounds, i.e., high PoA corresponds with low PoS, and vice versa.

tion guarantees irrelevant in many realistic multi-agent scenarios. For example, there may be an extremely large number of agents in the multi-agent system or the relevant situational parameters may be time-varying and volatile or there may be computational and run-time restrictions on the agents. In these instances, it is important to consider the system performance in the transient (i.e., within the time before a Nash equilibrium is reached), which is a perspective that is especially relevant in control theory.

To study the transient performance, we must fix the model of the agents' transient behaviour. One classic distributed algorithm for computing pure Nash equilibria is the best-response algorithm, where the agents are ordered sequentially, and at each step of the execution, a single agent best responds (i.e., optimizes her local utility unilaterally given the previous agents' decisions in the sequence). In general, the best-response algorithm is not guaranteed to find the globally optimal solution, but *is* guaranteed (under mild conditions) to converge to a pure Nash equilibrium in any congestion or resource allocation game if the sequence is repeated a sufficient number of times. Interestingly, in certain well-structured domains, researchers have derived guarantees on the approximation ratio of such algorithms; including set covering problems [116], extendible problems [117] and submodular maximization problems [118].

In this section, we study the transient behavior of the best-response algorithm in resource allocation games by imposing limitations on the time complexity of the sequential process. Specifically, we consider the agents' learning process to follow the κ *round-robin best-response algorithm* and study the approximation guarantees that result from various designs of utility functions in the context of the well-studied class of *resource allocation games.* We also compare the approximation guarantees achieved against the corresponding Price of Anarchy guarantees. This comparison is warranted because the set of pure Nash equilibria contains all globally attractive, reachable states of the κ *round-robin best-response algorithm* in any resource allocation game.

9.2.1 Model

We consider the class of resource-allocation games, and wish to measure the transient performance under the class of *(round-robin) best response processes*, where a certain agent (out of n agents) performs a best response in a round-robin fashion. For a given joint action $\bar{a} \in \mathcal{A}$, we say the action $\hat{a}_i$ is a *best response* for agent i if

$$\hat{a}_i \in \mathrm{BR}(\bar{a}_{-i}) = \arg\max_{a_i \in \mathcal{A}_i} U(a_i, \bar{a}_{-i}), \tag{9.10}$$

where BR may be non-singleton. We also assume that the best response process begins with none of the resources being utilized by any of the agents, denoted by the agents selecting the null joint action $\emptyset$ at time 0. The underlying algorithm is formalized in Algorithm 1.

To arrive at non-trivial efficiency guarantees, we define a *κ-round walk* as a best response process in which Algorithm 1 is run for $T = \kappa \cdot n$ steps. Here, the set of agents perform best

Algorithm 1 Best Response Process

Require: $a(0) \leftarrow \emptyset$, $\tau \leftarrow 0$, T
  while $\tau \leq T$ **do**
    $i \leftarrow (\tau + 1) \mod n$ ▷ Next agent is selected.
    $a_i(\tau + 1) \leftarrow \hat{a}_i \in \mathrm{BR}(a_{-i}(\tau))$ ▷ i best responds.
    $a_{-i}(\tau + 1) \leftarrow a_{-i}(\tau)$ ▷ No other agent moves.
    $\tau \leftarrow \tau + 1$
  end while

responses in succession κ times. We assume that during a κ-round walk, agent i selects its best response $a_i(\tau + 1)$ arbitrarily from $\mathrm{BR}(a_{-i}(\tau))$ if it is not unique. This induces a set of possible action trajectories of the form $(a(0) = \emptyset, a(1) \ldots, a(\kappa n - 1), a(\kappa n))$ selected by the agents throughout the best response process. The potential solution set that occurs after the agents run a κ-round walk is denoted by $\mathrm{sol}(\kappa) \subset \mathcal{A}$ with

$$\mathrm{sol}(\kappa) = \{a(\kappa n) \text{ for each trajectory starting at } a(0) = \emptyset\}.$$

The worst achievable efficiency at the end of the κ-round walk with respect to the best achievable system welfare is defined by the following competitive ratio, which we term *Price of Urgency*:

$$\mathrm{PoU}(G; \kappa) = \frac{\min_{a \in \mathrm{sol}(\kappa)} W(a)}{\max_{a \in \mathcal{A}} W(a)} \in [0, 1]. \tag{9.11}$$

Note that Price of Urgency closer to 1 implies the worst case efficiency after κ rounds is closer to optimal. We additionally extend the efficiency measure to a set of games $\mathcal{G}$ as

$$\mathrm{PoU}(\mathcal{G}; \kappa) = \inf_{G \in \mathcal{G}} \mathrm{PoU}(G; \kappa). \tag{9.12}$$

We wish to derive agent utility rules that optimize the Price of Urgency. More specifically, for a given class of welfare rules $\mathcal{W}$ and number of rounds $\kappa \geq 1$, the main goal is to characterize optimal performance guarantees of the form

$$\mathrm{PoU}^*(\mathcal{W};\kappa) = \sup_{\mathcal{U}:\mathcal{W}\to\mathbb{R}^{\mathbb{N}}_{\geq 0}} \mathrm{PoU}(\mathcal{G};\kappa). \tag{9.13}$$

We will explicitly address how these optimal efficiency guarantees change as a function of the set of possible welfare rules $\mathcal{W}$ as well as the number of rounds κ.

9.2.2 Optimal transient performance

Our first set of results characterizes the attainable performance guarantees for a single round of best response process. We focus on the performance of such "one-round walks" to describe the quickest non-trivial guarantees that can occur under the round-robin best response process, as each agent is required to perform only one best response to arrive at the resulting joint action $a \in \mathrm{sol}(1)$. Thus, we derive the one-round guarantees through a linear program construction that is a function of both the set of allowable welfare rules $\mathcal{W}$ and the utility rules $\mathcal{U}$. Moreover, we will restrict attention to welfare rules that are generated by the span of a given set of basis welfare rules.

Our characterization of the best achievable efficiency guarantees for a one-round walk through a utility design $\mathcal{U}$ is as follows:

Theorem 9.2.1. *Suppose that $\mathcal{W}$ is spanned from a set of welfare rules $\{w^1,\ldots,w^m\}$ where each w^j is submodular. Then the optimal efficiency guarantees achievable with a one-round best response process is given by*

$$\mathrm{PoU}^*(\mathcal{W};1) = \min_{1\leq j\leq m} \frac{1}{\beta^j} \tag{9.14}$$

where $\beta^j \in \mathbb{R}_{\geq 0}$ is the solution to the following program.

$$(u^j,\beta^j) \;\in\; \arg\min_{\beta,u\in\mathbb{R}^{\mathbb{N}}_{\geq 0}} \quad \beta \quad \textit{subject to:} \tag{9.15}$$

$$\beta w^j(y) \geq \sum_{i=1}^{y} u(i) - zu(y+1) + w^j(z) \quad \forall y,z \geq 1,$$

where we take $u^j(1) = 1$ *and* $y, z \in \mathbb{N}$. *Furthermore, a utility design* $\mathcal{U}$ *that achieves this optimal efficiency guarantee is linear and of the form*

$$\mathcal{U}\left(w = \sum_{j=1}^{m} \alpha^j w^j\right) = \sum_{j=1}^{m} \alpha^j u^j, \tag{9.16}$$

where u^j *is the corresponding solution in* (9.15).

The above theorem sets forth a prescriptive process by which to characterize the optimal efficiency guarantees achievable within a one-round best response process. [4] Acquiring β^j through the program in (9.15) may be computationally infeasible in general; however, by considering certain structured classes of welfare rules, we can derive closed form expressions for the one round performance guarantees. We therefore consider a natural restriction of submodular welfare rules centered around the idea of *curvature*, which is defined below.

Definition 9.2.1 (Curvature)**.** *A submodular welfare rule* w *has a curvature of* $c \in [0,1]$ *if* $c = 1 - \lim_{n\to\infty}(w(n+1) - w(n))/w(1)$.

In this sense, curvature characterizes the rate of diminishing returns associated with a welfare rule w.With this, we can arrive at a tight, closed-form characterization of the optimal one-round performance guarantees, as shown below.

Theorem 9.2.2. *Let the set* $\mathcal{W}$ *comprise of all welfare rules* w *such that* $w \in \mathbb{R}^{\mathbb{N}}_{>0}$ *has a curvature of at most* c. *Then the optimal efficiency guarantees achievable with a one-round best response process satisfies*

$$\mathrm{PoU}^*(\mathcal{W}; 1) = 1 - \frac{c}{2}\left(\geq \frac{1}{2}\right). \tag{9.17}$$

The optimal utility design that achieves the above efficiency guarantee also has a closed form expression, which can be found in the proof. The results in Theorem 9.2.2 suggests that, under the optimal utility design, running best response processes for these classes of games can result

[4] While we mostly focus on the class of submodular welfare rules in this paper, the linear program in (9.15) can be extended to consider other classes.

in the agents coordinating to a high quality joint action very quickly. If the curvature c is close to 0, we can even arrive at an approximation guarantee of nearly 1 after only a one-round walk.

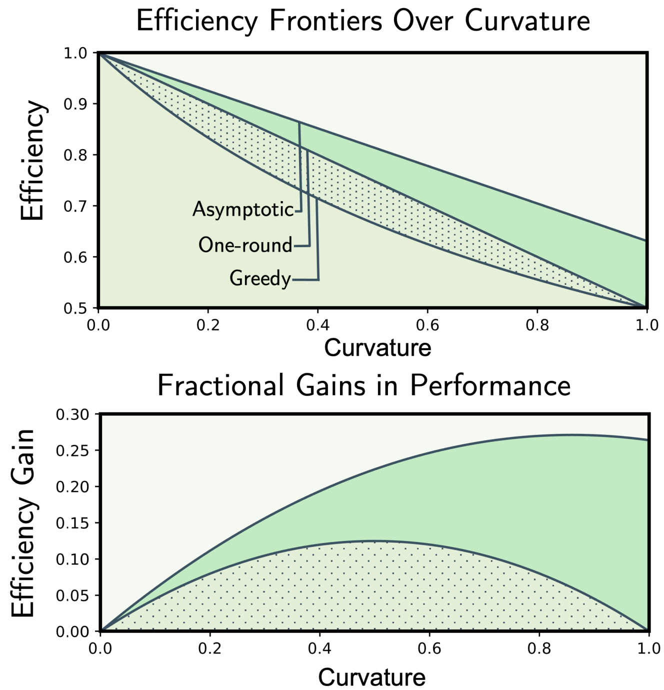

Figure 9.4: In the top figure, we visually depict the efficiency guarantees of Theorem 9.2.2 with respect to the optimal asymptotic guarantees. Additionally, the fractional gains in the performance when moving from the greedy solution to the optimal one-round and the asymptotic solutions are depicted in the bottom figure.

9.2.3 Muliple round walks

We now extend to κ-round walks, and study the resulting efficiency guarantees. Allowing the best response process to continue for more than $\kappa = 1$ rounds may appear to be a natural avenue to increase the performance guarantees. However, in the next theorem, we show that further rounds do not increase the relative efficiency guarantees. Specifically, with regards to the set of welfare rules of a certain curvature, we derive an upper bound for the efficiency of κ-round walk that exactly matches the efficiency guarantee of the one-round walk.

Theorem 9.2.3. *Let the set $\mathcal{W}$ comprise of all welfare rules w such that $w \in \mathbb{R}^{\mathbb{N}}_{>0}$ has a curvature of at most c. Then the efficiency guarantees of the optimal utility design with a*

κ-round best response process, for any $\kappa \geq 1$, is respectively upper bounded by

$$\mathrm{PoU}^*(\mathcal{W};\kappa) \leq 1 - \frac{c}{2}. \tag{9.18}$$

Notably, for any curvature $c \in [0,1]$, the above efficiency guarantees suggest that running the best response process for more than one round does not necessarily increase the performance. We also remark that the results in Theorem 9.2.3 is not endemic to the specific dynamics we consider in this paper. Allowing for different order of play apart from round-robin does not affect the resulting upper bounds. This is further elaborated on in Section 9.3. Therefore, in general, this suggests stark diminishing returns for running the best response process for further rounds.

9.2.4 The tradeoff between anarchy and urgency

In the next resut, we characterize the Price of Anarchy and Price of Urgency of the utility designs that maximize the Price of Anarchy and Price of Urgency.

Theorem 9.2.4. *Let $\mathcal{W}$ denote the set of all possible submodular welfare rules, U_{PoA} denote the utility design that maximizes the Price of Anarchy, and U_1^* denote the utility design that maximizes the efficiency guarantees of the one-round walk. Then the efficiency guarantees with a one-round best response process for both utility designs are*

$$\mathrm{PoU}(\mathcal{G}_{\mathcal{W},U_{\mathrm{PoA}}};1) = 0 \qquad \mathrm{PoU}(\mathcal{G}_{\mathcal{W},U_1^*};1) = \frac{1}{2}. \tag{9.19}$$

Furthermore, the Price of Anarchy guarantees of both utility designs are respectively

$$\mathrm{PoA}(\mathcal{G}_{\mathcal{W},U_{\mathrm{PoA}}}) = 1 - \frac{1}{e} \qquad \mathrm{PoA}(\mathcal{G}_{\mathcal{W},U_1^*}) = \frac{1}{2}. \tag{9.20}$$

We observe that while the asymptotic guarantees of U_1^* are equivalent to the corresponding transient guarantees, the transient guarantees of U_{PoA} unexpectedly degrade to 0. Interestingly, optimizing for asymptotic performance does not necessarily translate to good transient performance in our setting. To clarify the interplay between the transient and asymptotic guarantees, we would like to characterize the exact Pareto optimal frontier for these guarantees. While calculating this trade-off frontier is difficult to do in general, we restrict our analysis to the specific subset of resource allocation games known as *set covering games* [112] to arrive at an exact trade-off curve. Set covering games are characterized by the following welfare rule.

$$W^{\rm sc}(j) = \left\{ \begin{array}{ll} 1, & \text{for } j \geq 1 \\ 0, & \text{for } j = 0 \end{array} \right\}. \tag{9.21}$$

With this, we arrive at the following Pareto frontier characterization, depicted in Figure 9.5. Note that the end points of the trade-off curve matches (9.19) and (9.20) exactly.

Theorem 9.2.5. *Let* $\mathcal{W} = \{W^{\rm sc}\}$*, where* $W^{\rm sc}$*, defined in* (9.21)*, is the set covering welfare rule and* $\mathcal{U} = \{U\}$ *is the corresponding utility rule. Under the constraint that* $\mathrm{PoA}(\mathcal{G}_{W^{\rm sc},U}) = Q \in [\frac{1}{2}, 1 - \frac{1}{e}]$*, the optimal* $\max_U \mathrm{PoU}(\mathcal{G}_{W^{\rm sc},U}; 1)$ *is*

$$\left[\sum_{j=0}^{\infty} \max \left\{ j!(1 - \frac{1-Q}{Q} \sum_{\tau=1}^{j} \frac{1}{\tau!}), 0 \right\} + 1 \right]^{-1}. \tag{9.22}$$

Notably in Figure 9.5, we see a stark drop-off in transient guarantees when the Price of Anarchy is close to $1 - 1/e$. This extreme trade-off prompts a more careful interpretation of asymptotic results, especially in the setting of resource allocation games.

9.3 Chapter proofs

9.3.1 Proofs from Section 9.1

Proof of Theorem 9.1.1 We prove Statements i) and ii) of the claim separately, below:

Proof of Statement i). Given any rule T (not necessarily linear), we show that there exists some linear rule $T^{\rm lin}$ satisfying $\mathrm{PoA}(\mathcal{G}^n_{T^{\rm lin}}) \leq \mathrm{PoA}(\mathcal{G}^n_T)$. The linear rule we consider is generated from

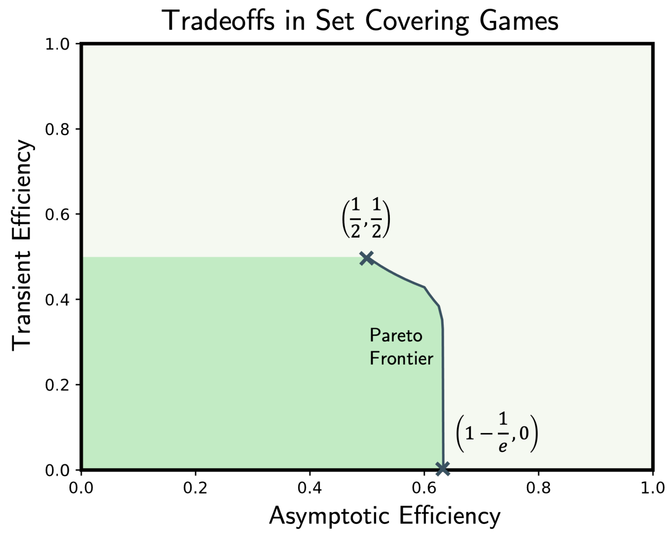

Figure 9.5: We depict the Pareto-optimal frontier of the one-round efficiency $(\text{PoU}(\mathcal{G}_{W^{\text{sc}},U}; 1))$ versus the asymptotic efficiency guarantees $(\text{PoA}(\mathcal{G}_{W^{\text{sc}},U}))$ that are possible with regards to the class of set-covering games. We note that the severe drop off in transient efficiency that results from optimizing the asymptotic efficiency.

the taxes $T(b_j)$, $j = 1, \ldots, m$, as follows: $T^{\text{lin}}(\sum_{j=1}^{m} \alpha_j b_j) = \sum_{j=1}^{m} \alpha_j T(b_j)$. As the same set of arguments also hold for the Price of Stability, the statement follows.

Let $\bar{\mathcal{G}}_T^n$ and $\bar{\mathcal{G}}_{T^{\text{lin}}}^n$ be the restricted families of congestion games with a maximum of n users in which every resource e has resource cost $\ell_e \in \{b_1, \ldots, b_m\}$. Within this restricted class of games, the Price of Anarchy of the local rule T must be equal to that of T^{lin} since the resulting taxes are equivalent. For linear rules such as T^{lin}, one can show that for any congestion game $G \in \mathcal{G}_{T^{\text{lin}}}^n$ there is another game $G' \in \bar{\mathcal{G}}_{T^{\text{lin}}}^n$ (possibly with many more resources) that has arbitrarily close Price of Anarchy following the proof of Theorem 5.6 in [100]. In other words, the Price of Anarchy of $\mathcal{G}_{T^{\text{lin}}}^n$ is equal to the Price of Anarchy of $\bar{\mathcal{G}}_{T^{\text{lin}}}^n$. Meanwhile, for general local rules such as T, we observe that the Price of Anarchy achieved within the restricted class of games $\bar{\mathcal{G}}_T^n$ can only be less than or equal to the Price of Anarchy achieved within $\mathcal{G}_T^n$. It immediately follows that $\text{PoA}(\mathcal{G}_{T^{\text{lin}}}^n) \leq \text{PoA}(\mathcal{G}_T^n)$.

Proof of Statement ii). Consider the following m linear programs:

$$\begin{aligned} \underset{F,\rho}{\text{maximize}} \quad & \rho \\ \text{subject to:} \quad & b_j(y)y - \rho b(x)x + \min\{x, n-y\}F(x) - \min\{y, n-x\}F(x+1) \geq 0 \\ & \qquad \forall (x,y) \in \{0,\ldots,n\} \times \{1,\ldots,n\} \cup (n,0). \end{aligned} \tag{9.23}$$

Observe that the above linear program is a relaxation of the linear program in (8.3) where we only consider the constraints $(x, y, z) \in \mathcal{I}(n)$ such that $(x, y) \in \{0, \ldots, n\} \times \{1, \ldots, n\} \cup (n, 0)$ and $z = \max\{0, x+y-n\}$. Reference [90] provides an expression for a set of optimal solutions $(F_j^{\text{opt}}, \rho_j^{\text{opt}})$, $j = 1, \ldots, m$, to the m linear programs above and show that these are also optimal solutions of the m linear programs in (8.3). As part of their proof, they show that the functions F_j^{opt} must be nondecreasing.

The rest of the proof is shown in two steps: a) we show that for the solutions $(F_j^{\text{opt}}, \rho_j^{\text{opt}})$, $j = 1, \ldots, m$, to the m linear programs in (8.3), the functions $F_1^{\text{opt}}, \ldots, F_m^{\text{opt}}$ are unique (up to rescaling); b) leveraging the fact that the functions $F_1^{\text{opt}}, \ldots, F_m^{\text{opt}}$ are nondecreasing, we construct a congestion game G that has $\text{PoS}(G) = \max_j\{1/\rho_j^{\text{opt}}\}$.

Part iia) – Proof that F_j^{opt} is the unique optimal solution. We must show that there is no other function F that yields a value of $\rho = \rho^{\text{opt}}$. By contradiction, let us assume that there exists a function $\hat{F}$ different from F_j^{opt} that also achieves ρ^{opt}. Let $k+1$ be the first index at which $\hat{F}(k+1) \neq F_j^{\text{opt}}(k+1)$. If $k = 0$, due to the constraint corresponding to $(x = 0, y = 1)$ in the linear program in (9.23), it holds that $\hat{F}(1) \leq b(1) = F_j^{\text{opt}}(1)$. Since $\hat{F}(1) \neq F_j^{\text{opt}}(1)$, it must hold that $\hat{F}(1) < F_j^{\text{opt}}(1)$. A similar argument holds for $k > 0$, since

$$\hat{F}(k+1) \leq \max_{y \in \{1,\ldots,n\}} \frac{b(y)y - \rho^{\text{opt}} b(k)k + \min\{k, n-y\}\hat{F}(k)}{\min\{y, n-k\}} = F_j^{\text{opt}}(k+1),$$

where the equality holds since $\hat{F}(k) = F_j^{\text{opt}}(k)$, by assumption. In short, at the first $k+1$ where $\hat{F}$ does not equal F_j^{opt}, the former is always strictly lower than the latter or else a constraint in the linear program would be violated. The contradiction follows from the constraints with

$(x = n, y = y_n^*)$:

$$\rho^{\text{opt}} \leq \frac{b(y_n^*)y_n^* + (n - y_n^*)\hat{F}(n)}{b(n)n} < \frac{b(y_n^*)y_n^* + (n - y_n^*)F_j^{\text{opt}}(n)}{b(n)n} = \rho_j^{\text{opt}},$$

where the strict inequality holds since $n - y_n^* > 0$. Observe that if $n - y_n^* = 0$, it holds that $\rho_j^{\text{opt}} = 1$, which violates the stated conditions for uniqueness in the theorem statement.

Part iib) – Game construction. Here we show that for each function in $\{F_j^{\text{opt}}\}$, we can construct a congestion game that has Price of Stability equal to $1/\rho_j^{\text{opt}}$. Without loss of generality, we assume that that all basis functions $b_1, \ldots, b_m$ are scaled such that $F_j(1) = 1$. Consider the active constraint correponding to $x = n$ for each F_j^{opt}, which – after some rearrangement – appears as follows:

$$\frac{1}{\rho_j^{\text{opt}}} = \max_{y \in \{0,1,\ldots,n\}} \frac{b_j(n)n}{(n-y)F_j^{\text{opt}}(n) + b_j(y)y}. \tag{9.24}$$

Define $y_n^* \in \{0, 1, \ldots, n\}$ as an argument that maximizes the right-hand side in the above expression. Consider a congestion game G with a set of n users $N = \{1, \ldots, n\}$ and $n - y_n^* + 1$ resources $\mathcal{E} = \{e_0, e_1, \ldots, e_{n-y_n^*}\}$. The users' action sets are defined as follows: Each user $i \in \{1, \ldots, n - y_n^*\}$ has action set $\mathcal{A}_i = \{a_i^{\text{ne}}, a_i^{\text{opt}}\}$ where $a_i^{\text{ne}} = \{e_0\}$ and $a_i^{\text{opt}} = \{e_i\}$; and, each user $i \in \{n - y_n^* + 1, \ldots, n\}$ has $\mathcal{A}_i = \{a_i = \{e_0\}\}$. The cost on resource e_0 is b_j, whereas each $e \in \{e_1, \ldots, e_{n-y_n^*}\}$ has cost $[F_j^{\text{opt}}(n) + \epsilon] \cdot b_j$ for some $\epsilon > 0$. Since $F_j^{\text{opt}}(1) = 1$ and F_j^{opt} is nondecreasing, it is straightforward to verify that the assignment $(a_1^{\text{ne}}, \ldots, a_{n-y_n^*}^{\text{ne}}, a_{n-y_n^*+1}, \ldots, a_n)$ is the unique pure Nash equilibrium of the game. Simply observe that for any assignment $a \in \Pi_i \mathcal{A}_i$, any user $i \in \{1, \ldots, n - y_n^*\}$ selecting its action a_i^{opt} can decrease its cost by selecting its action a_i^{ne} instead, since $F_j^{\text{opt}}(|a|_{e_0}) < F_j^{\text{opt}}(n) + \epsilon$. Thus, the constructed game has a unique pure Nash equilibrium, with system cost $b_j(n)n$. The assignment $(a_1^{\text{opt}}, \ldots, a_{n-y_n^*}^{\text{opt}}, a_{n-y_n^*+1}, \ldots, a_n)$ has system cost $(n - y_n^*)[F_j^{\text{opt}}(n) + \epsilon] + b_j(y_n^*)y_n^*$. Thus, taking the limit as $\epsilon \to 0^+$, the Price of Stability of the constructed game satisfies $\text{PoS}(G) \geq b_j(n)n/[(n-y_n^*)F(n)+b_j(y_n^*)y_n^*] = 1/\rho_j^{\text{opt}}$. Since the function F_j^{opt} has corresponding Price of Anarchy guarantee of $1/\rho_j^{\text{opt}}$, the Price of Stability must also be upper-bounded by

$1/\rho_j^{\text{opt}}$. Thus, $\text{PoS}(\mathcal{G}_T^n) = \max_j\{1/\rho_j^{\text{opt}}\} = \text{PoA}(\mathcal{G}_T^n)$, concluding the proof.

Proof of Theorem 9.1.2 Observe that the optimal upper bound achievable from the smoothness argument in Proposition 9.1.2 can be computed as the solution to the following fractional program:

$$\inf_{\zeta>0,\lambda>0,\mu} \left\{ \frac{\lambda}{1-\mu} \text{ s.t. } (\zeta,\lambda,\mu) \text{ satisfy (9.1) } \forall a,a' \in \mathcal{A}, \forall G \in \mathcal{G}_T^n \right\}.$$

To reduce the number of constraints, we introduce the following parameterization of any pair of assignments $a, a' \in \mathcal{A}$ in a game $G \in \mathcal{G}_T^n$: Consider each resource $e \in \mathcal{E}$ and recall that $\ell_e(x) = \sum_{j=1}^m \alpha_{e,j} \cdot b_j(x)$ with $\alpha_{e,j} \geq 0$ for all j. Let $x_e = |a|_e$, $y_e = |a'|_e$ and $z_e = |\{i \in N : e \in a_i\} \cap \{i \in N : e \in a_i'\}|$. It follows that (x_e, y_e, z_e) belongs to the set $\mathcal{I}(n)$ of all triplets $(x,y,z) \in \mathbb{N}^3$ that satisfy $1 \leq x+y-z \leq n$ and $z \leq \min\{x,y\}$. We define parameters $\theta(x,y,z,j) = \sum_{e \in \mathcal{E}_{x,y,z}} \alpha_{e,j}$ where $\mathcal{E}_{x,y,z} = \{e \in \mathcal{E} \text{ s.t. } (x_e,y_e,z_e) = (x,y,z)\}$. Under this parameterization, observe that the inequality in (9.1) can be rewritten as

$$\sum_{j=1}^m \sum_{x,y,z} \left[b_j(x)x + (y-z)F_j(x+1) - (x-z)F_j(x) + \zeta \left[\sum_{k=1}^x F_j(j) - \sum_{k=1}^y F_j(k) \right] \right] \theta(x,y,z,j)$$
$$\leq \sum_{j=1}^m \sum_{x,y,z} \left[\lambda b_j(y)y + \mu b_j(x)x \right] \theta(x,y,z,j).$$

We note that any (ζ,λ,μ) that satisfies the above constraint for each individual summand corresponding with the triplets $(x,y,z) \in \mathcal{I}(n)$ and $j = 1,\dots,n$ must satisfy the smoothness definition in Propositon 9.1.2 since we have shown that the inequalities governing the smoothness definition are a linear combination over the $|\mathcal{I}(n)| \times m$ summands with nonnegative coefficients $\theta(x,y,z,j)$. Based on this observation, we obtain the following linear program for computing an upper bound on the Price of Stability after the change of variables $\gamma = (1-\mu)/\lambda$,

$\nu = 1/\lambda$ and $\kappa = \zeta/\lambda$:

$$\underset{\gamma,\nu\geq 0,\kappa\geq 0}{\text{maximize}} \quad \gamma \quad \text{subject to:}$$

$$b_j(y)y - \gamma b_j(x)x + \nu[(x-z)F_j(x) - (y-z)F_j(x+1)] + \kappa\left[\sum_{k=1}^{x} F_j(k) - \sum_{k=1}^{y} F_j(k)\right] \geq 0,$$
$$\forall (x,y,z) \in \mathcal{I}(n), \forall j \in \{1,\dots,m\}.$$

Under this change of variables, it holds that $\mathrm{PoS}(\mathcal{G}_T^n) \leq 1/\gamma^{\mathrm{opt}}$ for optimal solutions (γ^{opt}, ν^{opt}, κ^{opt}). We note that the 'inf' objective can now be written as a 'maximize' since $\gamma \in [0,1]$ must hold.

We are interested in obtaining an upper bound on the best Price of Stability that can be achieved by introducing taxation rules. By including the functions F_j, $j = 1,\dots,m$, as decision variables in the dual program, we obtain the following bilinear program for computing a local taxation rule that minimizes the upper bound on the Price of Stability:

$$\underset{\{F_j\},\gamma,\kappa\geq 0}{\text{maximize}} \quad \gamma \quad \text{subject to:}$$

$$b_j(y)y - \gamma b_j(x)x + (x-z)F_j(x) - (y-z)F_j(x+1) + \kappa\left[\sum_{k=1}^{x} F_j(k) - \sum_{k=1}^{y} F_j(k)\right] \geq 0,$$
$$\forall (x,y,z) \in \mathcal{I}(n), \forall j \in \{1,\dots,m\}.$$

Then, for optimal solution $(\{F_j^{\mathrm{opt}}\}, \gamma^{\mathrm{opt}}, \kappa^{\mathrm{opt}})$, the taxation rule T^{opt} defined as $T^{\mathrm{opt}}(b_j)(x) = F_j^{\mathrm{opt}}(x) - b_j(x)$ for all j and x satisfies $\mathrm{PoS}(T^{\mathrm{opt}}) \leq 1/\gamma^{\mathrm{opt}}$. In the above bilinear program, we have imposed $\nu = 1$, which removes one bilinearity in the constraints. The only remaining bilinearity involves the decision variable κ and the functions $F_1,\dots,F_m$.

To obtain the local taxation rule T^{opt} that guarantees a particular Price of Anarchy $\bar{\Pi}$ while minimizing the upper bound on the Price of Stability, we add the constraints for the Price of Anarchy from (8.3) to the bilinear program. We require that $\bar{\Pi}$ be greater than or equal to $\mathrm{MinPoA}(n,\mathcal{L})$ for feasibility. We can then simultaneously minimize the upper bound on the

Price of Stability while guaranteeing the desired Price of Anarchy. After some rearrangement of decision variables, we obtain the bilinear program in the claim.

Proof of Theorem 9.1.3 Consider the set of games $\mathcal{G}_T^n$ with at most n users, family of latency functions $\mathcal{L} = \text{span}(b_1, \ldots, b_m)$ under basis functions $b_1, \ldots, b_m$ and local taxation rule T. Define $F_j(x) = b_j(x) + T(b_j)(x)$ for $x = 1, \ldots, n$, $j = 1, \ldots, m$. Without loss of generality, we normalize such that $F_j(1) = 1$ for $j = 1, \ldots, m$. Define a game $G \in \mathcal{G}$ with u users and $u - v + 1$ resources for v such that $0 \leq v < u \leq n$. We denote the user set as $N = \{1, \ldots, u\}$ and the resource set as $\mathcal{E} = \{e_0, e_1, \ldots, e_{u-v}\}$. The users' action sets are defined as follows: Each user $i \in \{1, \ldots, u-v\}$, has action set $\mathcal{A}_i = \{a_i^{\text{ne}}, a_i^{\text{opt}}\}$ with $a_i^{\text{ne}} = \{e_0\}$ and $a_i^{\text{opt}} = \{e_i\}$, while each user $i \in \{u - v + 1, u\}$ has action set $\mathcal{A}_i = \{a_i\}$ with $a_i = \{e_0\}$. Resource e_0 has resource cost function $\ell_0(x) = b_j(x)$, while each resource e_k, $k = 1, \ldots, u - v$, has resource cost function $\ell_k(x) = \alpha_k b_j(x)$ where $\alpha_k = \max_{v+k \leq x \leq u} F_j(x) + \epsilon$ for $\epsilon > 0$.

Next, we prove that the game G as defined above has a unique pure Nash equilibrium which corresponds with the assignment $a^{\text{ne}} = (a_1^{\text{ne}}, \ldots, a_{u-v}^{\text{ne}}, a_{u-v+1}, \ldots, a_u)$. Consider the choices of user $k \in \{1, \ldots, u - v\}$ with respect to any assignment in which all users $i \in \{1, \ldots, k - 1\}$ play the action a_i^{ne}. The remaining users $i \in \{k + 1, \ldots, u - v\}$ play either of their actions in $\mathcal{A}_i$. Observe that user k must select either the resource e_0 which is currently selected by at least $k + l - 1$ users, or the resource e_k which is currently not selected by any other user. It follows that user k selects $a_i^{\text{ne}} = \{e_0\}$ in this scenario, since $F_j(y) < \max_{v+k \leq x \leq u} F_j(x) + \epsilon$ with $\epsilon > 0$, for $y = v + k, \ldots, u$. Note that, starting from any assignment $a \in \mathcal{A}$, one can repeat this argument from user $k = 1$ to user $k = u - v$ to show that any sequence of best responses will settle on the assignment a^{ne} and, thus, that this is the unique pure Nash equilibrium. Note that the system cost associated with this assignment is $\text{SC}(a^{\text{ne}}) = b_j(u)u$. Meanwhile, the system cost of the assignment $a^{\text{opt}} = (a_1^{\text{opt}}, \ldots, a_{u-v}^{\text{opt}}, a_{u-v+1}, \ldots, a_u)$ is $\text{SC}(a^{\text{opt}}) = b_j(v)v + \sum_{k=1}^{u-v}[\max_{v+k \leq x \leq u} F(x) + \epsilon]$. Furthermore, it holds that $\text{MinCost}(G) \leq \text{SC}(a^{\text{opt}})$. Thus, for

$\epsilon \to 0^+$, the Price of Stability satisfies

$$\mathrm{PoS}() \geq \frac{b_j(u)u}{\mathrm{MinCost}(G)} \geq \frac{b_j(u)u}{b_j(v)v + \sum_{k=1}^{u-v}[\max_{v+k\leq x\leq u} F_j(x)]}.$$

We have shown that within the family of games $\mathcal{G}_T^n$, there exists a singleton game with a unique pure Nash equilibrium for any b_j, $j = 1, \ldots, m$, and any pair (u, v) such that $0 \leq v < u \leq n$. We also derived the lower bound on the Price of Stability for each of these games. Observe that the maximum value of this lower bound over all b_j and all valid pairs (u, v) represents a lower bound on the Price of Stability, i.e.,

$$\begin{aligned}\mathrm{PoS}(\mathcal{G}_T^n) &\geq \max_j \max_{0\leq v<u\leq n} \frac{b_j(u)u}{b_j(v)v + \sum_{k=1}^{u-v}[\max_{v+k\leq x\leq u} F_j(x)]}\\ &= \max_j \max_{0\leq v<u\leq n} \frac{b_j(u)u}{b_j(v)v + \sum_{k=1}^{u-v} F_j^{(u,v)}(k)},\end{aligned}$$

where we define $F_j^{(u,v)}(k) := \max_{v+k\leq x\leq u} F_j(x)$, for $k = 1, \ldots, u-v$, for conciseness. It follows that, given a family of congestion games $\mathcal{G}_T^n$ corresponding to maximum number of users n, basis functions $b_1, \ldots, b_m$ and local taxation rule T, a lower bound on the Price of Stability can be computed as $\mathrm{PoS}(\mathcal{G}_T^n) \geq \max_j\{1/\gamma_j^{\mathrm{opt}}\}$, where γ_j^{opt}, $j = 1, \ldots, m$, is the optimal value of the following linear program:

$$\begin{aligned}&\underset{\gamma}{\text{maximize}} \quad \gamma \quad \text{subject to:}\\ &\gamma b_j(u)u \leq b_j(v)v + \sum_{k=1}^{u-v} F^{(u,v)}(k), \quad \forall (u,v) \in \{(u,v)\in\mathbb{N}^2 \text{ s.t. } 0\leq v<u\leq n\},\\ &F^{(u,v)}(k) = \max_{v+k\leq x\leq u} F_j(x), \quad \forall k\in\{1,\ldots,u-v\}, \forall (u,v) \in \{(u,v)\in\mathbb{N}^2 \text{ s.t. } 0\leq v<u\leq n\}.\end{aligned}$$

It is critical to note that we assumed $F_j(1) = 1$, for $j = 1, \ldots, m$ in the derivation of this program.

By including the functions F_j, $j = 1, \ldots, m$, as decision variables in the above linear program, we obtain a (not necessarily convex) program for minimizing the lower bound on $\mathrm{PoS}(\mathcal{G}_T^n)$.

We can then write the following m programs (one for each b_j) for computing the minimum lower bound on the Price of Stability achievable for a maximum allowable Price of Anarchy $\bar{\Pi}$ greater than or equal to the minimum achievable Price of Anarchy in $\mathcal{G}_T^n$, where we include the Price of Anarchy constraints from the linear program in (8.3):

$$\begin{aligned}
&\underset{F,\nu^{-1},\rho\nu^{-1},\gamma}{\text{maximize}} \quad \gamma \quad \text{subject to:}\\
&\rho\nu^{-1} \geq \bar{\Pi}^{-1}\nu^{-1}, \quad F(1) = 1,\\
&\nu^{-1}b_j(y)y - \rho\nu^{-1}b_j(x)x + (x-z)F(x) - (y-z)F(x+1) \geq 0, \quad \forall (x,y,z) \in \mathcal{I}(n),\\
&\gamma b_j(u)u \leq b_j(v)v + \sum_{k=1}^{u-v} F^{(u,v)}(k), \quad \forall (u,v) \in \{(u,v) \in \mathbb{N}^2 \text{ s.t. } 0 \leq v < u \leq n\},\\
&F^{(u,v)}(k) = \max_{v+k \leq x \leq u} F(x), \quad \forall k \in \{1,\dots,u-v\}, \forall (u,v) \in \{(u,v) \in \mathbb{N}^2 \text{ s.t. } 0 \leq v < u \leq n\}.
\end{aligned} \tag{9.25}$$

Let $F_j^{\text{opt}}, \nu_j^{\text{opt}}, \rho_j^{\text{opt}}, \gamma_j^{\text{opt}}$ be the optimal values that solve the above program. The Price of Anarchy achieved by the corresponding local taxation rule is $\text{PoA}(\mathcal{G}_{T^{\text{opt}}}^n) = \max_j\{1/\rho_j^{\text{opt}}\}$ and the resulting optimal lower bound on the Price of Stability is $\text{PoS}(\mathcal{G}_{T^{\text{opt}}}^n) \geq \max_j\{1/\gamma_j^{\text{opt}}\}$.

Note that the above program is not a convex program because of the equality constraints for the values $F^{(u,v)}(k)$. Next, we show that solving the above program is equivalent to the problem of maximizing $\sum_{x=1}^n F(x)$ subject to the Price of Anarchy constraints. First, for each basis function b_j, observe that maximizing the value of γ is equivalent to maximizing the sum over values $F^{(u,v)}(k)$, for $k = 1,\dots,u-v$, for all (u,v) such that $0 \leq v < u \leq n$. Second, observe that the upper bound on $F(x+1)$ imposed by $F(x)$ and ρ in the constraints corresponding to the Price of Anarchy is increasing in the value of $F(x)$, i.e.,

$$\begin{aligned}
(y-z)F(x+1) \leq (x-z)F(x) - \rho\nu^{-1}b_j(x)x + \nu^{-1}b_j(y)y,&\\
\forall (y,z) \in \{(y,z) \in \mathbb{N}^2 \text{ s.t. } (x,y,z) \in \mathcal{I}(n)\},&
\end{aligned}$$

since it always holds that $x - z \geq 0$ since $z \leq \min\{x,y\}$ in the definition of $\mathcal{I}(n)$. Thus, maximizing the value of $F(x+1)$ corresponds with maximizing the value of $F(x)$. Finally, maximizing the sum over values $F^{(u,v)}(k)$, for $k = 1, \ldots, u-v$, for all (u,v) such that $0 \leq v < u \leq n$, is equivalent to maximizing the sum over values $\max_{v+k \leq x \leq u} F(x)$, by definition. We showed above that maximizing the feasible value of any given decision variable $F(\hat{x})$, $1 \leq \hat{x} \leq n$, corresponds with maximizing the values of $F(x)$, for $x = 1, \ldots, \hat{x}-1$. Thus, maximizing $\max_{v+k \leq x \leq u} F(x)$, for $k = 1, \ldots, u-v$, for all (u,v) such that $0 \leq v < u \leq n$, is equivalent to maximizing $F(k)$, $k = 1, \ldots, n$, which is equivalent to maximizing $\sum_{k=1}^{n} F(k)$. It follows that the nonconvex program above yields the same solution as the linear program in (9.3), concluding the proof.

Arbitrary number of users. Here, we seek to investigate the system performance that can be obtained using a local taxation rule in settings with an arbitrary number of users. We first observe that a lower bound on the Pareto curve follows immediately from the linear program in Section 9.1.2, as the Price of Anarchy and Price of Stability metrics are both nondecreasing in the maximum number of users n. As the same argument does not apply to upper bounds on the two metrics, we develop a method for computing an upper bound on the Pareto curve between the Price of Anarchy and Price of Stability in polynomial congestion games with any number of users.

Before presenting our next result, we define several parameters that are necessary for extending the solution of the upcoming finite dimensional bilinear program to an arbitrary number of users. Given integer $d \in \mathbb{N}_{\geq 1}$, even integer $\bar{n} \in \mathbb{N}_{\geq 2}$, optimal values $\nu^{\text{opt}}, \rho^{\text{opt}}, \gamma^{\text{opt}}, \kappa^{\text{opt}}$ and function $F^{\text{opt}} : \{1, \ldots, n\} \to \mathbb{R}$, we define:

$$
\begin{aligned}
F^{\infty}(x) &:= \begin{cases} F^{\text{opt}}(x) & \text{if } 0 \leq x \leq \bar{n}/2, \\ \beta[x^{d+1} - (x-1)^{d+1}] & \text{if } x > \bar{n}/2 \end{cases} \\
\rho^{\infty} &:= \min\Big\{\rho^{\text{opt}}, \beta\nu^{\text{opt}}\frac{\bar{n}}{2}\Big[1 - \Big(1 - \frac{2}{\bar{n}}\Big)^{d+1}\Big] - d\Big[\frac{\beta\nu^{\text{opt}}}{d+1}\frac{\bar{n}}{2}\Big(\Big(\frac{2}{\bar{n}} + 1\Big)^{d+1} - 1\Big)\Big]^{1+1/d}\Big\} \\
\gamma^{\infty} &:= \min\{\gamma^{\text{opt}}, \gamma_1, \gamma_2 + \beta\kappa^{\text{opt}}, \gamma_3\}
\end{aligned}
\tag{9.26}
$$

where $\beta := F^{\mathrm{opt}}(\bar{n}/2)/[(\bar{n}/2)^{d+1} - (\bar{n}/2-1)^{d+1}]$,

$$\gamma_1 := \min_{x\in\{1,\dots,\bar{n}/2-1\}}$$
$$\frac{1}{x^{d+1}}\Big[\hat{y}_1^{d+1}(x) + F^{\mathrm{opt}}(x)x - F^{\mathrm{opt}}(x+1)\hat{y}_1(x) - \kappa^{\mathrm{opt}}\sum_{k=x+1}^{\bar{n}/2} F^{\mathrm{opt}}(k) - \beta\kappa^{\mathrm{opt}}\Big[\hat{y}^{d+1}(x) - \Big(\frac{\bar{n}}{2}\Big)^{d+1}\Big]\Big] \tag{9.27}$$

where $\hat{y}_1(x) = \max\{\bar{n}/2+1, [F^{\mathrm{opt}}(x+1)/(d+1)/(1-\beta\kappa^{\mathrm{opt}})]^{1/d}\}$,

$$\gamma_2 := \min_{y\in\{0,\dots,\bar{n}/2-1\}}\ \min_{r\geq\bar{n}/2}$$
$$\beta y + \frac{1}{r^{d+1}}\Big[y^{d+1} + \beta[r^{d+1} - (r-1)^{d+1}]r - \beta(r+1)^{d+1}y + \kappa^{\mathrm{opt}}\sum_{k=y+1}^{\bar{n}/2-1} F^{\mathrm{opt}}(k) - \beta\kappa^{\mathrm{opt}}\Big(\frac{\bar{n}}{2}-1\Big)^{d+1}\Big] \tag{9.28}$$

and,

$$\gamma_3 := \min_{x\geq\bar{n}/2} \beta\kappa^{\mathrm{opt}} + \frac{1}{x^{d+1}}\Big[(1-\beta\kappa^{\mathrm{opt}})\hat{y}_3^{d+1}(x) + \beta[x^{d+1}-(x-1)^{d+1}]x - \beta[(x+1)^{d+1}-x^{d+1}]\hat{y}_3(x)\Big] \tag{9.29}$$

where $\hat{y}_3(x) = \max\{\bar{n}/2, [\beta((x+1)^{d+1} - x^{d+1})/(d+1)/(1-\beta\kappa^{\mathrm{opt}})]^{1/d}\}$.

The following theorem presents our upper bound on the Pareto frontier between the Price of Anarchy and Price of Stability metrics in polynomial congestion games for any number of users:

Theorem 9.3.1. *Consider the family of resource cost functions $\mathcal{L} = span(b_1)$ corresponding to basis function $b_1(x) = x^d$, of order $d \in \mathbb{N}_{\geq 1}$. Further, consider a maximum allowable Price of Anarchy $\bar{\Pi}^{-1} \geq \sup_{n\geq 1}\mathrm{MinPoA}(n,\mathcal{L})$. Let $F^{\mathrm{opt}}, \nu^{\mathrm{opt}}, \rho^{\mathrm{opt}}, \gamma^{\mathrm{opt}}, \kappa^{\mathrm{opt}}$ be optimal values that* *solve the following bilinear program for some even integer* $\bar{n} \in \mathbb{N}_{\geq 2}$ *and* $\epsilon > 0$*:*

$$\underset{F\geq 0,\nu^{-1}\geq 0,\nu^{-1}\rho,\gamma,\kappa\geq 0}{\text{maximize}} \quad \gamma \quad \text{subject to:}$$

$$\nu^{-1}\rho \geq (\mathrm{PoA}^*)^{-1}\nu^{-1} \tag{9.30}$$

$$F(x+1) \geq F(x), \qquad \forall x \in \{1,\dots,\bar{n}-1\} \tag{9.31}$$

$$F\Big(\frac{\bar{n}}{2}\Big) \leq \nu^{-1}\Big[\Big(\frac{\bar{n}}{2}+1\Big)^{d+1} - \Big(\frac{\bar{n}}{2}\Big)^{d+1}\Big] \tag{9.32}$$

$$F(1) + \frac{(d+1)(\bar{n}/2+1)^d\kappa}{(\bar{n}/2)^{d+1}-(\bar{n}/2-1)^{d+1}}F\Big(\frac{\bar{n}}{2}\Big) \leq (d+1)\Big(\frac{\bar{n}}{2}+1\Big)^d \tag{9.33}$$

$$1 - \frac{\kappa}{(\bar{n}/2)^{d+1}-(\bar{n}/2-1)^{d+1}}F\Big(\frac{\bar{n}}{2}\Big) \geq \epsilon \tag{9.34}$$

$$\nu^{-1}y^{d+1} - \nu^{-1}\rho x^{d+1} + xF(x) - yF(x+1) \geq 0, \qquad \forall (x,y) \in \mathcal{I}_{\leq \bar{n}}(\bar{n}) \tag{9.35}$$

$$y^{d+1} - \gamma x^{d+1} + xF(x) - yF(x+1) + \kappa\Big[\sum_{k=1}^{x} F(k) - \sum_{k=1}^{y} F(k)\Big] \geq 0, \quad \forall (x,y) \in \mathcal{I}_{\leq \bar{n}}(\bar{n}) \tag{9.36}$$

where $\mathcal{I}_{\leq n}(n)$ *is the set of all pairs* $x, y \in \{0,\dots,n\}$ *such that* $1 \leq x+y \leq n$*. Then, for* $F^\infty : \mathbb{N} \to \mathbb{R}$*,* ρ^∞ *and* γ^∞ *defined as in* (9.26)*, the local taxation rule* T^∞ *with* $T^\infty(x^d)(x) = F^\infty(x) - x^d$ *satisfies* $\mathrm{PoA}(\mathcal{G}_{T^\infty}) \leq 1/\rho^\infty$ *and* $\mathrm{PoS}(\mathcal{G}_{T^\infty}) \leq 1/\gamma^\infty$*.*

In the study of polynomial congestion games, we are often interested in the setting where resource cost functions have the form $\ell(x) = \sum_{j=0}^d \alpha_j x^j$ for $\alpha_j \geq 0$ for given order $d \geq 1$. As it is stated, the result in Theorem 9.3.1 can only accommodate resource cost functions corresponding to a single monomial basis function. However, consider the scenario where we solve the bilinear program program in Theorem 9.3.1 for each monomial basis functions $b_1, \dots, b_m$, under the same values $\kappa = \kappa^*$, $\nu = \nu^*$ and $\bar{\Pi}^{-1}$ greater than or equal to $\sup_{n\geq 1} \mathrm{MinPoA}(n, \mathcal{L} = \mathrm{span}(b_1,\dots,b_m))$. Then, the values $\mathrm{PoA}(\mathcal{G}_T) \leq \max_j\{1/\rho_j^\infty\}$ and $\mathrm{PoS}(\mathcal{G}_T) \leq \max_j\{1/\gamma_j^\infty\}$ must be valid upper bounds on the Pareto frontier between the Price of Anarchy and Price of Stability in $\mathcal{G}$, where ρ_j^∞ and γ_j^∞ are derived as in (9.26) for all $j = 1,\dots,m$. We state this consequence of Theorem 9.3.1 in the following corollary:

Corollary 9.3.1. *Consider the family of resource cost functions $\mathcal{L} = span(b_1, \ldots, b_m)$ corresponding to basis function $b_j(x) = x^{d_j}$ of order $d_j \in \mathbb{N}_{\geq 1}$. Further, consider a maximum allowable Price of Anarchy $\bar{\Pi} \geq \sup_{n \geq 1} \mathrm{MinPoA}(n, \mathcal{L})$. Given even integer $\bar{n} \in \mathbb{N}_{\geq 2}$, $\hat{\kappa} \geq 0$, $\hat{\nu} \geq 0$ and $\epsilon > 0$, for each basis functions b_j, $j = 1, \ldots, m$, let $F_j^{\mathrm{opt}}, \rho_j^{\mathrm{opt}}, \gamma_j^{\mathrm{opt}}$ denote optimal values corresponding to a solution of the corresponding bilinear program in Theorem 9.3.1 under additional equality constraints $\kappa_j = \hat{\kappa}$ and $\nu_j = \hat{\nu}$. Then, for $F_j^\infty : \mathbb{N} \to \mathbb{R}$, ρ_j^∞ and γ_j^∞ defined as in (9.26), the local taxation rule T^∞ with $T^\infty(b_j)(x) = F_j^\infty(x) - x^{d_j}$ satisfies $\mathrm{PoA}(\mathcal{G}_{T^\infty}) \leq \max_j\{1/\rho_j^\infty\}$ and $\mathrm{PoS}(\mathcal{G}_{T^\infty}) \leq \max_j\{1/\gamma_j^\infty\}$.*

Proof of Theorem 9.3.1. We first provide an informal outline for the reader's convenience. A similar approach can be used to compute upper bounds for any class of congestion games. For congestion games with resource costs in $\mathcal{L} = \mathrm{span}(x^d)$, $d \in \mathbb{N}_{\geq 1}$, and any number of users n, the proof amounts to showing that the values $\gamma^\infty, \rho^\infty, \nu^{\mathrm{opt}}, \kappa^{\mathrm{opt}}$ from the bilinear program in Theorem 9.3.1 satisfy the constraints of two linear programs governing upper bounds on the Price of Anarchy and the Price of Stability, respectively. The constraints of these two linear programs are parameterized by each pair $x, y \in \{0, \ldots, n\}$. We divide the proof as follows:

– *Upper bound on the Price of Anarchy:* In the first part, we show that the values $\rho^\infty, \nu^{\mathrm{opt}}$ are feasible points of the linear program in (8.3) for the function F^∞ and n users. We first focus on the values x, y such that $1 \leq x + y \leq n$. We show that the constraints parameterized by $0 \leq x < \bar{n}/2$ and $y \geq 0$ are equivalent to constraints from the bilinear program in Theorem 9.3.1, leveraging the fact that $F^\infty(x) = F^{\mathrm{opt}}(x)$ and $F^\infty(x+1) = F^{\mathrm{opt}}(x+1)$. Then, for $x \geq \bar{n}/2$ and $y \geq 0$, we prove that all the constraints are satisfied if ρ^∞ is less than or equal to the second expression in the minimum that governs the definition of ρ^∞ in (9.26). Finally, we show that the constraints with $x + y > n$ are redundant, as they are less strict than those with $1 \leq x + y \leq n$.

– *Upper bound on the Price of Stability:* Next, we show that the values $\gamma^\infty, \kappa^{\mathrm{opt}}$ are feasible points of the linear program in Section 9.1.2 for $\nu = 1$, the function F^∞ and n users. We first focus on the values x, y such that $1 \leq x + y \leq n$. We show that the constraints with

$0 \leq x < \bar{n}/2$ and $0 \leq y \leq \bar{n}/2$, and the constraints with $x = 0$ and $y > \bar{n}/2$, are equivalent to constraints from the bilinear program in Theorem 9.3.1. Then, we prove that $\gamma^\infty, \kappa^{\text{opt}}$ satisfy the constraints with $1 \leq x < \bar{n}/2$ and $y > \bar{n}/2$ because $\gamma^\infty \leq \gamma_1$, the constraints with $x \geq \bar{n}/2$ and $0 \leq y < \bar{n}/2$ because $\gamma^\infty \leq \gamma_2$, and the constraints with $x, y \geq \bar{n}/2$ because $\gamma^\infty \leq \gamma_3$, for $\gamma_1, \gamma_2, \gamma_3$ defined as in (9.27)–(9.29). Finally, we show that the constraints with $x + y > n$ are less strict than those with $1 \leq x + y \leq n$.

It is important to note that F^∞ is a nondecreasing function by the constraints in (9.31) and by its definition in (9.26). Furthermore, there is always a feasible point in the bilinear program since $\bar{\Pi}$ is greater than or equal to the minimum achievable Price of Anarchy and because monomials of order $d \geq 1$ are all convex and nondecreasing. One feasible point corresponds with $\kappa = 0$, $\nu = 1$, and F and ρ solving the linear program in Theorem 8.1.1 for $\bar{n}$ users. Observe that all the constraints in the bilinear program are satisfied because the function F is unique and nondecreasing by Theorem 9.1.1, $\rho \geq \bar{\Pi}^{-1}$ since $1/\rho$ is the minimum achievable Price of Anarchy (which is strictly greater than 1 for polynomials [72]) and because the tax $T(x^d)$ must be lower (pointwise) than the marginal contribution. The last statement can be shown by virtue of our result on the best achievable lower bound on the Price of Stability in Theorem 9.1.3 which showed that the lower bound decreases for larger F. Since the Price of Stability of T according to the lower bound will be $1/\rho > 1$ for $\bar{n} \geq 2$ and the Price of Stability of marginal contribution is 1, our statement must hold.

The following inequalities are useful for the proof:

Observe that, for any $x \geq \bar{n}/2$, it holds that

$$(x+1)^d = \sum_{k=0}^{d} \binom{d}{k} x^{d-k} \leq x^d + x^{d-1} \sum_{k=1}^{d} \binom{d}{k} \left(\frac{\bar{n}}{2}\right)^{k-1} = x^d + x^{d-1} \frac{\bar{n}}{2} \Big[\Big(\frac{2}{\bar{n}} + 1\Big)^d - 1 \Big]. \quad (9.37)$$

We will also use the following two inequalities for sums of polynomials with $d \geq 1$ and

$x \geq y > 0$:

$$\sum_{k=y+1}^{x} k^d \geq \frac{1}{d+1}(x^{d+1} - y^{d+1}) + \frac{1}{2}(x^d - y^d) \tag{9.38}$$

$$\sum_{k=y+1}^{x} k^d \leq (x-y)x^d \leq x^{d+1} - y^{d+1}. \tag{9.39}$$

The remainder of the proof is divided into two parts as in the informal outline above:

– *Upper bound on the Price of Anarchy:* Consider the following linear program for computing the Price of Anarchy for any arbitrary number of users n given a nondecreasing function F:

$$\begin{aligned}
&\underset{\rho,\nu \geq 0}{\text{maximize}} \quad \rho \quad \text{subject to:}\\
&y^{d+1} - \rho x^{d+1} + \nu[xF(x) - yF(x+1)] \geq 0, \quad \forall x,y \in \{0,\dots,n\} \text{ s.t. } 1 \leq x+y \leq n,\\
&y^{d+1} - \rho x^{d+1} + \nu[(n-y)F(x) - (n-x)F(x+1)] \geq 0, \quad \forall x,y \in \{0,\dots,n\} \text{ s.t. } x+y > n.
\end{aligned} \tag{9.40}$$

First, we show that the above linear program is identical to the linear program in (8.3) for F nondecreasing. Observe that the constraints from the linear program in Section 9.1.2 are

$$\begin{aligned}
&y^{d+1} - \gamma x^{d+1} + (x-z)F(x) - (y-z)F(x+1)\\
&= y^{d+1} - \gamma x^{d+1} + xF(x) - yF(x+1) + z[F(x+1) - F(x)].
\end{aligned}$$

Since $F(x+1) - F(x) \geq 0$, the above expression is minimized for the smallest value of z. If follows that $z = 0$ when $x + y \leq n$ and $z = x + y - n$ when $x + y > n$, since the triplets $(x,y,z) \in \mathcal{I}(n)$ satisfy $1 \leq x+y-z \leq n$ and $z \leq \min\{x,y\}$. Thus, for $x,y \in \{0,\dots,n\}$, $x - z = x$ and $y - z = y$ when $x + y \leq n$, while $x - z = n - y$ and $y - z = n - x$ when $x + y > n$.

We show that the values $(\rho,\nu) = (\rho^\infty, \nu^{\text{opt}})$ are feasible in the above linear program for $F = F^\infty$ as defined in the claim and arbitrary n. We dispense with the case where $x = 0$ as, in this case, the strictest constraint on $F^\infty(1) = F^{\text{opt}}(1)$ is at $(x,y) = (0,1)$, which is already included by the constraints in (9.35). We first consider the constraints with $1 \leq x + y \leq n$.

– In the region where $1 \leq x < \bar{n}/2$ and $0 \leq y \leq \bar{n}/2$, observe that $x + y < \bar{n}$, $F^{\infty}(x) = F^{\mathrm{opt}}(x)$ and $F^{\infty}(x+1) = F^{\mathrm{opt}}(x+1)$. Then, the values $(\rho^{\infty}, \nu^{\mathrm{opt}})$ are feasible in the linear program in (9.40) for $F = F^{\infty}$ because $(F^{\mathrm{opt}}, \rho^{\mathrm{opt}}, \nu^{\mathrm{opt}})$ satisfy the constraints in (9.35), $F^{\infty}(x) = F^{\mathrm{opt}}(x)$ for $x \leq \bar{n}/2$ and $\rho^{\infty} \leq \rho^{\mathrm{opt}}$.

– Observe that, in the region where $1 \leq x < \bar{n}/2$ and $y > \bar{n}/2$, the constraint are less strict than when $y = \bar{n}/2$ if it holds that

$$\begin{aligned}
&\left(\frac{\bar{n}}{2}\right)^{d+1} - \frac{\bar{n}}{2}\nu^{\mathrm{opt}}F^{\mathrm{opt}}(x+1) \leq y^{d+1} - y\nu^{\mathrm{opt}}F^{\mathrm{opt}}(x+1) \\
\iff &\frac{y^{d+1} - (\bar{n}/2)^{d+1}}{y - \bar{n}/2} \geq \nu^{\mathrm{opt}}F^{\mathrm{opt}}(x+1).
\end{aligned}$$

The left-hand side of the last line is minimized for $y = \bar{n}/2 + 1$ by convexity and is most constraining for $x = \bar{n}/2 - 1$ since F^{opt} is nondecreasing. Observe that this condition on $F^{\mathrm{opt}}(\bar{n}/2)$ holds by the constraint in (9.32).

– Consider the region where $x \geq \bar{n}/2$ and $y \geq 0$. In this scenario, the constraints read as

$$y^{d+1} - \rho^{\infty}x^{d+1} + \beta\nu^{\mathrm{opt}}[x^{d+1} - (x-1)^{d+1}]x - \beta\nu^{\mathrm{opt}}[(x+1)^{d+1} - x^{d+1}]y \geq 0.$$

Observe that the left-hand side in the above is convex in y and that it is minimized over the nonnegative reals $y \geq 0$ at $\hat{y} = [\beta\nu^{\mathrm{opt}}[(x+1)^{d+1} - x^{d+1}]/(d+1)]^{1/d}$. Thus, it is sufficient to show that the following holds:

$$\begin{aligned}
&\rho^{\infty} \leq \beta\nu^{\mathrm{opt}}\Big[x - \frac{(x-1)^{d+1}}{x^d}\Big] - d\frac{1}{(d+1)^{1+\frac{1}{d}}}[\beta\nu^{\mathrm{opt}}[(x+1)^{d+1} - x^{d+1}]]^{1+\frac{1}{d}}\frac{1}{x^{d+1}} \\
\Longleftarrow\ &\rho^{\infty} \leq \beta\nu^{\mathrm{opt}}x\Big[1 - \Big(1 - \frac{1}{x}\Big)^{d+1}\Big] - d\Big[\frac{\beta\nu^{\mathrm{opt}}}{d+1}\frac{\bar{n}}{2}\Big(\Big(\frac{2}{\bar{n}} + 1\Big)^{d+1} - 1\Big)\Big]^{1+\frac{1}{d}},
\end{aligned}$$

where the implication holds by the identity in (9.37). The above inequality is strictest for $x = \bar{n}/2$ and is satisfied by the definition of ρ^{∞} in (9.26).

Note that, in the above, we have shown that for any $x, y \geq 0$, the linear program constraints corresponding with $1 \leq x + y \leq n$ are satisfied, without ever explicitly using the fact that

$1 \leq x+y \leq n$. Here, we show that $(\rho^\infty, \nu^{\text{opt}})$ is feasible for the constraints with $x+y > n$ by observing that these are less strict than the constraints we have already shown to be satisfied. Observe that this amounts to showing that

$$\begin{aligned} &\nu x F^\infty(x) - \nu y F^\infty(x+1) \leq \nu(n-y)F^\infty(x) - \nu(n-x)F^\infty(x+1) \\ \iff &\nu(x+y-n)[F(x) - F(x+1)] \leq 0. \end{aligned}$$

This must hold, since $\nu \geq 0$, $x+y-n > 0$ and F^∞ is nondecreasing.

– *Upper bound on the Price of Stability:* We continue by proving that $(\gamma, \nu, \kappa) = (\gamma^\infty, 1, \kappa^{\text{opt}})$ are feasible in the following linear program for $F = F^\infty$ as defined in the claim and arbitrary n:

$$\begin{aligned} &\underset{\gamma,\nu\geq 0,\kappa\geq 0}{\text{maximize}} \quad \gamma \quad \text{subject to:} \\ &y^{d+1} - \gamma x^{d+1} + \nu[xF(x) - yF(x+1)] + \kappa\left[\sum_{k=1}^{x} F(k) - \sum_{k=1}^{y} F(k)\right] \geq 0, \\ &\qquad\qquad \forall (x,y) \in \mathcal{I}_{\leq n}(n), \\ &y^{d+1} - \gamma x^{d+1} + \nu[(n-y)F(x) - (n-x)F(x+1)] + \kappa\left[\sum_{k=1}^{x} F(k) - \sum_{k=1}^{y} F(k)\right] \geq 0, \\ &\qquad\qquad \forall (x,y) \in \mathcal{I}_{>n}(n). \end{aligned} \tag{9.41}$$

where $\mathcal{I}_{\leq n}(n)$ and $\mathcal{I}_{>n}(n)$ are the sets of pairs $x, y \in \{0, \ldots, n\}$ such that $x+y \leq n$ and $x+y > n$, respectively. Following an identical set of arguments as the ones we used to show the equivalence of the linear programs in (8.3) and (9.40) for F nondecreasing, one can verify that the linear program in Section 9.1.2 is identical to the above linear program for F nondecreasing.

We first show that that $(\gamma, \nu, \kappa) = (\gamma^\infty, 1, \kappa^{\text{opt}})$ is feasible for constraints with $1 \leq x+y \leq n$.

– In the region where $0 \leq x < \bar{n}/2$ and $0 \leq y \leq \bar{n}/2$, observe that $x+y < \bar{n}$, $F^\infty(x) = F^{\text{opt}}(x)$ and $F^\infty(x+1) = F^{\text{opt}}(x+1)$. Then, the values $(\gamma^\infty, 1, \kappa^{\text{opt}})$ are feasible in the linear program in (9.41) because $(F^{\text{opt}}, \gamma^{\text{opt}}, 1, \kappa^{\text{opt}})$ satisfy the constraints in (9.36) and $\gamma^\infty \leq \gamma^{\text{opt}}$.

–In the setting where $x = 0$ and $y \geq \bar{n}/2$, the following must hold:

$$(1 - \beta\kappa^{\text{opt}})y^{d+1} - F^{\text{opt}}(1)y - \kappa^{\text{opt}} \sum_{k=1}^{\bar{n}/2} F^{\text{opt}}(k) + \beta\kappa^{\text{opt}} \Big(\frac{\bar{n}}{2}\Big)^{d+1} \geq 0.$$

Since $1-\beta\kappa^{\text{opt}} > 0$, it follows that the left-hand side is convex and is minimized over nonnegative real values $y \geq \bar{n}/2$ by

$$\hat{y} = \max\Big\{\frac{\bar{n}}{2}, \Big[\frac{F^{\text{opt}}(1)}{(d+1)(1-\beta\kappa^{\text{opt}})}\Big]^{1/d}\Big\}.$$

By Constraint 9.33, it holds that $\hat{y} = \bar{n}/2$, and the corresponding linear program condition is covered in Constraint 9.36.

–Consider the scenario where $1 \leq x \leq \bar{n}/2 - 1$ and $y \geq \bar{n}/2 + 1$, where we require that

$$(1 - \beta\kappa^{\text{opt}})y^{d+1} - \gamma^{\infty}x^{d+1} + F^{\text{opt}}(x)x - F^{\text{opt}}(x+1)y - \kappa^{\text{opt}} \sum_{k=x+1}^{\bar{n}/2} F^{\text{opt}}(k) + \beta\kappa^{\text{opt}} \Big(\frac{\bar{n}}{2}\Big)^{d+1} \geq 0.$$

Observe that the left-hand side is convex since $1 - \beta\kappa^{\text{opt}} > 0$ by the condition in (9.34) and is minimized for nonnegative real values $y \geq \bar{n}/2 + 1$ at

$$\hat{y} = \max\Big\{\frac{\bar{n}}{2} + 1, \Big[\frac{F^{\text{opt}}(x+1)}{(d+1)(1-\beta\kappa^{\text{opt}})}\Big]^{1/d}\Big\}.$$

Observe that the resulting linear program conditions are satisfied for $y = \hat{y}$ and for all $1 \leq x \leq \bar{n}/2$ since $\gamma^{\infty} \leq \gamma_1$, for γ_1 as defined in (9.27).

–We now consider the setting where $x \geq \bar{n}/2$ and $0 \leq y < \bar{n}/2$. Here we require:

$$\begin{aligned}\gamma^{\infty}x^{d+1} \leq y^{d+1} &+ \beta[x^{d+1} - (x-1)^{d+1}]x - \beta[(x+1)^{d+1} - x^{d+1}]y \\ &+ \kappa^{\text{opt}} \sum_{k=y+1}^{\bar{n}/2-1} F^{\text{opt}}(k) + \beta\kappa^{\text{opt}}\Big(x^{d+1} - \Big(\frac{\bar{n}}{2} - 1\Big)^{d+1}\Big).\end{aligned}$$

Observe that the resulting linear program constraints are satisfied for all $0 \leq y < \bar{n}/2$ since $\gamma^{\infty} \leq \gamma_2$ as defined in (9.28).

–For $x, y \geq \bar{n}/2$, we require

$$\gamma^{\infty} x^{d+1} \leq y^{d+1} + \beta[x^{d+1} - (x-1)^{d+1}]x - \beta[(x+1)^{d+1} - x^{d+1}]y + \beta\kappa^{\text{opt}}(x^{d+1} - y^{d+1}).$$

The left-hand side is convex in y as $1 - \beta\kappa^{\text{opt}} > 0$ and is minimized over the nonnegative reals $y \geq \bar{n}/2$ by

$$\hat{y} = \max\Big\{\frac{\bar{n}}{2}, \Big[\frac{\beta[(x+1)^{d+1} - x^{d+1}]}{(d+1)(1-\beta\kappa^{\text{opt}})}\Big]^{\frac{1}{d}}\Big\}.$$

The resulting linear program constraints are satisfied as $\gamma^{\infty} \leq \gamma_3$ as defined in (9.29).

Observe that in the above, we have not explicitly used the fact that $1 \leq x + y \leq n$. Thus, as we did for the upper bound on the Price of Anarchy, here we prove that the constraints with $x + y > n$ are less strict that the constraints with $1 \leq x + y \leq n$ for $(gamma^{\text{opt}}, 1, \kappa^{\text{opt}})$. In fact, this amounts to showing once more that

$$\begin{aligned} & \nu x F^{\infty}(x) - \nu y F^{\infty}(x+1) \leq \nu(n-y)F^{\infty}(x) - \nu(n-x)F^{\infty}(x+1) \\ \iff & \nu(x+y-n)[F(x) - F(x+1)] \leq 0. \end{aligned}$$

This must hold, since, $\nu = 1 > 0$, $x + y - n > 0$ and F^{∞} is nondecreasing.

Computing a lower bound on the attainable joint performance guarantees. Given the maximum number of users n, consider a game parameterization corresponding with a $(8^n) \times (3n)$ table $\mathcal{R}$ whose rows are all the unique permutations of $3n$-long binary vectors $\mathbf{r} \in \{0,1\}^{3n}$. Under this parameterization, each row $\mathbf{r} \in \mathcal{R}$ corresponds with a different resource, and, collectively, the rows encode the users' actions $a_i^{\text{w-ne}}, a_i^{\text{b-ne}}, a_i^{\text{opt}}$ as follows:

Consider the resource e corresponding with the row $\mathbf{r} \in \mathcal{R}$. For any $i \in \{1, \dots, n\}$, if $\mathbf{r}_i = 1$, then $e \in a_i^{\text{w-ne}}$, else $e \notin a_i^{\text{w-ne}}$; if $\mathbf{r}_{n+i} = 1$, then $e \in a_i^{\text{b-ne}}$, else $e \notin a_i^{\text{b-ne}}$; and, if $\mathbf{r}_{2n+i} = 1$, , then $e \in a_i^{\text{opt}}$, else $e \notin a_i^{\text{opt}}$. The coefficients in the basis representation of the 8^n resource cost functions will be the decision variables of our final linear program (i.e., there are $8^n \times m$ decision variables).

Recall that a lower bound on the Price of Stability of $\mathcal{G}_T^\tau$ can be computed as

$$\begin{aligned}
\underset{G\in\mathcal{G}_T^\tau}{\text{maximize}} \quad & \mathrm{SC}(a^{\text{b-ne}}) \quad \text{subject to:} \\
& \mathrm{SC}(a^{\text{opt}}) = 1, \quad \mathrm{SC}(a^{\text{w-ne}}) = \tau, \\
& C_i(a^{\text{w-ne}}) - C_i(a_i', a_{-i}^{\text{w-ne}}) \le 0, \quad \forall a_i' \in \{a_i^{\text{b-ne}}, a_i^{\text{opt}}\}, \quad \forall i \in \{1,\dots,n\}, \\
& C_i(a_{1:i}^{\text{b-ne}}, a_{i+1:n}') - C_i(a_{1:i-1}^{\text{b-ne}}, a_{i:n}') < 0, \quad \forall a' \in \mathcal{A}\setminus\{a^{\text{w-ne}}\}, \quad \forall i \in \{1,\dots,n\}.
\end{aligned} \tag{9.42}$$

Recast under the above parameterization, we observe that the problem of computing the lower bound in (9.42) can be formulated as the following linear program:

$$\begin{aligned}
\underset{\alpha_{e,j}\ge 0}{\text{maximize}} \quad & \sum_{e\in\mathcal{E}}\sum_{j=1}^{m} \alpha_{e,j}\cdot |a^{\text{b-ne}}|_e b_j(|a^{\text{b-ne}}|_e) \quad \text{subject to:} \\
& \sum_{e\in\mathcal{E}}\sum_{j=1}^{m} \alpha_{e,j}\cdot |a^{\text{opt}}|_e b_j(|a^{\text{opt}}|_e) = 1, \quad \sum_{e\in\mathcal{E}}\sum_{j=1}^{m} \alpha_{e,j}\cdot |a^{\text{w-ne}}|_e b_j(|a^{\text{w-ne}}|_e) = \tau, \\
& \sum_{e\in a_i^{\text{w-ne}}\setminus a_i'}\sum_{j=1}^{m}\Big[\alpha_{e,j}\cdot F_j(|a^{\text{w-ne}}|_e)\Big] - \sum_{e\in a_i'\setminus a_i^{\text{w-ne}}}\sum_{j=1}^{m}\Big[\alpha_{e,j}\cdot F_j(|a^{\text{w-ne}}|_e+1)\Big] \le 0, \\
& \qquad\qquad \forall a_i' \in \{a_i^{\text{b-ne}}, a_i^{\text{opt}}\}, \quad \forall i\in\{1,\dots,n\}, \\
& \sum_{e\in a_i^{\text{b-ne}}\setminus a_i'}\sum_{j=1}^{m}\Big[\alpha_{e,j}\cdot F_j(|a_{1:i}^{\text{b-ne}}, a_{i+1:n}'|_e)\Big] - \sum_{e\in a_i'\setminus a_i^{\text{b-ne}}}\sum_{j=1}^{m}\Big[\alpha_{e,j}\cdot F_j(|a_{1:i-1}^{\text{b-ne}}, a_{i:n}'|_e)\Big] < 0, \\
& \qquad\qquad \forall a' \in \mathcal{A}\setminus\{a^{\text{w-ne}}\}, \quad \forall i\in\{1,\dots,n\},
\end{aligned} \tag{9.43}$$

where we use $\mathcal{E}$ to denote the set of resources corresponding with the rows of $\mathcal{R}$, and $a_i'\setminus a_i = a_i'\setminus(a_i'\cap a_i)$ denotes the resources that user i selects in a_i' and not in a_i, e.g., if $e\in a_i^{\text{w-ne}}\setminus a_i^{\text{b-ne}}$, then the corresponding binary vector $\mathbf{r}$ must have $\mathbf{r}_i = 1$ and $\mathbf{r}_{n+i} = 0$. For given joint action a, observe that the quantity $|a|_e$ is simply the number of 1's in the appropriate columns of the binary vector $\mathbf{r}$ corresponding with e, i.e., $|a|_e = \sum_{i=1}^n \mathbf{r}_{nj_i+i}$ where $j_i = 0$ if $a_i = a_i^{\text{w-ne}}$, $j_i = 1$ if $a_i = a_i^{\text{b-ne}}$ and $j_i = 2$ if $a_i = a_i^{\text{opt}}$. Note that the objective and constraints in (9.42) and (9.43) coincide for games with a maximum number of users n.

As mentioned, the number of decision variables in (9.43) grows as $8^n \times m$ in the maximum number of users n and number of basis functions m, and, thus, the linear program is not tractable for large n. Nonetheless, lower bounds for small n can be computed within a reasonable amount of time, as provided in Figure 9.3.

9.3.2 Proofs from Section 9.2

Notation. Given a set $\mathcal{S}$, $|\mathcal{S}|$ represents its cardinality and $\mathbb{1}_{\mathcal{S}}$ describes the corresponding indicator function. ($\mathbb{1}_{\mathcal{S}}(e) = 1$ if $e \in \mathcal{S}$, 0 otherwise). We denote the index of the j'th component of a vector $\mathbf{v}$ with $\mathbf{v}_j$ or $\mathbf{v}(j)$ interchangeably. We use $\mathbf{1}$ to denote a vector of all ones and $\mathbf{0}$ to denote a vector of all zeros. We sometimes use the denotation $w(0) = u(0) = 0$ for any welfare or utility rule.

9.3.3 Linear Program Formulation of One Round Walk

We first give a linear program that computes the efficiency $\mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}}n; 1)$ that is based on a search for a worst case game construction $G \in \mathcal{G}_{\mathcal{W},\mathcal{U}}n$ that achieves the worst efficiency ratio for one-round. Here, $\mathcal{G}_{\mathcal{W},\mathcal{U}}n$ denotes the set of games with a fixed n number of agents, set of welfare rules $\mathcal{W}$ and utility design $\mathcal{U}$. A comparable primal-dual approach was also explored in [115] and [104] for different settings. We note that it is possible to extend the linear program for $\mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}}n, \kappa)$ for $\kappa > 1$ rounds, but the program becomes intractable in general.

First, we apply a key observation that for a game G, truncating the action set of each agent i to $\mathcal{A}_i = \{a_i^{\varnothing}, a_i^{\mathrm{br}}, a_i^{\mathrm{opt}}\}$ does not affect the efficiency metric $\mathrm{PoU}(G; 1)$. Here, $a_i^{\varnothing}$ is the null action that does not select any resources, a_i^{br} is the action that agent i plays after the one-round walk is completed with $a^{\mathrm{br}} \in \mathrm{sol}(1)$, and a_i^{opt} is the action that agent i plays in a joint action that optimizes the welfare $a^{\mathrm{opt}} = \arg\max_{a \in \mathcal{A}} W(a)$.[5] Therefore, we can restrict attention to the class of games $\mathcal{G}_{\mathcal{W},\mathcal{U}}^{n,3}$ where agents only have these three actions available without loss of generality. Furthermore, scaling W uniformly does not affect the ratio $\mathrm{PoU}(G; 1) = \frac{W(a^{\mathrm{br}})}{W(a^{\mathrm{opt}})}$, and we can assume that $W(a^{\mathrm{br}}) = 1$ without loss of generality. So we aim to find a game that

[5]Note that a_i^{br} and a_i^{opt} may be the same action, but using separate denotations does not affect the game structure. Additionally, if a^{br} is not unique, then the one that performs the worst with respect to W is selected.

maximizes the optimal welfare $W(a^{\text{opt}})$ to provide the lowest ratio. Consolidating the previous observations results in the following optimization problem

$$\text{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}}n;1)^{-1} = \tag{9.44}$$

$$\max_{G\in\mathcal{G}^{n,3}_{\mathcal{W},\mathcal{U}}} \quad W(a^{\text{opt}}) \quad \text{subject to:} \tag{9.45}$$

$$W(a^{\text{br}}) = 1, \tag{9.46}$$

$$U_i(a^{\text{br}}_{j\le i}, a^{\varnothing}_{j>i}) \ge U_i(a^{\text{br}}_{j<i}, a^{\text{opt}}_i, a^{\varnothing}_{j>i}) \quad \forall i \in \mathcal{I}, \tag{9.47}$$

The constraint inequality in (9.47) maintains that the joint action a^{br} is indeed the joint action that results after a one-round walk, where each agent i's best response is a^{br}_i (over a^{opt}_i) given that the previous $j \le i$ agents have also played a^{br}_j. To formulate the linear program from the optimization problem in (9.44), some necessary definitions are introduced. The possible resource allocations is enumerated by the following product set

$$\mathcal{P} = \prod_{i\in\mathcal{I}}\{\emptyset, \{a^{\text{br}}_i\}, \{a^{\text{opt}}_i\}, \{a^{\text{br}}_i, a^{\text{opt}}_i\}\},$$

where each resource is classified with the set of actions that select it by each agent. Then some corresponding vectors in $\{0,1\}^n$ can be defined.

$$b^p_i = \Big\{1 \text{ if } a^{\text{br}}_i \in p_i,\ 0 \text{ otherwise}\Big\},$$
$$o^p_i = \Big\{1 \text{ if } a^{\text{opt}}_i \in p_i,\ 0 \text{ otherwise}\Big\},$$

where $p \in \mathcal{P}$ describes a resource type. We define the norm of b^p to be $|b^p| = \sum_{i\in\mathcal{I}} b^p_i$ and denote the number of nonzero elements before index i as $|b^p|_{<i} = \sum_{1\le j<i} b^p_j$ (similarly for $|o^p| = \sum_{i\in\mathcal{I}} o^p_i$). With this, we describe the linear program in the following lemma.

Lemma 9.3.1. *Consider the welfare set $\mathcal{W} = \{w^1, \dots, w^m\}$ with $w^\ell(1) = 1$, and the corresponding utility design $\mathcal{U}(w^\ell) = u^\ell$ with $u^\ell(1) = 1$ for all ℓ. For n agents, the one-round walk* *efficiency is* $\mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}}n;1) = \min_{1\leq \ell \leq m} \frac{1}{\beta^\ell}$ [6], *where* $\beta^\ell \in \mathbb{R}$ *is the solution to*

$$\begin{aligned} \beta^\ell = \min_{\{\lambda_i\}_{i\in\mathcal{I}},\beta} \quad & \beta \quad \textit{subject to:} \\ & \beta w^\ell(|b^p|) \geq w^\ell(|o^p|) + \\ & \sum_{i\in\mathcal{I}} \lambda_i \Big[(b_i^p - o_i^p) u^\ell(|b^p|_{<i}+1) \Big] \quad \forall p \in \mathcal{P} \\ & \lambda_i \geq 0 \quad \forall i \in \mathcal{I}. \end{aligned} \tag{9.48}$$

Proof. First we show the equivalence of the optimization program proposed in (9.44) and the primal linear program described below. We later show that the dual of the program below is exactly the linear program described in the lemma. Here, we use $\beta = \mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}}n;1)^{-1}$ to denote the efficiency guarantee.

$$\beta = \max_{\{\eta_p^\ell\}_{\ell,p\in\mathcal{P}}} \sum_{\substack{1\leq \ell\leq m,\\ p\in\mathcal{P}}} w^\ell(|o^p|)\cdot \eta_p^\ell \quad \text{subject to:} \tag{9.49}$$

$$\sum_{\substack{1\leq \ell\leq m,\\ p\in\mathcal{P}}} w^\ell(|b^p|)\cdot \eta_p^\ell = 1 \tag{9.50}$$

$$\sum_{\substack{1\leq \ell\leq m,\\ p\in\mathcal{P}}} \Big[(b_i^p - o_i^p) u^\ell(|b^p|_{<i}+1) \Big]\cdot \eta_p^\ell \geq 0 \quad \forall i \in \mathcal{I} \tag{9.51}$$

$$\eta_p^\ell \geq 0 \quad \forall p \in \mathcal{P},\ 1\leq \ell \leq m. \tag{9.52}$$

Here, each decision variable $\eta_p^\ell \in \mathbb{R}_{\geq 0}$ is a real non-negative number. We define a vector label for each resource r as $\ell_r(i) = \{a_i \in \mathcal{A}_i :\ \text{if } r \in a_i\}$. This function describes in what actions is the resource selected by each agent i, with $\ell_r \in \mathcal{P}$. Furthermore, we denote the specific partition of the resource set with $\mathcal{R}sub = \{r \in \mathcal{R} : \ell_r = p, w_r = w^\ell\}$. Now we show that $W(a^{\mathrm{opt}})$ in

[6] Here, we assume that $\mathcal{W}$ is a finite set for ease of exposition, but it is straightforward to extend the efficiency result to an uncountable set.

(9.45) matches (9.49).

$$\begin{aligned}
W(a^{\text{opt}}) &= \sum_{r\in\mathcal{R}} w_r(|a^{\text{opt}}|_r) \\
&= \sum_{\substack{1\le\ell\le m,\\ p\in\mathcal{P}}} \sum_{r\in\mathcal{R}sub} w^\ell(|a^{\text{opt}}|_r) \\
&= \sum_{\substack{1\le\ell\le m,\\ p\in\mathcal{P}}} w^\ell(|o^p|)\cdot\eta_p^\ell,
\end{aligned}$$

where $\eta_p^\ell = |\mathcal{R}sub| \in \mathbb{N}$. The first equality is from the definition of the welfare function. The second equality results from partitioning the resource set. The third equality occurs by the fact that $|a^{\text{opt}}|_r = \sum_{j\in\mathcal{I}} \mathbb{1}_{a_j^{\text{opt}}}(r) = |o^p|$ if $r \in \mathcal{R}sub$; additionally, the value $w^\ell(|o^p|)$ is constant for any $r \in \mathcal{R}sub$. A similar argument can be made about the welfare of the best response action $W(a^{\text{br}})$, so (9.46) matches (9.50) as well.

Now we show the utility constraint in (9.47) matches the constraint in (9.51). For conciseness, let $a^1 = (a^{\text{br}}_{j<i}, a^{\text{br}}_i, a^{\varnothing}_{j>i})$ and $a^2 = (a^{\text{br}}_{j<i}, a^{\text{opt}}_i, a^{\varnothing}_{j>i})$. The utility difference can be written as

$$\begin{aligned}
U_i(a^1) - U_i(a^2) &= \sum_{r\in a_i^{\text{br}}} u_r(|a^1|_r) - \sum_{r\in a_i^{\text{opt}}} u_r(|a^2|_r) \\
&= \sum_{r\in\mathcal{R}} \Big(\mathbb{1}_{a_i^{\text{br}}}(r) u_r(|a^1|_r) - \mathbb{1}_{a_i^{\text{opt}}}(r) u_r(|a^2|_r)\Big) \\
&= \sum_{\substack{1\le\ell\le m,\\ p\in\mathcal{P}}} \sum_{r\in\mathcal{R}sub} \Big(\mathbb{1}_{a_i^{\text{br}}}(r) u_r(|a^1|_r) - \mathbb{1}_{a_i^{\text{opt}}}(r) u_r(|a^2|_r)\Big) \\
&= \sum_{\substack{1\le\ell\le m,\\ p\in\mathcal{P}}} \sum_{r\in\mathcal{R}sub} \Big[\big(b_i^p - o_i^p\big) u^\ell(|b^p|_{<i}+1)\Big] \\
&= \sum_{\substack{1\le\ell\le m,\\ p\in\mathcal{P}}} \Big[\big(b_i^p - o_i^p\big) u^\ell(|b^p|_{<i}+1)\Big]\eta_p^\ell.
\end{aligned}$$

The first equality is from the definitions of the utility functions. The second and third equalities comes from rewriting the sum using indicator functions and partitioning the resource set along $\mathcal{P}$. The fourth equality is a result of three facts: that $\mathbb{1}_{a_i^{\text{br}}}(r) = b_i^p$; that $\mathbb{1}_{a_i^{\text{opt}}}(r) = o_i^p$; that $|a^1|_r = \sum_{j<i} \mathbb{1}_{a_j^{\text{br}}}(r) + 1 = |b^p|_{<i} + 1$ if $r \in a_i^{\text{br}}$ (similarly for $|a^2|_r$). The fifth equality comes from sliding out the relevant terms of the first sum.

The constraint in (9.52) ensures a well-defined non-degenerate game parametrization. Observe that in the primal program in (9.49), we have relaxed $\eta_p^\ell \in \mathbb{N}$ to $\eta_p^\ell \in \mathbb{R}$, where we have normalized the number of resources in each partition so that $W(a^{\text{br}}) = 1$. This is done without loss of generality, since we can scale up the optimal arguments $\{\eta_p^\ell\}_{\ell,p\in\mathcal{P}}$ uniformly and round to derive a corresponding valid game construction that achieves an efficiency ratio $\text{PoU}(G;1)$ that is arbitrarily close to the solution of the primal program.

We now verify that the dual of the program in (9.49) is the one in (9.48). Note that primal program in (9.49) can be concisely written as

$$\begin{aligned}
&\max_{\eta} \quad c^T\eta \quad \text{subject to:}\\
&A\eta = 1\\
&\begin{bmatrix} H \\ I_{m\cdot 4^n} \end{bmatrix} \eta \succeq 0,
\end{aligned}$$

where η is the vector of $\{\eta_p^\ell\}_{\ell,p\in\mathcal{P}}$, $I_{m\cdot 4^n}$ corresponds to the identity matrix of dimension $m \cdot 4^n \times m \cdot 4^n$, and c, A, H are the compactly written vectors in equations (9.49), (9.50), and (9.51) respectively. Writing the dual linear program gives

$$\begin{aligned}
&\max_{\lambda\succeq 0,\ \xi\succeq 0,\ \beta} \quad -\beta \quad \text{subject to:}\\
&A_\ell^T\beta - \begin{bmatrix} H_\ell^T, I_{4^n} \end{bmatrix} \begin{bmatrix} \lambda \\ \xi \end{bmatrix} - c_\ell = \mathbf{0} \quad \forall 1 \le \ell \le m,
\end{aligned}$$

where $\mathbf{0}$ is a vector of zeros, and $c = (c_1^T, \ldots, c_m^T)^T$ associated with each $1 \le \ell \le m$ (likewise for A and H). Observing that $A_\ell^T\beta - \begin{bmatrix} H_\ell^T, I_{4^n} \end{bmatrix} \begin{bmatrix} \lambda \\ \xi \end{bmatrix} - c_\ell = \mathbf{0}$ is equivalently written as

$A_\ell^T\beta - H_\ell^T\lambda - c_\ell = \xi$ and as $A_\ell^T\beta + H_\ell^T\lambda + c_\ell \succeq 0$ and substituting back c_ℓ, A_ℓ, H_ℓ gives the result. □

While we have an exact characterization of the one-round walk efficiency, we cannot use this program directly to derive efficiency bounds tractably. However, by reasoning about the tight constraints in dual program, we can arrive at a more tractable program when the number of agents is not fixed.

Lemma 9.3.2. *Consider the welfare set $\mathcal{W} = \{w^1, \ldots, w^m\}$ with $w^\ell(1) = 1$, and the corresponding utility design $\mathcal{U}(w^\ell) = u^\ell$ with $u^\ell(1) = 1$ for all ℓ. The one round walk efficiency is* $\mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}}; 1) = \min_{1\leq\ell\leq m}\frac{1}{\beta^\ell}$*, where $\beta^\ell \in \mathbb{R}\cup\{\infty\}$ is the solution to*

$$\begin{aligned}
&\beta^\ell = \min \quad \beta \quad \textit{subject to:} \\
&\beta w^\ell(y) \geq \mathrm{H}^\ell\left(\sum_{\mathrm{i}=1}^{\mathrm{y}} \mathrm{u}^\ell(\mathrm{i}) - \mathrm{z}\min_{1\leq \mathrm{i}\leq \mathrm{y}+1}\mathrm{u}^\ell(\mathrm{i})\right) + \mathrm{w}^\ell(\mathrm{z}) \\
&\textit{for all } z, y \in \mathbb{N} \textit{ s.t. } z \geq 0 \textit{ and } y \geq 1,
\end{aligned} \tag{9.53}$$

and $\mathrm{H}^\ell = \sup_{\mathrm{i}} \mathrm{w}^\ell(\mathrm{i})/\mathrm{i}$.

Proof. The dual program in (9.48) provides a solution for β^ℓ for fixed n agents. We first show the solution is upper bounded by $\beta^\ell \leq \tilde{\beta}^\ell$ for any n, where $\tilde{\beta}^\ell$ is the solution to the program in (9.53).

Let n be the number of agents. For a given $p \in \mathcal{P}$, we denote $y_p = |b^p|$ and $z_p = |o^p|$ for ease of notation. Additionally, to convey which indices the resource type p are non-zero in and in what order, we define vectors B^p for a^{br} and O^p for a^{opt}. Formally, $\mathrm{B}^p : \{1, \ldots, y_p\} \to \{1, \ldots, n\}$ and $\mathrm{O}^p : \{1, \ldots, z_p\} \to \{1, \ldots, n\}$ with

$$\mathrm{B}^p(j) = i \text{ if } b_i^p = 1 \text{ and } |b^p|_{\leq i} = j,$$

$$\mathrm{O}^p(j) = i \text{ if } o_i^p = 1 \text{ and } |o^p|_{\leq i} = j.$$

Considering the dual program in (9.48), we add the constraint $\lambda_i = \mathrm{H}^\ell = \max_{1\leq \mathrm{j}\leq \mathrm{n}} \mathrm{w}^\ell(\mathrm{j})/\mathrm{j}$ explicitly. Since we shrink the feasible region, the optimal solution to (9.48) potentially increases. We verify that the resulting feasible region is nonempty. Consider the constraints according to p such that $b^p = \mathbf{0}$. The corresponding dual constraint takes the form

$$0 \geq w^\ell(z_p) - \sum_{j=1}^{z_p} \lambda_{\mathrm{O}^p(j)} u^\ell(1).$$

Simplifying the expression gives $\sum_{j=1}^{z_p} \lambda_{\mathrm{O}^p(j)} \geq w^\ell(z_p)$, which is always satisfied if $\lambda_i = \mathrm{H}^\ell$ for all i. If the constraints according p are such that $b^p \neq \mathbf{0}$, then $\beta w^\ell(y)$ is present and strictly positive in the inequality (9.48) and β can be taken as high as needed to satisfy the constraint. Therefore the feasible region is nonempty.

For any $p \in \mathcal{P}$ such that $b^p \neq \mathbf{0}$, we can simplify the dual constraint in (9.48) to

$$\beta w^\ell(y_p) \geq w^\ell(z_p) + \sum_{i=1}^{y_p} \mathrm{H}^\ell \mathrm{u}^\ell(\mathrm{i}) - \sum_{\mathrm{i}\in\mathcal{I}} \mathrm{H}^\ell \mathrm{o}_\mathrm{i}^\mathrm{p} \mathrm{u}^\ell(|\mathrm{b}^\mathrm{p}|_{<\mathrm{i}} + 1).$$

Furthermore, for any $p \in \mathcal{P}$, we observe that $\sum_{i\in\mathcal{I}} o_i^p u^\ell(|b^p|_{<i} + 1) \geq z_p \min_{1\leq i\leq y_p+1} u^\ell(i)$. Thus, for any $p \in \mathcal{P}$, we can replace the corresponding dual constraint with a more binding constraint

$$\beta w^\ell(y) \geq w^\ell(z) + \sum_{i=1}^{y} \mathrm{H}^\ell \mathrm{u}^\ell(\mathrm{i}) - \sum_{\mathrm{i}\in\mathcal{I}} \mathrm{H}^\ell \mathrm{z} \min_{1\leq \mathrm{i}\leq \mathrm{y}+1} \mathrm{u}^\ell(\mathrm{i}),$$

for some $0 \leq z \equiv z_p \leq n$ and $1 \leq y \equiv y_p \leq n$. Therefore, replacing the dual constraints gives an upper bound for $\beta^\ell \leq \tilde{\beta}^\ell$. Limiting the number of agents $n \to \infty$ results in the program in (9.53).

Now we show that the solution is lower bounded by $\beta^\ell \geq \tilde{\beta}^\ell$, where $\tilde{\beta}^\ell$ is the solution to the program in (9.53). We show that when we remove dual constraints, we arrive at the program in (9.53). Since the feasible region expands, the optimal solution potentially decreases. Let the set of agents be $\mathcal{I} = \mathbb{N}$ and $j^p = \arg\min_{1\leq j\leq y_p+1} u^\ell(j)$. We remove all the dual constraints barring the constraints that correspond to $p \in \mathcal{P}$ with either (**a**) $y_p = 0$ and $z_p = z^* = \arg\max w^\ell(j)/j$ or (**b**) $y_p > 0$ and $\mathrm{B}^p(j^p - 1) < \mathrm{O}^p(1)$ and $\mathrm{O}^p(z_p) < \mathrm{B}^p(j^p)$. The first property refers to all resource types where a^{br} is never selected but a^{opt} is by z^* agents. The second property refers to all resource types where the indices of the agents selecting a^{opt} are between the agents with index $\mathrm{B}^p(j^p-1)$ and $\mathrm{B}^p(j^p)$.

Assume property (a). Then the corresponding dual constraint in (9.48) can be written as

$$0 \geq w^\ell(z^*) - \sum_{j=1}^{z^*} \lambda_{\mathrm{O}^p(j)} u^\ell(1),$$

for any resource type $p \in \mathcal{P}$ that satisfies property (a). Therefore, for any $j \in \mathbb{N}$, except for at most z^*-1 values, observe that $\lambda_j \geq \mathrm{H}^\ell$ must hold.

Now assume property (b). With respect to a resource type $p \in \mathcal{P}$ that satisfies property (b), we observe that $u^\ell(|b^p|_{<i}+1) = u^\ell(j^p)$ for any agent with index $i = \mathrm{O}^p(j)$ for some j. Therefore, under the two previous observations, we can rewrite the relaxed dual program as

$$\begin{aligned}
&\min_{\lambda \succeq \mathbf{0}} \quad \beta \quad \text{subject to:} \\
&\beta w^\ell(y_p) \geq \sum_{j=1}^{y_p} \lambda_{\mathrm{B}^p(j)} u^\ell(j) - \sum_{j=1}^{z_p} \lambda_{\mathrm{O}^p(j)} u^\ell(j^p) + w^\ell(z_p) \\
&\text{for all } p \in \mathcal{P}', \\
&\lambda_i \geq \mathrm{H}^\ell \ \text{ for all } i \in \mathbb{N} \text{ but at most } z^*-1 \text{ values,}
\end{aligned} \tag{9.54}$$

where $\mathcal{P}' = \{p \in \mathcal{P} : p \text{ satisfies property (b)}\}$. Assuming that the optimal dual variable is $\lambda_i = \mathrm{H}^\ell$ for all $i \in \mathbb{N}$, observe that we recover the proposed program given in (9.53). To show this claim, we confirm that the binding constraint for β in (9.54) is larger when considering a different sequence of lambdas $\lambda \neq \mathrm{H}^\ell \mathbf{1}$. In other words for a given $y \geq 1$ and $z \geq 0$, we show that for the resulting dual variables,

$$\begin{aligned}
\beta_\lambda &= \max_{p \in \mathcal{P}'} \Big\{ \frac{1}{w^\ell(y_p)} \big(\sum_{j=1}^{y_p} \lambda_{\mathrm{B}^p(j)} u^\ell(j) - \sum_{j=1}^{z_p} \lambda_{\mathrm{O}^p(j)} u^\ell(j^p) \big) \Big\} \\
&\geq \frac{\mathrm{H}^\ell}{w^\ell(y)} \left(\sum_{j=1}^{y} u^\ell(j) - \sum_{j=1}^{z} u^\ell(j^p) \right) = \beta_{y,z}
\end{aligned} \tag{9.55}$$

For any $\lambda \neq \mathrm{H}^\ell \mathbf{1}$, consider two cases where either λ is a divergent sequence, or it is bounded above. In the first case, since λ must satisfy $\lambda_j \geq 0$ for all $j \in \mathbb{N}$, the limit $\lim_{j\to\infty} \lambda_j = \infty$. If $u^\ell(j) = 0$ for all j, note that $\beta_{y,z} = 0$ for any $y \geq 1$ and $z \geq 0$. Since β_λ must also be greater than 0, the inequality in (9.55) holds in this case. If $u^\ell(J) > 0$ for some $J \in \mathbb{N}$, consider a constraint with p such that $y_p > J$ and $z_p = 0$. For any $M > 0$, we can choose B^p, such that $\lambda_{\mathrm{B}^p(j)} > M$ for all $1 \leq j \leq y_p$. Thus $\beta_\lambda \geq \frac{1}{w(y_p)} \sum_{j=1}^{y_p} Mu(j)$. Since M is arbitrary, $\beta_\lambda = \infty \geq \beta_{y,z}$ for any $y \geq 1$ and $z \geq 0$ as well.

In the second case, since λ is also bounded below by H^ℓ, for all but a finite set of values, there exists a convergent sub-sequence λ_{ss} that converges to a value $V \geq \mathrm{H}^\ell$ by the Bolzano-Weierstrauss theorem. Let $M_u^y = \max_{1\leq j\leq y+1} u^\ell(i)$, $x = \max(y, z)$, and $\varepsilon > 0$. Since λ_{ss} converges, there exists a $J \in \mathbb{N}$ such that for any $j \geq J$, $|\lambda_{ss}(j) - V| \leq \frac{\varepsilon}{2M_u^y x}$.

For a given y and z, consider any constraint with $p \in \mathcal{P}'$ such that $y_p = y$ and $z_p = z$. Additionally, B^p and O^p can be chosen to ensure that $|\lambda_{\mathrm{B}^p(j)} - V| \leq \frac{\varepsilon}{2M_u^y x}$ and$|\lambda_{\mathrm{O}^p(j)} - V| \leq \frac{\varepsilon}{2M_u^y x}$ for all j. Therefore

$$
\begin{aligned}
\beta_\lambda &\geq \frac{1}{w^\ell(y_p)} \big(\sum_{j=1}^{y_p} \lambda_{\mathrm{B}^p(j)} u^\ell(i) - \sum_{j=1}^{z_p} \lambda_{\mathrm{O}^p(j)} u^\ell(j^p) \big) \\
&\geq \frac{V}{w^\ell(y)} \big(\sum_{j=1}^{y} u^\ell(i) - \sum_{j=1}^{z} u^\ell(j^p) \big) - \frac{\varepsilon}{2} - \frac{\varepsilon}{2} \\
&\geq \beta_{y,z} - \varepsilon.
\end{aligned}
$$

Since ε is arbitrary, we have that $\beta_\lambda \geq \beta_{y,z}$ for any y and z and we show the claim. Therefore the proposed program is an upper bound and we have shown the equality $\beta^\ell = \tilde{\beta}^\ell$. □

9.3.4 Proof of Theorem 9.2.1

Given a set of welfare rules and utility rules, Lemma 9.3.2 provides an exact characterization of the one-round walk efficiency through a linear program. We modify the linear program in (9.53) to compute the utility rules that optimize the one-round walk efficiency. If a given welfare rule w is submodular, note that $\sup_i w(i)/i = 1$ and so $\mathrm{H}^\ell = 1$. Furthermore, if the utility rule u is assumed to be non-increasing, then $\min_{1 \leq i \leq y+1} u(i) = u(y+1)$. Additionally, $w(1) - 1 \cdot u(y+1) \geq w(0) - 0 \cdot u(y+1) = 0$ for any $y \geq 1$, so $z = 0$ is a nonbinding constraint. We lastly note that the values $\{u(i)\}_{i \in \mathcal{I}}$ can be established as decision variables for the program in (9.53) to produce the linear program in (9.15), rewritten below.

$$
\begin{aligned}
&(\beta^*, u^*) \in \arg \min_{\beta, \{u(i)\}_{i \in \mathcal{I}}} \quad \beta \quad \text{subject to:} \\
&\beta w(y) \geq \sum_{i=1}^{y} u(i) - z u(y+1) + w(z) \quad \forall y, z \geq 1 \\
&u(1) = 1,
\end{aligned} \tag{9.56}
$$

where β^* is a tight characterization of the efficiency guarantee only if the resulting optimal utility rule u^* is non-increasing and a lower bound if not. We now verify that the optimal utility rule u^* is indeed non-increasing. First, rearranging the terms in the constraint in (9.56) gives that for any $y \geq 1$,

$$
u^*(y+1) \geq \sup_{z \geq 1} \Big(\frac{1}{z} \big(\sum_{i=1}^{y} u^*(i) + w(z) - \beta^* w(y) \big) \Big). \tag{9.57}
$$

We verify $u^*(y+1)$ is well-defined. Note that since u^* is optimal, the efficiency bound $\beta^* < \infty$ is nontrivial (as u_{mc} guarantees an efficiency guarantee greater than 1/2). Then, by recursion and the fact that $\frac{w(z)}{z} \leq 1$ for all z, there exists a solution for $u^*(y+1)$ such that (9.57) holds with equality and the resulting value is finite for all $y \geq 1$. Additionally $u^*(y)$ must be non-negative for all $y \geq 1$, since limiting $z \to \infty$ in (9.57) gives that $u(y+1) \geq 0$.

Now we show that the solution u^* is non-increasing. Suppose for contradiction that for some $y \geq 1$, that $u^*(y) < u^*(y+1)$. Let $z_{y+1} \in \arg\max_{z \geq 1} w(z) - z u(y+1)$ be the number that achieves the maximum.

We verify that z_{y+1} is well-defined. Suppose for contradiction that $w(z) - z u^*(y+1)$ is always increasing in z, so z_{y+1} is not well defined. Since $\beta^* < \infty$, the limit $\lim_{z \to \infty} w(z) - z u^*(y+1)$ must converge and therefore $u^*(y+1)$ must be equal to $Q = \lim_{z \to \infty} \Delta w(z)$, where we denote

$\Delta w(z) = w(z) - w(z-1)$ for conciseness. From the original contradiction assumption then $u^*(y) < u^*(y+1) = Q$. Then taking the constraint in (9.56), with $y-1$ and $z \to \infty$ gives $\beta w(y-1) \geq \lim_{z\to\infty} w(z) - zu^*(y) \geq \infty$, which is a contradiction.

Now, substituting z_{y+1} into (9.57) for y and $y+1$ produces the following expressions

$$u^*(y+1) = \frac{1}{z_{y+1}}\big(\sum_{i=1}^{y} u^*(i) + w(z_{y+1}) - \beta^* w(y)\big)$$
$$u^*(y) \geq \frac{1}{z_{y+1}}\big(\sum_{i=1}^{y-1} u^*(i) + w(z_{y+1}) - \beta^* w(y-1)\big).$$

Inputting these expressions into the assumption $u^*(y) < u^*(y+1)$ reduces to the inequality $u(y) > \beta^*\Delta w(y)$. Similarly, for some $j \geq 1$, substituting z_{y+j} into (9.56) for $y+j$ and $y+j+1$ gives

$$u^*(y+j+1) \geq \frac{1}{z_{y+j}}\big(\sum_{i=1}^{y+j} u^*(i) + w(z_{y+j}) - \beta^* w(y+j)\big)$$
$$u^*(y+j) = \frac{1}{z_{y+j}}\big(\sum_{i=1}^{y+j-1} u^*(i) + w(z_{y+j}) - \beta^* w(y+j-1)\big).$$

Thus by substituting the second expression into first, the following inequality holds

$$u^*(y+j+1) \geq u^*(y+j) + \frac{u^*(y+j) - \beta^*\Delta w(y+j)}{z_{y+j}}. \tag{9.58}$$

We show, by induction, that the following expression holds for any $j \geq 1$,

$$\frac{u^*(y+j) - \beta^*\Delta w(y+j)}{z_{y+j}} \geq \frac{u^*(y+1) - \beta^*\Delta w(y+1)}{z_{y+1}} > 0. \tag{9.59}$$

The base case holds for $j = 1$, since

$$u^*(y+1) - \beta^*\Delta w(y+1) > u^*(y) - \beta^*\Delta w(y) > 0.$$

This comes from the assumption that $u^*(y+1) > u^*(y)$, $\Delta w(y+1) \leq \Delta w(y)$ by submodularity of w, and that $u^*(y) - \beta^*\Delta w(y) > 0$ from the previous argument. For the inductive case for $J \geq 2$, assume that the inequality holds for all $j < J$. Then, by applying the induction assumption to (9.58) and subsequently to the definition of z_{y+J}, we have that

$$u^*(y+J) > u^*(y+J-1) > \cdots > u^*(y+1)$$
$$z_{y+J} \leq z_{y+J-1} \leq \cdots \leq z_{y+1}.$$

Therefore the statement in (9.59) holds due to the aforementioned inequalities and the fact that $\Delta w(y+J) \leq \Delta w(y+1)$ due to submodularity of w. Therefore (9.59) holds and we have that $u^*(y+j+1) \geq u^*(y+j) + D$, where $D = \frac{u^*(y+1)-\beta^*\Delta w(y+1)}{z_{y+1}} > 0$. Following this, $u^*(y+j) \geq u^*(y+1) + D(j-1)$.

Now consider the constraint in (9.56) where $y \to \infty$ and $z = 0$. Since $w(y) \leq y$,

$$\beta^* \geq \lim_{y\to\infty} \frac{1}{y}\sum_{i=1}^{y} u^*(i) \geq \infty, \tag{9.60}$$

where the last inequality results from the fact that $u^*(y) \sim y$ is of linear order by the previous argument. Since β^* must be finite, contradiction ensues and the solution u^* must be non-increasing and the efficiency guarantees are tight for each linear program.

Thus, so far, we have shown the statement in (9.14) with regards to the welfare set $\mathcal{W} = \{w^1, \ldots, w^m\}$. We lastly show that the results extend linearly to a span of welfare rules as claimed in (9.16). Note that for the welfare set ww_{span} spanned from $\{w^1, \ldots, w^m\}$, the resulting optimal guarantees $\text{PoU}^*(\mathcal{W}; 1) \geq \text{PoU}^*(ww_{\text{span}}; 1)$, since ww_{span} is a larger set of welfare rules. We show that the utility design as in (9.16) achieves $\text{PoU}^*(ww_{\text{span}}; 1) = \text{PoU}^*(\mathcal{W}; 1)$ and therefore is optimal with respect to ww_{span}. Consider $w = \sum_{\ell=1}^{m} \alpha^\ell w^\ell$ for any non-negative $\{\alpha^\ell\}_{1\leq\ell\leq m}$. Let the corresponding utility design be $\mathcal{U}_{\text{lin}}(w) = \sum_{\ell=1}^{m} \alpha^\ell u^\ell$, where u^ℓ is the corresponding solution to (9.15) for w^ℓ and $\beta^* = \min_{1\leq\ell\leq m} \frac{1}{\beta^\ell} = \text{PoU}^*(\mathcal{W}; 1)$. From the characterization program in (9.53) with respect to ww_{span} and $\mathcal{U}_{\text{lin}}$, the dual constraint for any

$y \geq 1$, $z \geq 0$, $1 \leq \ell \leq m$, and $\{\alpha^\ell\}_{1 \leq \ell \leq m}$ can be rewritten as

$$\sum_{\ell=1}^{m} \alpha^\ell \cdot \left[\beta^* w^\ell(y) - \sum_{j=1}^{y} u^\ell(j) + z u^\ell(y+1) - w^\ell(z)\right] \geq 0.$$

This constraint will always be satisfied for any non-negative $\{\alpha^\ell\}_{1 \leq \ell \leq m}$, as the inner terms is non-negative by definition of β^* and u^ℓ. Therefore, we have that under the linear utility design $\mathcal{U}_{\text{lin}}$, we have that $\text{PoU}(\mathcal{G}_{ww_{\text{span}}, \mathcal{U}_{\text{lin}}}) \geq \text{PoU}^*(\mathcal{W}; 1)$ and is optimal.

9.3.5 Proof of Theorem 9.2.2

Given a curvature $\mathcal{C}$, let $\mathcal{W}$ be the set of welfare rules that have curvature of at most $\mathcal{C}$. From Lemma 8.3.4, we know there exists a basis set of welfare rules, such that for any $w \in \mathcal{W}$, we can come up with a decomposition $w = \sum_{b \in \mathbb{N}} \alpha^b w^b$, with $\alpha^b = (2w(b) - w(b-1) - w(b+1))/\mathcal{C}$ and

$$w^b(j) = \begin{cases} j, & \text{if } 0 \leq j \leq b \\ b + (1 - \mathcal{C}) \cdot (j - b) & \text{if } j > b. \end{cases} \tag{9.61}$$

We refer to these welfare rules as *b-covering* welfare rules. We note that for any $b \in \mathbb{N}$, the welfare rule w^b has a curvature of $\mathcal{C}$. For each welfare rule w^b, we claim that the corresponding optimal utility rule from running the program in (9.15) is

$$u^b(j) = \begin{cases} (1 - \beta^b)(\frac{b+1}{b})^{j-1} + \beta^b & \text{if } j \leq b+1 \\ (1 - \mathcal{C})\beta^b & \text{if } j \geq b+1, \end{cases} \tag{9.62}$$

where and $\beta^b = \frac{(\frac{b+1}{b})^b}{(\frac{b+1}{b})^b - \mathcal{C}}$ is the resulting optimal efficiency. Taking the minimum across b, we have that $\min_{b \in \mathbb{N}} \frac{1}{\beta^b} = 1 - \mathcal{C}/2$ for $b = 1$. Therefore, using Theorem 9.2.1, the optimal efficiency guarantee is $\text{PoU}^*(\mathcal{W}; 1) = 1 - \mathcal{C}/2$.

Now we verify that u^b and β^b are indeed the optimal solutions. We first remove all constraints in (9.15) apart from the ones that satisfy $z = b$ for any $y \geq 1$. This results in a lower bound for β^b that we claim later to be tight.

Rearranging the terms in the constraint in (9.56) gives that for any $y \geq 1$, the optimal solution satisfies

$$u^*(y+1) = \sup_{z\geq 1}\Big(\frac{1}{z}\big(\sum_{i=1}^{y} u^*(i) + w(z) - \beta^* w(y)\big)\Big). \tag{9.63}$$

Substituting in for w and the binding constraint $z = b$, the recursive equation for u^b is then

$$\begin{aligned} u^b(1) &= 1 \\ u^b(j+1) &= \frac{1}{b}\sum_{i=1}^{j} u^b(i) + 1 - \frac{1}{b}\beta^* w^b(j), \end{aligned}$$

for some optimal $\beta^* \geq 1$. To solve for the closed form expression for u^b, a corresponding linear, time-invariant, discrete time system is constructed as follows.

$$\begin{aligned} x_1(t+1) &= x_1(t) + x_2(t) \\ x_2(t+1) &= \frac{1}{b}(x_1(t) + x_2(t)) + s(t) \\ s(t) &= 1 - \frac{1}{b}\beta^* w^b(t). \end{aligned}$$

For the initial condition $(x_1(1), x_2(1)) = (0, 1)$, the corresponding solution $x_2(t) \equiv u^b(j)$. Then using the state transition matrix, we can solve for the explicit solution for $x_2(t)$ as

$$\begin{aligned} x_2(1) =& 1 \\ x_2(t) =& \frac{1}{b}B^{t-2} + \sum_{\tau=1}^{t-2}\frac{1}{b}B^{t-2-\tau}(1-\beta^* w^b(\tau)) \\ & + (1-\beta^* w^b(t-1)) \quad t > 1, \end{aligned} \tag{9.64}$$

where $B = \frac{b+1}{b}$. Simplifying the expression for $x_2(t)$ for $t-1 > b$ and substituting $w^b(t) =$

$(1-\mathcal{C})t + \mathcal{C}\min(t,b)$ results in the following

$$\begin{aligned}
x_2(t) =& \frac{1}{b}B^{t-2}\Big(1 + \sum_{\tau=1}^{b} B^{-\tau}(1-\beta^*\tau) \\
&+ \sum_{\tau=b+1}^{t-2} B^{-\tau}(1-\beta^*((1-\mathcal{C})\tau + \mathcal{C}b))\Big) \\
&+ (1-\beta^*(t-1-\mathcal{C}(t-1)+\mathcal{C}b)).
\end{aligned}$$

Now we can use the series identities $\sum_{j=1}^{d} p^j = \frac{p-p^{d-1}}{1-p}$ and $\sum_{j=1}^{d} jp^j = \frac{p-(d+1)p^{d+1}+dp^{d+2}}{(1-p)^2}$ and simplify the terms to

$$x_2(t) = B^{t-2}(\beta^*(\mathcal{C}B^{1-b} - B) + B) + (1-\mathcal{C})\beta^*.$$

Thus, the above expression is the closed form solution for u^b when $j-1 > b$. We have already shown that the optimal utility rule u^b must be non-increasing in the proof of Theorem 9.2.1. This is only possible when $\beta^* \geq \frac{B^b}{B^b - \mathcal{C}}$. Therefore the optimal solution must be $\beta^* = \beta^b = \frac{B^b}{B^b-\mathcal{C}}$. Substituting for β^* in the expression in (9.64) and simplifying results in the closed form expression in (9.62) for u^b. It can be seen that u^b defined in (9.62) is indeed non-increasing. We lastly verify that the binding constraint for u^b is indeed when $z = b$ for any $y \geq 1$ and so β^b is tight. In (9.15), we examine the terms $w^b(z) - zu^b(y+1)$ for any $y \geq 1$. Note that $1 = w^b(z) - w^b(z-1) \geq u^b(y+1)$ when $z \leq b$ and $(1-\mathcal{C}) = w^b(z) - w^b(z-1) \leq u^b(y+1)$ when $z \geq b$ for any y. Thus the maximum $\max_z w^b(z) - zu^b(y+1)$ occurs when $z = b$, and we have shown the claim.

9.3.6 Proof of Theorem 9.2.3

In this section, we first provide upper bounds on the efficiency metric $\mathrm{PoU}^*(\mathcal{W};\kappa)$. To do this, we construct a game G such that for any utility design $\mathcal{U}$, rounds $\kappa \geq 1$, and curvature $\mathcal{C}$, we have that $\mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}};\kappa) \leq \mathrm{PoU}(G;\kappa) \leq 1 - \mathcal{C}/2$. Let $\mathcal{C}$ be the curvature and consider the b-covering rule w^b with $b = 1$ as in (9.61) with $w^b(2) = 2 - \mathcal{C}$. Additionally, let $u = \mathcal{U}(w^b)$ be the corresponding utility rule for a given utility design. A two-agent game G is constructed as follows. Let the resource set be $\mathcal{R} = \mathcal{R}_1 \cup \mathcal{R}_2 \cup \mathcal{R}_3$, where $\mathcal{R}_j$ is a set of resources such that the ratio of resources satisfies $|\mathcal{R}_1| = |\mathcal{R}_2| = u(2) \cdot |\mathcal{R}_3|$. If $u(2)$ is not a whole number, we can scale up $|\mathcal{R}_j|$ uniformly and round $u(2) \cdot |\mathcal{R}_3|$ to get arbitrarily close to the given ratio. Let $x = |\mathcal{R}_1|$. The action sets for the game construction the agents will be determined by u according to the following three cases: (**a**) $0 \leq u(2) \leq (1 - \mathcal{C})$, (**b**) $(1 - \mathcal{C}) \leq u(2) \leq 1$, and (**c**) $u(2) \geq 1$.

For case (a), Agent 1's actions are $\mathcal{A}_1 = \{a_1^{\varnothing}, a_1^1 = \mathcal{R}_1, a_1^2 = \mathcal{R}_2\}$. Agent 2's actions are $\mathcal{A}_2 = \{a_2^{\varnothing}, a_2^1 = \mathcal{R}_3, a_2^2 = \mathcal{R}_1\}$. The optimal allocation is $a^{\mathrm{opt}} = \{a_1^2, a_2^2\}$ resulting in a welfare of $2x$. An allocation that can occur after a one round walk is $a^{\mathrm{br}} = \{a_1^1, a_2^1\}$ resulting in a welfare of $(1 + u(2))x$. Therefore, $\mathrm{PoU}(G; 1) \leq \frac{(1+u(2))x}{2x} \leq 1 - \frac{\mathcal{C}}{2}$ by assumption of $u \leq 1 - \mathcal{C}$. Additionally, observe that a^{br} is a Nash equilibrium and therefore is still the resulting allocation after any number of additional rounds $\kappa \geq 1$. Therefore $\mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}}; \kappa) \leq \mathrm{PoU}(G; \kappa) \leq 1 - \frac{\mathcal{C}}{2}$ for this case of utility design.

For case (b), Agent 1's actions are $\mathcal{A}_1 = \{a_1^{\varnothing}, a_1^1 = \mathcal{R}_1, a_1^2 = \mathcal{R}_2\}$. Agent 2's actions are $\mathcal{A}_2 = \{a_2^{\varnothing}, a_2^1 = \mathcal{R}_3, a_2^2 = \mathcal{R}_1\}$. The optimal allocation is $a^{\mathrm{opt}} = \{a_1^2, a_2^2\}$ resulting in a welfare of $2x$. An allocation that can occur after a one-round walk is $a^{\mathrm{br}} = \{a_1^1, a_2^2\}$ resulting in a welfare of $w^b(2) \cdot x$. Therefore, $\mathrm{PoU}(G; 1) \leq \frac{w^b(2) \cdot x}{2x} = 1 - \frac{\mathcal{C}}{2}$. For $\kappa \geq 2$, there is a best response path that leads to the end state a^{br}. This is achieved by reaching $a' = \{a_1^1, a_2^1\}$ in the first round. As a' is a Nash action, the best response process can remains at a' for $\kappa - 1$ rounds and in the last round, switch to a^{br}. Therefore $\mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}}; \kappa) \leq \mathrm{PoU}(G; \kappa) \leq 1 - \frac{\mathcal{C}}{2}$ for this case.

For case (c), Agent 1's actions are $\mathcal{A}_1 = \{a_1^{\varnothing}, a_1^1 = \mathcal{R}_1, a_1^2 = \mathcal{R}_2\}$. Agent 2's actions are $\mathcal{A}_2 = \{a_2^{\varnothing}, a_2^1 = \mathcal{R}_1, a_2^2 = \mathcal{R}_3\}$. The optimal allocation is $a^{\mathrm{opt}} = \{a_1^2, a_2^2\}$ resulting in a welfare of $(1 + u(2))x$. An allocation that can occur after a one round walk is $a^{\mathrm{br}} = \{a_1^1, a_2^1\}$ resulting in a welfare of $w^b(2) \cdot x$. Therefore, $\mathrm{PoU}(G; 1) = \frac{w^b(2) \cdot x}{(1+u(2))x} \leq 1 - \frac{\mathcal{C}}{2}$ by assumption of $u(2) > 1$. Additionally, observe that a^{br} is a Nash equilibrium and therefore is still the resulting allocation after any number of additional rounds. Therefore $\mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}}; \kappa) \leq \mathrm{PoU}(G; \kappa) \leq 1 - \frac{\mathcal{C}}{2}$ for this case.

Since $u = \mathcal{U}(w^b)$ was chosen arbitrarily, we have that the upper bound holds for any utility design and we have shown that $\mathrm{PoU}^*(\mathcal{W};\kappa) \leq 1 - \mathcal{C}/2$. Furthermore, based on our game construction, the efficiency bounds hold even when we relax the class of best response dynamics that we consider. Since the game construction comprises of only two agents, allowing agents to best respond multiple times during a round or best respond out of order of round-robin does not improve the efficiency guarantees that result from the given game G.

Now we show that the upper bound $\mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathrm{CI}};\kappa) \leq (1+\mathcal{C})^{-1}$. As before, a game G is constructed such that under the common interest design CI, $\kappa \geq 1$, and curvature $\mathcal{C}$, we have that $\mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathrm{CI}};\kappa) \leq \mathrm{PoU}(G;\kappa) \leq (1+\mathcal{C})^{-1}$. Let G have n players with a resource set $\mathcal{R} = \mathcal{R}^{\mathrm{opt}} \cup \mathcal{R}^{\mathrm{both}} \cup \{r^n\}$ with $|\mathcal{R}^{\mathrm{opt}}| = n$ and $|\mathcal{R}^{\mathrm{both}}| = n-1$. Each agent i has three actions in its action set $\mathcal{A}_i = \{a_i^{\varnothing}, a_i^{\mathrm{br}}, a_i^{\mathrm{opt}}\}$. The resources are selected by the agents in the following manner: each resource $r_j^{\mathrm{opt}} \in \mathcal{R}^{\mathrm{opt}}$ is selected by agent j in action $a_j^{\mathrm{opt}} \ni r_j^{\mathrm{opt}}$ for all $1 \leq j \leq n$; each resource $r_j^{\mathrm{both}} \in \mathcal{R}^{\mathrm{both}}$ is selected by agent $j+1$ in action $a_{j+1}^{\mathrm{opt}} \ni r_j^{\mathrm{both}}$ and by agent j in action $a_j^{\mathrm{br}} \ni r_j^{\mathrm{both}}$ for all $1 \leq j \leq n-1$; agent n selects the resource in action $a_n^{\mathrm{br}} \in r^n$. Given a curvature $\mathcal{C}$, consider two b-covering welfare rules $w^b, w_2^b \in \mathcal{W}$ with curvature $\mathcal{C}$ such that $w^b(1) = 1$ and $w^b(2) = 1 - \mathcal{C}$, and $w_2^b(1) = \mathcal{C}$ and $w_2^b(2) = \mathcal{C}(1-\mathcal{C})$. For any $r \in \mathcal{R}^{\mathrm{both}} \cup \{r^n\}$, let the corresponding welfare rule be $w_r = w^b$ and for any $r \in \mathcal{R}^{\mathrm{opt}}$, let the corresponding welfare rule be $w_r = w_2^b$. Under this game construction it can be seen that under a^{br}, each resource $r \in \mathcal{R}^{\mathrm{both}} \cup \{r^n\}$ is selected by exactly one agent, resulting in a welfare of $W(a^{\mathrm{br}}) = n$; also, under a^{opt}, each resource $r \in \mathcal{R}^{\mathrm{both}} \cup \mathcal{R}^{\mathrm{opt}}$ is selected by exactly one agent, resulting in a welfare of $W(a^{\mathrm{br}}) = (n-1)(1+\mathcal{C}) + \mathcal{C}$. Assuming that a^{br} is the joint action that results after κ rounds, we have that $\mathrm{PoU}(G,\kappa) \leq \frac{n}{(n-1)(1+\mathcal{C})+\mathcal{C}}$. Limiting the number of agents $n \to \infty$ to infinity gives the result. To verify that a^{br} can result after κ rounds, observe that for agent 1 selecting a_1^{br} over a_1^{opt} results in a higher system welfare. After that, agents 2 through $n-1$ are indifferent between a_j^{br} and a_j^{opt} given that the previous $i < j$ players have selected a_i^{br}. Therefore, a^{br} is the resulting allocation after one round. Additionally, a^{br} can be seen to be Nash equilibrium for the common interest utility, so after any number of rounds κ, a^{br} is still a possible joint allocation that can result.

9.3.7 Proof of Theorem 9.2.4

We show the trade-offs in Theorem 9.2.4 that result from considering utility designs that maximize the one-round walk efficiency versus the price of anarchy. That $\mathrm{PoA}(\mathcal{G}_{\mathcal{W},\mathcal{U}_{\mathrm{PoA}}}) = 1 - \frac{1}{e}$ comes from Theorem 8.2.1 and $\mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}_1^*}; 1) = \frac{1}{2}$ comes from setting $\mathcal{C} = 1$ in Theorem 9.2.2. We show that $\mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}_{\mathrm{PoA}}}; 1) = 0$ in Lemma 9.3.3. From Lemma 9.3.4, we have that $\mathrm{PoA}(\mathcal{G}_{\mathcal{W},\mathcal{U}_1^*}) \geq \mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}_1^*}; 1) = \frac{1}{2}$, since $\mathcal{U}_1^*$ must be a non-increasing utility design, as shown in Section 9.3.4. To show that this lower bound is tight, consider the set covering welfare wsc. As seen in Theorem 9.2.5, the price of anarchy guarantee is $\frac{1}{2}$, and so $\mathrm{PoA}(\mathcal{G}_{\mathcal{W},\mathcal{U}_1^*}) \leq \mathrm{PoA}(\mathcal{G}_{wsc,\mathcal{U}_1^*(wsc)}) = \frac{1}{2}$ as well. Now we outline Lemma 9.3.3 and Lemma 9.3.4 below.

Lemma 9.3.3. *Suppose that $\mathcal{W}$ is the set of all possible submodular welfare rules and consider the utility design $\mathcal{U}_{\mathrm{PoA}}$. Then $\mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}_{\mathrm{PoA}}}; 1) = 0$.*

Proof. We construct a game $G \in \mathcal{G}_{\mathcal{W},\mathcal{U}_{\mathrm{PoA}}}$ to upper bound the one-round walk efficiency such that $\mathrm{PoU}(G; 1) = 0$. Since $\mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}_{\mathrm{PoA}}}; 1)$ is defined to be greater than 0, we have equality. Consider a game with n players as follows. We partition the resource set as $\mathcal{R} = \bigcup_{1 \leq j \leq n+1} \mathcal{R}_j$. Every resource $r \in \mathcal{R}$ is endowed the local welfare rule $w_r = w^b$ as the b-covering welfare rule with curvature of $\mathcal{C} = 1$ for some fixed $b \geq 1$, as defined in (9.61). The corresponding utility rule is $u_{\mathrm{PoA}} = \mathcal{U}_{\mathrm{PoA}}(w^b)$ is the following recursive expression from Lemma 8.3.3,

$$
\begin{aligned}
u_{\mathrm{PoA}}(1) &= 1 \\
u_{\mathrm{PoA}}(j+1) &= \frac{1}{b}[j u_{\mathrm{PoA}}(j) - \rho^b \min\{j, b\}] + 1,
\end{aligned}
$$

with $\rho^b = (1 - \frac{b^b e^{-b}}{b!})^{-1}$. The number of resources in each set is $|\mathcal{R}_1| = v$ and $|\mathcal{R}_{j+1}| \sim v \cdot u_{\mathrm{PoA}}(j)$ for $1 \leq j \leq n$ and for some $v \geq 0$. If $u_{\mathrm{PoA}}(j)$ is not a whole number, we can scale v up and round to get arbitrarily close to the correct ratio of resources. Agent i selects $\mathcal{R}_1 = a_i^{\mathrm{br}}$ and $\mathcal{R}_{i+1} = a_i^{\mathrm{opt}}$ in each of its actions. It can be verified that a^{br} is a joint action that can result after a one round walk. Therefore, the efficiency is upper bounded by

$$\mathrm{PoU}(G;1) \leq \frac{W(a^{\mathrm{br}})}{W(a^{\mathrm{opt}})} = \frac{vb}{v\sum_{1\leq i\leq n} u_{\mathrm{PoA}}(i)}.$$

Now we show that as we increase n, the series $\sum_{1\leq i\leq n} u_{\mathrm{PoA}}(i)$ diverges, and the efficiency can get arbitrarily bad as the number of agents increase. To construct the closed form expression of $u_{\mathrm{PoA}}(j)$, we construct the following LTV state space system with $u_{\mathrm{PoA}}(j) := x(t)$

$$\begin{aligned} x(t+1) &= A(t)x(t) + s(t) \qquad A(t) = \frac{t}{b} \\ s(t) &= 1 - \frac{\rho^b}{b}\min(t,b) \end{aligned}$$

Solving for the solution $x(t)$ using the state transition matrix with the initial condition $x(1) = 1$ results in the following expression

$$\begin{aligned} x(t) &= \prod_{\tau=1}^{t}\frac{\tau}{b} + \sum_{T=1}^{t-1}\big[\big(1 - \frac{\rho^b}{b}\min(t,b)\big)\prod_{\tau=T+1}^{t-1}\frac{\tau}{b}\big] \\ &= \frac{t!}{b^t}\left(1 + \sum_{T=1}^{t}\frac{b^T}{T!}\big(1 - \frac{\rho^b}{b}\min(t,b)\big)\right) \end{aligned}$$

If $t \geq b$, then

$$\begin{aligned} x(t) &= \frac{t!}{b^t}\Bigg(1 - (e^b - 1)(\rho - 1) + \sum_{T=t+1}^{\infty}\frac{b^T}{T!}\big(\rho^b - 1\big) + \\ &\sum_{T=1}^{b}\frac{b^T}{T!}\frac{\rho^b(b-T)}{b}\Bigg) \\ &= \frac{t!}{b^t}\sum_{T=t+1}^{\infty}\frac{b^T}{T!}\big(\rho^b - 1\big) \\ &\geq (\rho^b - 1)\frac{b}{t+1} \\ &\sim \mathrm{O}(\frac{1}{t}) \end{aligned}$$

The first equality results from splitting the summation and the second equality will be shown later. Since $x(t)$ is on the order of $\frac{1}{t}$, the series $\sum_{i=1}^{N} u_{\text{PoA}}(i)$ diverges and the claim is shown. Now we verify the equality

$$\begin{aligned}
\sum_{T=1}^{b} \frac{b^T}{T!} \frac{\rho^b(b-T)}{b} &= (e^b-1)(\rho^b-1)-1 \\
\sum_{T=1}^{b} \frac{b^T(b-T)}{bT!} &= \frac{1}{\rho^b}\left(e^b\rho^b - e^b - \rho^b\right) \\
\sum_{T=1}^{b} \frac{b^T}{T!} - \sum_{T=1}^{b} \frac{b^{T-1}}{(T-1)!} &= \left(e^b - 1 - e^b(1-\frac{b^b e^{-b}}{b!})\right) \\
\frac{b^b}{b!} - 1 &= \frac{b^b}{b!} - 1
\end{aligned}$$

The last equality results from recognizing the terms on the left hand side as a telescoping sum. $\square$

Lemma 9.3.4. *Let $\mathcal{W} = \{w^1, \ldots, w^m\}$ be a set of welfare rules and $\mathcal{U}$ be a utility design such that $u^\ell = \mathcal{U}(w^\ell)$ is non-increasing for any $1 \leq \ell \leq m$. Then $\text{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}}; \kappa) \leq \text{PoA}(\mathcal{G}_{\mathcal{W},\mathcal{U}})$ for any $\kappa \geq 1$.*

Proof. We show this claim by a game construction, where a Nash equilibrium with the efficiency arbitrarily close to $\text{PoA}(\mathcal{G}_{\mathcal{W},\mathcal{U}})$ is reachable by a one-round walk. Let $\varepsilon_1 > 0$. Note that $\text{PoA}(\mathcal{G}_{\mathcal{W},\mathcal{U}}n)$ is non-increasing in n and lower bounded by 0. Therefore $\text{PoA}(\mathcal{G}_{\mathcal{W},\mathcal{U}}n)$ is a convergent sequence in n and for any ε_1, there exists an $N_1 \in \mathbb{N}$ such that $\text{PoA}(\mathcal{G}^{N_1}_{\mathcal{W},\mathcal{U}}) - \text{PoA}(\mathcal{G}_{\mathcal{W},\mathcal{U}}) \leq \varepsilon_1$.

Generalizing [115, Theorem 2] to a set of welfare rules provides a characterization of the price of anarchy as $\mathrm{PoA}(\mathcal{G}^{N_1}_{\mathcal{W},\mathcal{U}}) = \min_{1\leq \ell \leq m} \frac{1}{Q^\ell}$ with

$$
\begin{aligned}
Q^\ell = & \max_{\theta(y,x,z)} \sum_{y,x,z} w^\ell(z+x)\theta(y,x,z) \\
\text{s.t.} \ & \sum_{y,x,z} [y u^\ell(y+x) - z u^\ell(y+x+1)]\theta(y,x,z) \geq 0 \\
& \sum_{y,x,z} w^\ell(y+x)\theta(y,x,z) = 1 \\
& \theta(y,x,z) \geq 0,
\end{aligned}
\tag{9.65}
$$

where $y, x, z, \in \mathbb{N}$ with $1 \leq y+x+z \leq N_1$. For the $\ell^* = \arg\min_{1\leq \ell\leq m} \frac{1}{Q^\ell}$ that achieves the minimum, we refer to $w \equiv w^{\ell^*}$, $u \equiv u^{\ell^*}$ for ease of notation and refer to $\Theta(y,x,z)$ to denote the corresponding optimal variables for $\theta(y,x,z)$ of the linear program. We construct a matching game G as follows. Let $N_2 > N_1$ be the number of agents in the game and $D = N_2 + y + x - 1$. For each y, x, z pair and $1 \leq k \leq D$, we construct a set of resources $\mathcal{R}axb$ with $|\mathcal{R}axb| = \Theta(y,x,z)/D$. Each agent i has three actions in its action set $\mathcal{A}_i = \{a_i^{\varnothing}, a_i^{\mathrm{ne}}, a_i^{\mathrm{opt}}\}$. Each agent i selects $\{\mathcal{R}axb\}_{i\leq k\leq y+x+i-1}$ in a_i^{ne} for each pair y, z, x. If $y+z+x \leq i \leq N_2$, agent i selects $\{\mathcal{R}axb\}_{i-z\leq k\leq x+i-1}$ in a_i^{opt} for each pair y, z, x. Otherwise for $1 \leq i \leq y+z+x-1$, $a_i^{\mathrm{opt}} = a_i^{\varnothing}$ and agent i doesn't select any resources in a_i^{opt}.

We first confirm that the action a^{ne} is indeed a Nash equilibrium. Showing this for the first $y+x+z-1$ agents is trivial, since no resources are selected in a_i^{opt}. For the rest of the agents, the utility difference of a unilateral deviation to a_i^{opt} from a_i^{ne} is

$$\begin{aligned}
&U_i(a^{\text{ne}}) - U_i(a_i^{\text{opt}}, a_{-i}^{\text{ne}}) \\
&\geq \sum_{r \in a_i^{\text{ne}}} u_r(|a^{\text{ne}}|_r) - \sum_{r \in a_i^{\text{opt}}} u_r(|(a_i^{\text{opt}}, a_{-i}^{\text{ne}})|_r) \\
&\geq \sum_{y,x,z} [(y+x)u(|a^{\text{ne}}|_r) - \\
&\quad xu_r(y+x) - zu(y+x+1)] \cdot |\mathcal{R}axb| \\
&\geq \frac{1}{D} \sum_{y,x,z} [yu(y+x) - zu(y+x+1)]\Theta(y,x,z) \\
&\geq 0.
\end{aligned}$$

The first inequality comes from the definitions of the utility function. The second inequality comes from counting the resources that are selected in the either a_i^{ne} or a_i^{opt} by the agent in each set of resources in $\mathcal{R}axb$. The third inequality arises from the fact that $|a^{\text{ne}}|_r \leq y+x$, and since u is assumed to be non-increasing, $u(|a^{\text{ne}}|_r) \geq u(y+x)$. The fourth inequality comes from the fact that since $\Theta(y,x,z)$ has to satisfy the inequality constraint in (9.65) to be feasible. Similarly, in a one-round walk, the best response for the first $y+x+z-1$ is a_i^{ne}. The best response for the other agents during the one-round walk is also a_i^{ne}, since

$$\begin{aligned}
&U_i(a_{j<i}^{\text{ne}}, a_i^{\text{ne}}, a_{j>i}^{\varnothing}) - U_i(a_{j<i}^{\text{ne}}, a_i^{\text{opt}}, a_{j>i}^{\varnothing}) \\
&= \sum_{y,x,z} [\sum_{j=1}^{y+x} u(i) - xu(y+x) - zu(y+x+1)]|\mathcal{R}axb| \\
&\geq \frac{1}{D} \sum_{y,x,z} [yu(y+x) - zu(y+x+1)]\Theta(y,x,z) \\
&\geq 0.
\end{aligned}$$

Therefore, the Nash equilibrium a^{ne} is reached from an empty configuration in one-round. Additionally, since a^{ne} is a Nash equilibrium, the resulting action state after κ rounds can also be a^{ne}. Therefore in this game, $\text{PoU}(G;\kappa) \leq \text{PoA}(G)$. Now we calculate the efficiency of the

Nash equilibrium $W(a^{\mathrm{ne}})$ with respect to $W(a^{\mathrm{opt}})$. We have that

$$W(a^{\mathrm{ne}}) = \sum_{y,x,z} w(y+x)\cdot\Theta(y,x,z)\frac{N_2-2(y+x-1)}{N_2}+$$
$$2\sum_{i=1}^{y+x-1} w(i)\frac{\Theta(y,x,z)}{N_2} = 1+\mathrm{O}(\frac{1}{N_2}),$$

where, since $\Theta(y,x,z)$ is feasible, then it satisfies the equality constraint that $\sum_{y,x,z} w(y+x)\Theta(y,x,z) = 1$. $\mathrm{O}(\frac{1}{N_2})$ reflects that the rest of the terms are on order of $1/N_2$. Similarly,

$$W(a^{\mathrm{opt}}) = \sum_{y,x,z} w(z+x)\cdot\Theta(y,x,z)\frac{N_2-3(z+x-1)}{N_2}+$$
$$2\sum_{i=1}^{z+x-1} w(i)\frac{\Theta(y,x,z)}{N_2} = \mathrm{PoA}(\mathcal{G}^{N_1}_{\mathcal{W},\mathcal{U}})^{-1}+\mathrm{O}(\frac{1}{N_2}),$$

where, since $\Theta(y,x,z)$ is optimal, then $\sum_{y,x,z} w(z+x)\Theta(y,x,z) = \mathrm{PoA}(\mathcal{G}^{N_1}_{\mathcal{W},\mathcal{U}})^{-1}$. For any ε_2, we can choose N_2, such that $\mathrm{O}(\frac{1}{N_2}) \leq \varepsilon_2$, so $\mathrm{PoA}(G) \leq \mathrm{PoA}(\mathcal{G}^{N_1}_{\mathcal{W},\mathcal{U}})^{-1}+\varepsilon_2$. To put everything together, we have that

$$\mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}};\kappa) \leq \mathrm{PoU}(G;\kappa) \leq \mathrm{PoA}(G)$$
$$\leq \mathrm{PoA}(\mathcal{G}^{N_1}_{\mathcal{W},\mathcal{U}})^{-1}+\varepsilon_2 \leq \mathrm{PoA}(\mathcal{G}_{\mathcal{W},\mathcal{U}})+\varepsilon_1+\varepsilon_2,$$

and since ε_1 and ε_2 are arbitrary, we have the result for any rounds $\kappa \geq 1$. □

9.3.8 Proof of Theorem 9.2.5

To characterize the Pareto optimal frontier in (9.22), we first simplify the linear program in (9.53) with respect to the set covering welfare rule.

Lemma 9.3.5. *Let $\mathcal{W} = \{wsc\}$, where wsc is the set covering welfare rule defined in* (9.21), *and $\mathcal{U} = \{u\}$ be the corresponding utility rule. Then the one-round walk efficiency guarantee is*

$$\mathrm{PoU}(\mathcal{G}_{wsc,u};1)^{-1} = \sum_{i\in\mathbb{N}} u(i) - \min_{i\in\mathbb{N}} u(i) + 1. \tag{9.66}$$

Proof. Examine the dual program in (9.53) with substituting the set covering welfare defined in (9.21). Under the substitution, the dual constraint for a given z, y simplifies to

$$\beta \geq \sum_{i=1}^{y} u(i) - z \min_{1\leq i\leq y+1} u(i) + \min(1,z).$$

We have applied the fact $wsc(j) = \min(1,j) = 1$ when $j \geq 1$ and $\mathrm{H} = \max_{\mathrm{j}\in\mathbb{N}} \mathrm{wsc(j)/j} = 1$ to the dual constraint. Observe that the binding constraint occurs when we limit $y \to \infty$ and set $z = 1$ (and not $z = 0$ since $u(1) = 1$, the term $1 - \min_j u(j) \geq 0$). Under those binding constraints, $\mathrm{PoU}(\mathcal{G}_{wsc,u};1)^{-1} = \beta$, where β matches the given expression in (9.66). □

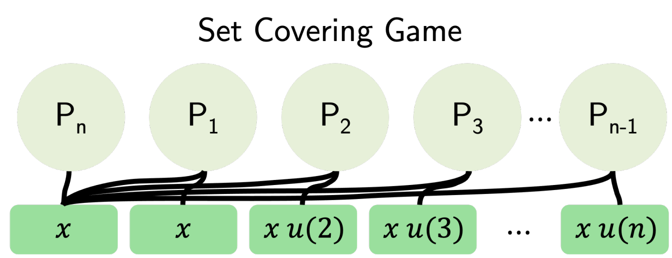

Figure 9.6: The worst case game construction achieving the one-round walk guarantee dictated by Lemma 9.3.5. In this game, all the agents can either stack on the first resource set or spread out.

To characterize the trade-off, we now provide an explicit expression of Pareto optimal utility rules, i.e., the utility rules u that satisfy either $\mathrm{PoU}(\mathcal{G}_{wsc,u};1) \geq \mathrm{PoU}(\mathcal{G}_{wsc,u'};1)$ or $\mathrm{PoA}(\mathcal{G}_{wsc,u}) \geq \mathrm{PoA}(\mathcal{G}_{wsc,u'})$ for all $u' \neq u$.

Lemma 9.3.6. *For a given* $\mathcal{X} \geq 0$*, a utility rule* $u^{\mathcal{X}}$ *that satisfies* $\mathrm{PoA}(\mathcal{G}_{wsc,u}) \geq 1/(1+\mathcal{X})$ *while maximizing* $\mathrm{PoU}(\mathcal{G}_{wsc,u};1)$ *is defined as in the following recursive formula:*

$$u^{\mathcal{X}}(1) = 1 \tag{9.67}$$
$$u^{\mathcal{X}}(j+1) = \max\{ju^{\mathcal{X}}(j) - \mathcal{X}, 0\}.$$

Proof. According to [115, Theorem 2], the price of anarchy can be written as

$$\frac{1}{\text{PoA}(\mathcal{G}_{wsc,u})} = 1 + \max_{1 \le j \le n-1} \{(j+1)u(j+1) - 1,$$
$$ju(j) - u(j+1), ju(j+1)\}. \tag{9.68}$$

We first show that if u is Pareto optimal, then it must also be non-increasing. Otherwise, we show that another u' exists that achieves at least the same one round efficiency, but a higher price of anarchy, contradicting our assumption that u is Pareto optimal. Assume, by contradiction, that there exists u that is Pareto optimal and not non-increasing, i.e., there exists a $J \ge 1$, in which $u(J) < u(J+1)$. Notice that switching the value $u(J)$ with $u(J+1)$ results in an unchanged one round efficiency according to (9.66) in Lemma 9.3.5 if $J > 1$. We show that u' with the values at J and $J+1$ switched has a higher price of anarchy than u.

For any $1 \le J \le n-1$, the expressions from (9.68) that include $u(J)$ or $u(J+1)$ are

$$Ju(J), \quad (J+1)u(J+1), \quad (J-1)u(J-1) - u(J) + 1,$$
$$Ju(J) - u(J+1) + 1, \quad (J+1)u(J+1) - u(J+2) + 1,$$
$$(J-1)u(J) + 1, \quad Ju(J+1) + 1.$$

After switching, the relevant expressions for u' are

$$Ju(J+1),\ (J+1)u(J),\ (J-1)u(J-1) - u(J+1) + 1,$$
$$Ju(J+1) - u(J) + 1, \quad (J+1)u(J) - u(J+2) + 1,$$
$$(J-1)u(J+1) + 1, \quad Ju(J) + 1.$$

Since $u(J) < u(J+1)$, switching the values results in a strictly looser set of constraints, and the value of the binding constraint in (9.68) for u' is at most the value of the binding constraint for u. Therefore $\text{PoA}(u) \le \text{PoA}(u')$. Note that if $J = 1$, switching J and $J+1$ and scaling down appropriately so $u'(1) = 1$, then $\text{PoU}(u') > \text{PoU}(u)$ as well. This contradicts our assumption that u is Pareto optimal.

Now we restrict our focus u that are non-increasing. Under this assumption, the price of anarchy is

$$\frac{1}{\mathrm{PoA}(wsc, u)} = 1 + \max_{1\leq j\leq n-1}\{ju(j) - u(j+1), (n-1)u(n)\},$$

as detailed in Corollary 2 in [115]. Let

$$\mathcal{X}_u = \max_{1\leq j\leq n-1}\{ju(j) - u(j+1), (n-1)u(n)\}. \tag{9.69}$$

For u to be Pareto optimal, we claim that $ju(j) - u(j+1) = \mathcal{X}_u$ must hold for all j. Consider any other u' with $\mathcal{X}_u = \mathcal{X}_{u'}$. It follows that $\mathrm{PoA}(u) = \mathrm{PoA}(u') = 1/(1+\mathcal{X}_u)$. By induction, we show that $u(j) \leq u'(j)$ for all j. The base case is satisfied, as $1 = f(1) \leq f'(1) = 1$. Under the assumption $u(j) \leq u'(j)$, we also have that

$$ju(j) - \mathcal{X}_u = u(j+1) \leq u'(j+1) = ju'(j) - \mathcal{X}_u^j, \tag{9.70}$$

where $\mathcal{X}_{u'}^j = ju'(j) - u'(j+1) \leq \mathcal{X}_{u'}$ by definition in (9.69), and so $u(j) \leq u'(j)$ for all j. Therefore the summation $\sum_{i\in\mathbb{N}} u(i) - \min_{i\in\mathbb{N}} u(i)$ in (9.66) is diminished and $\mathrm{PoU}(u) \geq \mathrm{PoU}(u')$, proving our claim. As u must satisfy $u(j) \geq 0$ for all j to be a valid utility rule, $u(j+1)$ is set to be $\max\{ju(j) - \mathcal{X}, 0\}$. Then we get the recursive definition for the maximal $u^{\mathcal{X}}$. Finally, we note that for infinite n, $\mathcal{X} \leq \frac{1}{e-1}$ is not achievable, as shown in [112]. $\square$

With the two previous lemmas, we can move to proving Theorem 9.2.5 We first characterize a closed form expression of the maximal utility rule $u^{\mathcal{X}}$, which is given in Lemma 9.3.6. We fix $\mathcal{X}$ so that $\mathrm{PoA}(u^{\mathcal{X}}) = \frac{1}{\mathcal{X}+1} = Q$. To calculate the expression for $u^{\mathcal{X}}$ for a given $\mathcal{X}$, a corresponding time varying, discrete time system to (9.67) is constructed as follows.

$$\begin{aligned} x(t+1) &= tx(t) - \mathcal{X}, \\ y(t) &= \max\{x(t), 0\}, \\ x(1) &= 1, \end{aligned}$$

where $y(t) \equiv u^{\mathcal{X}}(j)$ corresponds to the utility rule. Solving for the explicit solution for $y(t)$ using the state-transition matrix gives

$$\begin{aligned} &y(1) = 1 \\ &y(t) = \max\Big[\prod_{\ell=1}^{t-1} \ell - \mathcal{X}\big(\sum_{\tau=1}^{t-2} \prod_{\ell=\tau+1}^{t-1} \ell\big) - \mathcal{X}, 0\Big] \quad t > 1. \end{aligned}$$

Simplifying the expression and substituting for $u^{\mathcal{X}}(j)$ gives

$$u^{\mathcal{X}}(j) = \max\Big[(j-1)!(1 - \mathcal{X}\sum_{\tau=1}^{j-1} \frac{1}{\tau!}), 0\Big] \quad j \geq 1.$$

Substituting the expression for the maximal $u^{\mathcal{X}}$ into (9.66) gives the one round efficiency. Notice that for $\mathcal{X} \geq \frac{1}{e-1}$, $\lim_{j\to\infty} u^{\mathcal{X}}(j) = 0$, and therefore $\min_j u^{\mathcal{X}}(j) = 0$. Shifting the variables $j' = j + 1$, we get the statement in (9.22).

9.3.9 Supermodular welfare rules

In this section, efficiency of one-round walks are examined for classes of *supermodular games*. Supermodular games are an important sub-class of resource allocation games, in which there is a surplus of added system welfare when a resource is covered by more than one agent. Applications of supermodular games include clustering and power allocation in networks [119]. A welfare rule w is deemed to be supermodular if $w(j) - w(j-1)$ is increasing and non-negative for all $j \geq 1$. Interestingly, for supermodular games, the utility designs that both maximize the one-round efficiency and price of anarchy include the constant utility design $\mathcal{U}(w) = u_{\mathbf{1}}$ in which $u_{\mathbf{1}}(j) = 1$

for all $j \in \mathbb{N}$, and the *Shapley* utility design $\mathcal{U}(w) = u_{\text{shap}}$ in which $u_{\text{shap}}(j) = w(j)/j$ for all $j \in \mathbb{N}$. Furthermore, the optimal one-round and price of anarchy guarantees are equivalent, as seen below.

Proposition 9.3.1. *Consider a set of supermodular welfare rules $\mathcal{W} = \{w^1, \dots, w^m\}$ with $w^\ell(1) = 1$. If the number of agents is fixed to n, then the optimal one-round and price of anarchy guarantees are as follows*

$$\sup_{\mathcal{U}} \text{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}} n; 1) = \sup_{\mathcal{U}} \text{PoA}(\mathcal{G}_{\mathcal{W},\mathcal{U}} n) = \min_{1 \le \ell \le m} \frac{n}{w^\ell(n)}. \tag{9.71}$$

Furthermore any utility design in which $u^\ell = \mathcal{U}(w^\ell)$ is non-decreasing and satisfies $u^\ell(1) = 1$ and $\sum_{i=1}^{j} u^\ell(i)/w(j) \le 1$ for all $1 \le j \le n$ achieves the optimal one round efficiency guarantee.

Proof. The fact that $\sup_{\mathcal{U}} \text{PoA}(\mathcal{G}_{\mathcal{W},\mathcal{U}} n) = \min_{1 \le \ell \le m} \frac{n}{w^\ell(n)}$ comes from applying the result in [119, Theorem 4] to a class of welfare rules. Thus, we first show that $\sup_{\mathcal{U}} \text{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}} n; 1) \le \min_{1 \le \ell \le m} \frac{n}{w^\ell(n)}$ through a game construction that is valid for any utility design $\mathcal{U}$. Let $w^* = \arg\min_{1 \le \ell \le m} \frac{n}{w^\ell(n)}$ be the welfare rule that attains the minimum. Let the game G have n agents with agent i having the action set $\mathcal{A}_i = \{a_i^\varnothing, a_i^{\text{br}}, a_i^{\text{opt}}\}$. There are $n+1$ resources which are all endowed with the welfare rule $w_r = w^*$ for all $r \in \mathcal{R}$, with agent i either selecting $a_i^{\text{br}} = \{r_{i+1}\}$ or $a_i^{\text{opt}} = \{r_1\}$. Under any utility rule u, each agent i is indifferent to choosing a_i^{br} or a_i^{opt} if no other agents $j \ne i$ have selected r_1 through a_j^{opt}. Thus a^{br} can result after a κ-round walk with a welfare of $W(a^{\text{br}}) = n$. The welfare of the optimal allocation a^{opt} is $W(a^{\text{opt}}) = w(n)$. Therefore, for any utility design $u = \mathcal{U}(w^*)$, the efficiency is bounded by $\text{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}} n; 1) \le \text{PoU}(G; 1) = \frac{n}{w^*(n)}$. We remark that $\text{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}} n; \kappa) \le \text{PoU}(G; \kappa) = \frac{n}{w^*(n)}$ as well, since a^{br} is a Nash equilibrium.

Now we show that for a utility design $\mathcal{U}$, such that the utility rule $u^\ell = \mathcal{U}(w^\ell)$ is non-decreasing and satisfies $\sum_{i=1}^{j} u^\ell(i)/w(j) \le 1$ and $u^\ell(1) = 1$ for every j and ℓ, the one-round efficiency is lower bounded by $\text{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}} n; 1) \ge \min_{1 \le \ell \le m} \frac{n}{w^\ell(n)}$. To do this, we can use a modified version of the linear program in (9.53) in which $\text{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}} n; 1) \ge \min_{1 \le \ell \le m} \frac{1}{\beta^\ell}$, where

$\beta^\ell \in \mathbb{R}$ is the solution to

$$\begin{aligned}
&\beta^\ell = \min \quad \beta \quad \text{subject to:} \\
&\beta w^\ell(y) \geq \mathrm{H}^\ell \left(\sum_{\mathrm{i=1}}^{\mathrm{y}} \mathrm{u}^\ell(\mathrm{i}) - \mathrm{z} \min_{1 \leq \mathrm{i} \leq \mathrm{y}+1} \mathrm{u}^\ell(\mathrm{i}) \right) + \mathrm{w}^\ell(\mathrm{z}) \\
&\text{for all } 0 \leq z \leq n \text{ and } 1 \leq y \leq n,
\end{aligned}$$

where the linear program is a lower bound since we consider tighter constraints that allow y and z to to range from 1 to n. Since w^ℓ is supermodular, $\mathrm{H}^\ell = \mathrm{w}^\ell(\mathrm{n})/\mathrm{n}$ and assuming u^ℓ is non-decreasing, $\min_{1 \leq i \leq y+1} u^\ell(i) = u^\ell(1) = 1$. Thus, we can simplify the constraint as

$$\beta w^\ell(y) \geq \frac{w^\ell(n)}{n} \sum_{i=1}^{y} u^\ell(i) - \frac{w^\ell(n)}{n} z + w^\ell(z) \tag{9.72}$$

With this, we observe that $w^\ell(z) - z \cdot w^\ell(n)/n$ is convex in z. So the binding constraint for z occurs at either the end point $z = 0$ or $z = n$ and the terms can be cancelled out. Additionally, $\max_y \sum_{i=1}^{y} u^\ell(i)/w^\ell(y) = 1$ occurs at the binding constraint $y = 1$, by assumption that $\sum_{i=1}^{j} u^\ell(i)/w(j) \leq 1$ for all $1 \leq j \leq n$. Therefore, $\beta^\ell = \mathrm{H}^\ell = \mathrm{w}^\ell(\mathrm{n})/\mathrm{n}$ under the binding constraint of $y = 1$ and $z = 0$ (or n) and we indeed have that $\mathrm{PoU}(\mathcal{G}_{\mathcal{W},\mathcal{U}} n; 1) \geq \min_{1 \leq \ell \leq m} \frac{n}{w^\ell(n)}$. □

We remark that the utility rules $u_{\mathbf{1}}$ and u_{shap} both satisfy the assumptions in Proposition 9.3.1. Furthermore, we remark that the optimal one-round guarantees match the optimal κ-round and price of anarchy guarantees, and so running the best-response process for further rounds does not increase the resulting efficiency guarantees.

Part III

Conclusion

Chapter 10

Conclusions and future outlook

10.1 Decision-based mechanisms

In the first part of this dissertation, we considered mechanisms that modify the decision space in multi-agent systems. Our main focus was on understanding the role that information plays in competitive interactions, particularly its impact on the strategic behaviour of the competitors. We based our analysis on variations of the Colonel Blotto game [120], a popular framework for studying competitive resource allocation in adversarial settings. The majority of our results treat a player's equilibrium payoff in the nominal game model as the benchmark and compare against the player's equilibrium payoff in a perturbed model to analyze the impact of such perturbations on the strategic outcomes of the system.

In our first set of results within this context, we sought to identify the intrinsic value that competitors should associate with information on different parameters of the game, namely the battlefield values and the opponent's budget. We did so under the framework of General Lotto games by comparing a player's equilibrium payoff in the complete information game, and in an incomplete information game where the player only knows the prior distribution of either the battlefield values or the opponent's budget. Our analysis offers the following intuitive observations: first, that information on battlefield values is more valuable to the weaker competitor than the stronger one; and, second, that in several parameter regimes we characterize information on the opponent's budget offers no benefit to a competitor.

Our second set of results consider whether a competitor can strategically use information to shift the strategic outcome in her favour. We do so by studying two-stage formulations of two- and three-player General Lotto games where players have the opportunity to pre-emptively announce strategic information to their opponents. Such opportunities contrast with previously studies of pre-emption in the literature on General Lotto games in that they are viable even in the absence of mechanisms for coordination between players like alliances. We first studied concessions, and showed that a player can improve her equilibrium payoff by conceding battlefields to her opponent in three-player General Lotto games. We then considered pre-allocations, and observed that pre-allocated resources are never more than half as effective as real-time resources in two-player General Lotto games.

10.1.1 Future directions

Networks of conflict. Our analysis of three-player General Lotto games in which two competitors face a common adversary prompts the study of more general networks of conflict. A possible representation of such a conflict could be a graph where the nodes are the competitors and edges represent conflicts between competitors. Already, preliminary studies of networks of conflicts have been considered, e.g., in []. An interesting study of such problem settings should analyze the importance of coordination between allies in such networks, by comparing the strategic outcomes and resulting cumulative payoff across members of an alliance under coordinated and uncoordinated decision making. Such a study would be at the intersection of Colonel Blotto games and *network games* [].

Another important direction related to networks of conflict examines the effect of overlapping responsibilities in the coordination of teams against a common adversary. Existing results in Colonel Blotto games [] show that the decision space is complex, as – contrary to the intuition – the team's equilibrium payoff is non-monotonic in the team members' endowment, and the probability distribution of the team's overall allocation is a convolution over the probability distributions selected by its members. Future studies of tractable game models may allow for a more nuanced study of this problem setting.

Heterogeneous resources and conflicts. As previously discussed, another interpretation of our results on General Lotto games with pre-allocations compares the effectiveness of *fixed* and *mobile* resources, i.e., those resources that cannot move freely at the time of engagement, and those that can. Future work should explore other models of resource heterogeneity (e.g., rock-paper-scissors forces that vary in effectiveness against one another, decoy or "paper tiger" forces that are only revealed during the time of engagement), and heterogeneous battlefields (e.g., simultaneous conflicts with different contest success functions).

Beyond Bayesian formulations. Any study based on Bayesian formulations and Bayesian outcomes must consider the complexity of the emergent solution concepts. Indeed, it is strongly debated whether human reasoning and Bayesian reasoning is consistent. Furthermore, Bayesian formulations rely on known prior distributions over the system state, which may be unrealistic in many relevant problem settings. Future work should seek to relax the Bayesian assumption, perhaps exploring prior-free, Stackleberg formulations as in [], or examining more tractable equilibrium notions as in [].

10.2 Preference-based mechanisms

In the second part of this dissertation, we considered mechanisms that modify the preference structure in multi-agent systems. Our main focus was on studying how the system-level performance can be optimized through structured interventions that influence the agents' utility functions. Our analytical framework for this research direction was based on the congestion game model [75] (as well as the analogous resource allocation game model). We considered taxes implemented at the resource level that are anonymous and depend only on locally available information (i.e., the resource cost function and the number of agents selecting the resource).

Our first research direction within this context was to optimize the worst-case equilibrium efficiency (Price of Anarchy) across families of games. Inspired by established smoothness arguments and primal-dual techniques, we derived a methodology based on tractable linear programs to compute the optimal taxation mechanism and its associated Price of Anarchy guarantees. We also provided techniques to move beyond linear programming to more tractable methodologies while retaining near-optimal Price of Anarchy guarantees. This last set of results also provides universal guarantees on the optimal Price of Anarchy, which – for the problems we consider – are competitive with even the best achievable approximation ratios among polynomial time, centralized algorithms.

The second research direction in this part explored the potential consequences of optimizing for the worst-case equilibrium efficiency. We considered the impact of such a design approach on the best-case equilibrium efficiency (Price of Stability) and the transient efficiency (Price of Urgency). In the respective settings considered, our results show that the taxes that optimize the Price of Anarchy have corresponding Price of Stability guarantees equal to the Price of Anarchy, as well as arbitrarily poor Price of Urgency. We then proposed techniques for characterizing the Anarchy-Stability and Anarchy-Urgency trade-off curves.

10.2.1 Future directions

Global, local and other levels of information. Parallel research efforts have investigated the design of global taxation rules where the design of tax functions is conditioned on all parameters of a game instance. Interestingly, our results suggest that transitioning from local to global taxation rules may not provide significant reductions of the achievable Price of Anarchy. Future work should seek to establish whether local taxation rules remain competitive under other performance metrics including the Price of Stability and Price of Urgency. Furthermore, analyzing whether a tension between the Price of Anarchy and other performance metrics persists under global taxation rules is an open and interesting problem. Finally, global and local information represent only two extreme perspectives on the design of taxes. A fruitful direction for future study is in understanding the value of different pieces of information under this paradigm.

Persistence of trade-offs beyond congestion games. [121] and [122] also investigate trade-offs between the Price of Anarchy and Price of Stability, albeit in distinct classes of problems. However, they all report findings analagous to ours: When the Price of Anarchy is optimized, the Price of Anarchy and Price of Stability are equal. While it is immediate that the Price of Stability can never exceed the Price of Anarchy, it is unclear whether these two metrics must always be in tension with one another. A relevant research direction is to understand the broader class of problems for which the Price of Anarchy can only be optimized to the detriment of the Price of Stability, and likewise for the Price of Urgency.

Beyond worst-case analysis. Many of the results in this part represent analyses that venture beyond the worst-case perspective, and ask whether designing systems under such a pessimistic approach is advisable. A recent and significant thrust in computer science advocates for more analyses beyond the worst case, as worst-case examples can often be esoteric and non-representative of the underlying problems of interest. To date, numerous positive results exist in this direction, including the canonical result on the run-time of the simplex method for solving linear programs. Further thrusts beyond the worst case in our context might include restricting the family of games to measure equilibrium efficiency guarantees under "realistic" network structures, measure-theoretic analyses of equilibrium efficiency (e.g., smoothed analysis) and introducing data-driven techniques into the analysis (e.g., Monte Carlo methods).

Bibliography

[1] M. Broom and J. Rychtár, *Game-theoretical models in biology.* CRC Press, 2013.

[2] J. W. Weibull, *Evolutionary game theory.* MIT press, 1997.

[3] K. Leyton-Brown and Y. Shoham, *Essentials of game theory: A concise multidisciplinary introduction*, *Synthesis lectures on artificial intelligence and machine learning* **2** (2008), no. 1 1–88.

[4] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic game theory.* Cambridge University Press, 2007.

[5] A. M. Colman, *Game theory and its applications: In the social and biological sciences.* Psychology Press, 2013.

[6] J. D. Morrow, *Game theory for political scientists.* Princeton University Press, 1994.

[7] T. Börgers and D. Krahmer, *An introduction to the theory of mechanism design.* Oxford University Press, USA, 2015.

[8] T. Başar and G. J. Olsder, *Dynamic noncooperative game theory.* SIAM, 1998.

[9] S. H. Strogatz, *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering.* CRC press, 2018.

[10] F. Dörfler and F. Bullo, *On the critical coupling for kuramoto oscillators*, *SIAM Journal on Applied Dynamical Systems* **10** (2011), no. 3 1070–1099.

[11] F. Dorfler and F. Bullo, *Synchronization and transient stability in power networks and nonuniform kuramoto oscillators*, *SIAM Journal on Control and Optimization* **50** (2012), no. 3 1616–1642.

[12] N. Chopra and M. W. Spong, *On exponential synchronization of kuramoto oscillators*, *IEEE transactions on Automatic Control* **54** (2009), no. 2 353–357.

[13] R. Xiao, J. Li, and T. Chen, *Modeling and intelligent optimization of social collective behavior with online public opinion synchronization*, *International Journal of Machine Learning and Cybernetics* **10** (2019), no. 8 1979–1996.

[14] J. R. Marden and J. S. Shamma, *Revisiting log-linear learning: Asynchrony, completeness and payoff-based implementation*, *Games and Economic Behavior* **75** (2012), no. 2 788–808.

[15] H. P. Young, *Strategic learning and its limits.* Oxford University Press, 2004.

[16] R. Srikant and T. Başar, *The mathematics of Internet congestion control.* Springer, 2004.

[17] D. K. Molzahn, F. Dörfler, H. Sandberg, S. H. Low, S. Chakrabarti, R. Baldick, and J. Lavaei, *A survey of distributed optimization and control algorithms for electric power systems*, *IEEE Transactions on Smart Grid* **8** (2017), no. 6 2941–2962.

[18] K. Scaman, F. Bach, S. Bubeck, Y. T. Lee, and L. Massoulié, *Optimal algorithms for smooth and strongly convex distributed optimization in networks*, in *International Conference on Machine Learning (ICML)*, pp. 3027–3036, PMLR, 2017.

[19] R. B. Myerson, *Optimal auction design*, *Mathematics of operations research* **6** (1981), no. 1 58–73.

[20] D. Kovenock and B. Roberson, *Generalizations of the general lotto and colonel blotto games*, *Economic Theory* **71** (2021), no. 3 997–1032.

[21] S. Hart, *Discrete colonel blotto and general lotto games*, *International Journal of Game Theory* **36** (2008), no. 3 441–460.

[22] Y. Mo, T. H.-J. Kim, K. Brancik, D. Dickinson, H. Lee, A. Perrig, and B. Sinopoli, *Cyber–physical security of a smart grid infrastructure*, *Proceedings of the IEEE* **100** (2011), no. 1 195–209.

[23] D. Shishika, J. Paulos, and V. Kumar, *Cooperative team strategies for multi-player perimeter-defense games*, *IEEE Robotics and Automation Letters* **5** (2020), no. 2 2738–2745.

[24] R. Meira-Góes, E. Kang, R. H. Kwong, and S. Lafortune, *Synthesis of sensor deception attacks at the supervisory layer of cyber–physical systems*, *Automatica* **121** (2020) 109172.

[25] T.-Y. Zhang and D. Ye, *False data injection attacks with complete stealthiness in cyber–physical systems: A self-generated approach*, *Automatica* **120** (2020) 109117.

[26] G. Liang, S. R. Weller, F. Luo, J. Zhao, and Z. Y. Dong, *Distributed blockchain-based data protection framework for modern power systems against cyber attacks*, *IEEE Transactions on Smart Grid* **10** (2018), no. 3 3162–3173.

[27] R. Rajkumar, I. Lee, L. Sha, and J. Stankovic, *Cyber-physical systems: the next computing revolution*, in *Design automation conference*, pp. 731–736, IEEE, 2010.

[28] J. Xu, B. Liu, H. Mo, and D. Dong, *Bayesian adversarial multi-node bandit for optimal smart grid protection against cyber attacks*, *Automatica* **128** (2021) 109551.

[29] H. Rezaee, T. Parisini, and M. M. Polycarpou, *Resiliency in dynamic leader–follower multiagent systems*, *Automatica* **125** (2021) 109384.

[30] S. S. Kumkov, S. Le Ménec, and V. S. Patsko, *Zero-sum pursuit-evasion differential games with many objects: survey of publications*, *Dynamic games and applications* **7** (2017), no. 4 609–633.

[31] M. M. Tidball, O. Pourtallier, and E. Altman, *Approximations in dynamic zero-sum games ii*, *SIAM journal on control and optimization* **35** (1997), no. 6 2101–2117.

[32] O. Tsemogne, Y. Hayel, C. Kamhoua, and G. Deugoue, *A partially observable stochastic zero-sum game for a network epidemic control problem*, *Dynamic Games and Applications* **12** (2022), no. 1 82–109.

[33] H. Sandberg, S. Amin, and K. H. Johansson, *Cyberphysical security in networked control systems: An introduction to the issue*, *IEEE Control Systems Magazine* **35** (2015), no. 1 20–23.

[34] F. Miao, Q. Zhu, M. Pajic, and G. J. Pappas, *A hybrid stochastic game for secure control of cyber-physical systems*, *Automatica* **93** (2018) 55–63.

[35] C. Wu, X. Li, W. Pan, J. Liu, and L. Wu, *Zero-sum game-based optimal secure control under actuator attacks*, *IEEE Transactions on Automatic Control* **66** (2020), no. 8 3773–3780.

[36] E. Boix-Adserà, B. L. Edelman, and S. Jayanti, *The multiplayer colonel blotto game*, *Games and Economic Behavior* **129** (2021) 15–31.

[37] O. Gross and R. Wagner, *A continuous colonel blotto game*, tech. rep., RAND Project, Air Force, Santa Monica, CA, 1950.

[38] S. T. Macdonell and N. Mastronardi, *Waging simple wars: a complete characterization of two-battlefield blotto equilibria*, *Economic Theory* **58** (2015), no. 1 183–216.

[39] B. Roberson, *The colonel blotto game*, *Economic Theory* **29** (2006), no. 1 1–24.

[40] G. Schwartz, P. Loiseau, and S. S. Sastry, *The heterogeneous colonel blotto game*, in *2014 7th international conference on NETwork Games, COntrol and OPtimization (NetGCoop)*, IEEE, 2014.

[41] C. Thomas, *N-dimensional blotto game with heterogeneous battlefield values*, *Economic Theory* **65** (2018), no. 3 509–544.

[42] R. M. Bell and T. M. Cover, *Competitive optimality of logarithmic investment*, *Mathematics of Operations Research* **5** (1980), no. 2 161–166.

[43] R. B. Myerson, *Incentives to cultivate favored minorities under alternative electoral systems*, *American Political Science Review* **87** (1993), no. 4 856–869.

[44] Z. E. Fuchs and P. P. Khargonekar, *A sequential colonel blotto game with a sensor network*, in *2012 American Control Conference (ACC)*, pp. 1851–1857, IEEE, 2012.

[45] S. Guan, J. Wang, H. Yao, C. Jiang, Z. Han, and Y. Ren, *Colonel blotto games in network systems: Models, strategies, and applications*, *IEEE Transactions on Network Science and Engineering* **7** (2019), no. 2 637–649.

[46] E. M. Shahrivar and S. Sundaram, *Multi-layer network formation via a colonel blotto game*, in *2014 IEEE global conference on signal and information processing (GlobalSIP)*, pp. 838–841, IEEE, 2014.

[47] A. Ferdowsi, W. Saad, and N. B. Mandayam, *Colonel blotto game for sensor protection in interdependent critical infrastructure*, *IEEE Internet of Things Journal* **8** (2020), no. 4 2857–2874.

[48] A. Gupta, G. Schwartz, C. Langbort, S. S. Sastry, and T. Bařar, *A three-stage colonel blotto game with applications to cyberphysical security*, in *2014 American Control Conference*, pp. 3820–3825, IEEE, 2014.

[49] A. Gupta, T. Başar, and G. A. Schwartz, *A three-stage colonel blotto game: when to provide more information to an adversary*, in *International Conference on Decision and Game Theory for Security*, pp. 216–233, Springer, 2014.

[50] D. Kovenock and B. Roberson, *Coalitional colonel blotto games with application to the economics of alliances*, *Journal of Public Economic Theory* **14** (2012), no. 4 653–676.

[51] D. Q. Vu and P. Loiseau, *Colonel blotto games with favoritism: Competitions with pre-allocations and asymmetric effectiveness*, in *Proceedings of the 22nd ACM Conference on Economics and Computation*, pp. 862–863, 2021.

[52] S. Goyal and A. Vigier, *Attack, defence, and contagion in networks*, *The Review of Economic Studies* **81** (2014), no. 4 1518–1542.

[53] R. Powell, *Sequential, nonzero-sum "blotto": Allocating defensive resources prior to attack*, *Games and Economic Behavior* **67** (2009), no. 2 611–615.

[54] Y. Rinott, M. Scarsini, and Y. Yu, *A colonel blotto gladiator game*, *Mathematics of Operations Research* **37** (2012), no. 4 574–590.

[55] T. Adamo and A. Matros, *A blotto game with incomplete information*, *Economics Letters* **105** (2009), no. 1 100–102.

[56] D. Kovenock and B. Roberson, *A blotto game with multi-dimensional incomplete information*, *Economics Letters* **113** (2011), no. 3 273–275.

[57] J. Kim and B. Kim, *An asymmetric lottery blotto game with a possible budget surplus and incomplete information*, *Economics Letters* **152** (2017) 31–35.

[58] C. Ewerhart and D. Kovenock, *A class of n-player colonel blotto games with multidimensional private information*, *Operations Research Letters* **49** (2021), no. 3 418–425.

[59] R. Siegel, *Asymmetric all-pay auctions with interdependent valuations*, *Journal of Economic Theory* **153** (2014) 684–702.

[60] K. Paarporn, R. Chandan, M. Alizadeh, and J. R. Marden, *Asymmetric battlefield uncertainty in general lotto games*, *IEEE Control Systems Letters* **6** (2022) 2822–2827.

[61] L. Rentschler and T. L. Turocy, *Two-bidder all-pay auctions with interdependent valuations, including the highly competitive case*, *Journal of Economic Theory* **163** (2016) 435–466.

[62] M. R. Baye, D. Kovenock, and C. G. De Vries, *The all-pay auction with complete information*, *Economic Theory* **8** (1996), no. 2 291–305.

[63] Q. Zhu and T. Başar, *Game-theoretic methods for robustness, security, and resilience of cyberphysical control systems: games-in-games principle for optimal cross-layer resilient control systems*, *IEEE Control Systems Magazine* **35** (2015), no. 1 46–65.

[64] S. R. Etesami and T. Başar, *Dynamic games in cyber-physical security: An overview*, *Dynamic Games and Applications* **9** (2019), no. 4 884–913.

[65] W. Xing, X. Zhao, T. Basar, and W. Xia, *Security investment in cyber-physical systems: Stochastic games with asymmetric information and resource constrained players*, *IEEE Transactions on Automatic Control* (2021).

[66] J. Pita, M. Jain, J. Marecki, F. Ordóñez, C. Portway, M. Tambe, C. Western, P. Paruchuri, and S. Kraus, *Deployed armor protection: the application of a game theoretic model for security at the los angeles international airport*, in *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems: industrial track*, pp. 125–132, 2008.

[67] M. Tambe, *Security and game theory: algorithms, deployed systems, lessons learned.* Cambridge university press, 2011.

[68] R. Yang, B. J. Ford, M. Tambe, and A. Lemieux, *Adaptive resource allocation for wildlife protection against illegal poachers.*, in *Aamas*, pp. 453–460, 2014.

[69] D. Mattioli and C. Lombardo, *Amazon met with startups about investing, then launched competing products*, *Wall Street Journal* (2020).

[70] J. M. Snyder, *Election goals and the allocation of campaign resources*, *Econometrica: Journal of the Econometric Society* (1989) 637–660.

[71] R. Chandan, K. Paarporn, and J. R. Marden, *When showing your hand pays off: Announcing strategic intentions in colonel blotto games*, in *2020 American Control Conference (ACC)*, pp. 4632–4637, IEEE, 2020.

[72] V. Bilò and C. Vinci, *Dynamic taxes for polynomial congestion games*, *ACM Transactions on Economics and Computation (TEAC)* **7** (2019), no. 3 1–36.

[73] R. Cole, Y. Dodis, and T. Roughgarden, *How much can taxes help selfish routing?*, *Journal of Computer and System Sciences* **72** (2006), no. 3 444–467.

[74] A. Pigou, *The economics of welfare.* Macmillan, 1920.

[75] R. W. Rosenthal, *A class of games possessing pure-strategy nash equilibria*, *International Journal of Game Theory* **2** (1973), no. 1 65–67.

[76] J. R. Marden and A. Wierman, *Distributed welfare games*, *Operations Research* **61** (2013), no. 1 155–168.

[77] R. Cookson, C. Propper, M. Asaria, and R. Raine, *Socio-economic inequalities in health care in england*, *Fiscal Studies* **37** (2016), no. 3-4 371–403.

[78] P. D'Odorico, J. A. Carr, K. F. Davis, J. Dell'Angelo, and D. A. Seekell, *Food inequality, injustice, and rights*, *BioScience* **69** (2019), no. 3 180–190.

[79] G. Hardin, *The tragedy of the commons*, *Science* **162** (1968), no. 3859 1243–1248.

[80] N. Kaida and K. Kaida, *Facilitating pro-environmental behavior: The role of pessimism and anthropocentric environmental values*, *Social Indicators Research* **126** (2016), no. 3 1243–1260.

[81] T. Cabannes, M. A. S. Vincentelli, A. Sundt, H. Signargout, E. Porter, V. Fighiera, J. Ugirumurera, and A. M. Bayen, *The impact of gps-enabled shortest path routing on mobility: a game theoretic approach*, in *Transportation Research Board 97th Annual Meeting. Washington DC, USA*, pp. 7–11, 2018.

[82] J. Macfarlane, *When apps rule the road: The proliferation of navigation apps is causing traffic chaos. it's time to restore order*, *IEEE Spectrum* **56** (2019), no. 10 22–27.

[83] E. Koutsoupias and C. Papadimitriou, *Worst-case equilibria*, in *Annual Symposium on Theoretical Aspects of Computer Science*, pp. 404–413, Springer, 1999.

[84] B. Awerbuch, Y. Azar, and A. Epstein, *The price of routing unsplittable flow*, *SIAM Journal on Computing* **42** (2013), no. 1 160–177.

[85] G. Christodoulou and E. Koutsoupias, *The price of anarchy of finite congestion games*, in *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pp. 67–73, ACM, 2005.

[86] S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann, *Exact price of anarchy for polynomial congestion games*, *SIAM Journal on Computing* **40** (2011), no. 5 1211–1233.

[87] G. Christodoulou, E. Koutsoupias, and A. Nanavati, *Coordination mechanisms*, in *International Colloquium on Automata, Languages, and Programming*, pp. 345–357, Springer, 2004.

[88] I. Caragiannis, C. Kaklamanis, and P. Kanellopoulos, *Taxes for linear atomic congestion games*, *ACM Transactions on Algorithms (TALG)* **7** (2010), no. 1 1–31.

[89] D. Paccagnan and M. Gairing, *In congestion games, taxes achieve optimal approximation*, in *Proceedings of the 22nd ACM Conference on Economics and Computation*, pp. 743–744, ACM, 2021.

[90] D. Paccagnan, R. Chandan, B. L. Ferguson, and J. R. Marden, *Optimal taxes in atomic congestion games*, *ACM Transactions on Economics and Computation (TEAC)* **9** (2021), no. 3 1–33.

[91] A. Bjelde, M. Klimm, and D. Schmand, *Brief announcement: Approximation algorithms for unsplittable resource allocation problems with diseconomies of scale*, in *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures*, pp. 227–229, ACM, 2017.

[92] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden, *The price of stability for network design with fair cost allocation*, *SIAM Journal on Computing* **38** (2008), no. 4 1602–1623.

[93] A. S. Schulz and N. S. Moses, *On the performance of user equilibria in traffic networks*, in *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pp. 86–87, 2003.

[94] I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli, *Tight bounds for selfish and greedy load balancing*, *Algorithmica* **61** (2011), no. 3 606–637.

[95] G. Christodoulou and E. Koutsoupias, *On the price of anarchy and stability of correlated equilibria of linear congestion games*, in *European Symposium on Algorithms*, pp. 59–70, Springer, 2005.

[96] G. Christodoulou and M. Gairing, *Price of stability in polynomial congestion games*, *ACM Transactions on Economics and Computation (TEAC)* **4** (2015), no. 2 1–17.

[97] P. Kleer and G. Schäfer, *Tight inefficiency bounds for perception-parameterized affine congestion games*, *Theoretical Computer Science* **754** (2019) 65–87.

[98] V. Conitzer and T. Sandholm, *New complexity results about nash equilibria*, *Games and Economic Behavior* **63** (2008), no. 2 621–641.

[99] A. Fabrikant, C. Papadimitriou, and K. Talwar, *The complexity of pure nash equilibria*, in *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pp. 604–612, ACM, 2004.

[100] T. Roughgarden, *Intrinsic robustness of the price of anarchy*, *Journal of the ACM (JACM)* **62** (2015), no. 5 32.

[101] V. Gkatzelis, K. Kollias, and T. Roughgarden, *Optimal cost-sharing in general resource selection games*, *Operations Research* **64** (2016), no. 6 1230–1238.

[102] J. Kleinberg and S. Oren, *Mechanisms for (mis) allocating scientific credit*, in *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pp. 529–538, ACM, 2011.

[103] U. Nadav and T. Roughgarden, *The limits of smoothness: A primal-dual framework for price of anarchy bounds*, in *International Workshop on Internet and Network Economics*, pp. 319–326, Springer, 2010.

[104] V. Bilò, *A unifying tool for bounding the quality of non-cooperative solutions in weighted congestion games*, *Theory of Computing Systems* **62** (2018), no. 5 1288–1317.

[105] R. Chandan, D. Paccagnan, and J. R. Marden, "Matlab and Python packages to compute and optimize the price of anarchy." GitHub Repository: `https://github.com/rahul-chandan/resalloc-poa`, 2019.

[106] S. Suri, C. D. Tóth, and Y. Zhou, *Selfish load balancing and atomic congestion games*, *Algorithmica* **47** (2007), no. 1 79–96.

[107] M. Sviridenko, J. Vondrák, and J. Ward, *Optimal approximation for submodular and supermodular optimization with bounded curvature*, *Mathematics of Operations Research* **42** (2017), no. 4 1197–1218.

[108] U. Feige, *A threshold of ln n for approximating set cover*, *Journal of the ACM (JACM)* **45** (1998), no. 4 634–652.

[109] S. Barman, O. Fawzi, and P. Fermé, *Tight approximation guarantees for concave coverage problems*, in *38th International Symposium on Theoretical Aspects of Computer Science*, 2021.

[110] A. Vetta, *Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions*, in *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pp. 416–425, IEEE, 2002.

[111] M. Conforti and G. Cornuéjols, *Submodular set functions, matroids and the greedy algorithm: tight worst-case bounds and some generalizations of the rado-edmonds theorem*, *Discrete applied mathematics* **7** (1984), no. 3 251–274.

[112] M. Gairing, *Covering games: Approximation through non-cooperation*, in *International Workshop on Internet and Network Economics*, pp. 184–195, Springer, 2009.

[113] R. Chandan, D. Paccagnan, and J. R. Marden, *Optimal price of anarchy in cost-sharing games*, in *2019 American Control Conference (ACC)*, pp. 2277–2282, IEEE, 2019.

[114] G. Arslan, J. R. Marden, and J. S. Shamma, *Autonomous vehicle-target assignment: A game-theoretical formulation*, *Journal of Dynamic Systems, Measurement, and Control* **129** (2007), no. 5 584–596.

[115] D. Paccagnan, R. Chandan, and J. R. Marden, *Utility design for distributed resource allocation—part i: Characterizing and optimizing the exact price of anarchy*, *IEEE Transactions on Automatic Control* **65** (2019), no. 11 4616–4631.

[116] V. Chvatal, *A greedy heuristic for the set-covering problem*, *Mathematics of operations research* **4** (1979), no. 3 233–235.

[117] J. Mestre, *Greedy in approximation algorithms*, in *European Symposium on Algorithms*, pp. 528–539, Springer, 2006.

[118] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, *An analysis of approximations for maximizing submodular set functions—i*, *Mathematical programming* **14** (1978), no. 1 265–294.

[119] D. Paccagnan and J. R. Marden, *Utility design for distributed resource allocation–part ii: Applications to submodular, covering, and supermodular problems*, *IEEE Transactions on Automatic Control* **67** (2021), no. 2 618–632.

[120] É. Borel, *La théorie du jeu et les équations intégrales à noyau symétrique*, *Comptes Rendus de l'Académie* **173** (1921).

[121] A. Filos-Ratsikas, Y. Giannakopoulos, and P. Lazos, *The pareto frontier of inefficiency in mechanism design*, in *International Conference on Web and Internet Economics*, pp. 186–199, Springer, 2019.

[122] V. Ramaswamy, D. Paccagnan, and J. R. Marden, *Multiagent maximum coverage problems: The tradeoff between anarchy and stability*, *IEEE Transactions on Automatic Control* **67** (2021), no. 4 1698–1712.