

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Optimal Designs in Multi-Agent Systems and Industrial Refrigeration

Permalink

<https://escholarship.org/uc/item/38b5k4t3>

Author

Konda, Rohit

Publication Date

2024

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Optimal Designs in Multi-Agent Systems and Industrial Refrigeration

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

Doctor of Philosophy
in
Electrical and Computer Engineering

by

Rohit Konda

Committee in charge:

Professor Jason R. Marden, Chair
Professor Mahnoosh Alizadeh
Professor João P. Hespanha
Professor James B. Rawlings

September 2024

The Dissertation of Rohit Konda is approved.

Professor Mahnoosh Alizadeh

Professor João P. Hespanha

Professor James B. Rawlings

Professor Jason R. Marden, Committee Chair

June 2024

Optimal Designs in Multi-Agent Systems and Industrial Refrigeration

Copyright © 2024

by

Rohit Konda

To my family and friends.

Acknowledgements

First and foremost, I would like to give thanks to my advisor, Professor Jason Marden. In totality, my graduate studies would not be possible without the tremendous support and guidance that I have received from Jason. Throughout, I have enjoyed our technical discussions and I have matured as a presenter and a writer directly from Jason's advice. As an advisor, Jason has been nothing but approachable, friendly, and supportive - for that I am forever grateful. Apart from my advisor I'd like to also thank the members of my thesis committee: Professor Mahnoosh Alizadeh, Professor Joao Hespanha, Professor James B. Rawlings, and Jesse Crossno. Your valuable feedback and advice has been of tremendous help. Additionally, I would like to extend my gratitude to Jesse and the CrossnoKaye team for an excellent internship and collaborative experience, from which I have learned an immense amount.

I'm also very much grateful for the collaborative atmosphere that was present during my period at Harold Frank Hall. First, I'd like to thoroughly thank the lab members that have come before me: Keith Paarporn, David Grimsman, Rahul Chandan, Adel Aghajan, and Bryce Ferguson. Your collective mentorship and friendship is very much appreciated. I am also thankful that I can pass on the torch to the lab members that have come after me: Gilberto Diaz-Garcia, Vade Shah, Maxwell Crisafulli, and Jordan Prescott. It's been extremely rewarding to be able to work and grow alongside my peers during this process. My heartfelt gratitude also extends to not only the other wonderful student members at CCDC, but also to the peers that joined alongside our Whiteboard seminars.

Of course, the lows of my PhD would have been much drearier without the support and good times had with the many friends I have had the pleasure of meeting along the way in Santa Barbara. I am super grateful to have met so many wonderful people along the way. Additionally, I'd like to especially thank my friends from undergraduate

and high school for keeping me sane through lockdown. My gratitude goes to the many people that I have had the pleasure of meeting in the past five years.

Finally, I would like to express my heartfelt gratitude to my family, whose support and encouragement have been instrumental for my studies. My parents, Vijay and Vindhya, have always been my guiding lights, offering invaluable advice and constant motivation. Their dedication and sacrifices have provided me with the foundation I needed to succeed. Additionally, my brother Roshan has been a tremendous source of help and support. I am deeply appreciative of their collective efforts and could not have reached this milestone without their love and support.

Curriculum Vitæ

Rohit Konda

Education

2024	Ph.D. in Electrical and Computer Engineering, University of California, Santa Barbara.
2019	M.S. in Electrical and Computer Engineering, Georgia Institute of Technology
2018	B.S. in Biomedical Engineering, Georgia Institute of Technology

Publications

Journal Publications

- Konda, Rohit, et al. "Optimal Utility Design of Greedy Algorithms in Resource Allocation Games." *IEEE Transactions on Automatic Control* (2024).
- Squires, Eric, et al. "Composition of safety constraints for fixed-wing collision avoidance amidst limited communications." *Journal of Guidance, Control, and Dynamics* 45.4 (2022): 714-725.
- Konda, Rohit, Aaron D. Ames, and Samuel Coogan. "Characterizing safety: Minimal control barrier functions from scalar comparison systems." *IEEE Control Systems Letters* 5.2 (2020): 523-528.

Conference Publications

- Konda, Rohit, et al. "Utilizing Load Shifting for Optimal Compressor Sequencing in Industrial Refrigeration." 2024 American Control Conference (ACC). IEEE, 2024.
- Konda, Rohit, Rahul Chandan, and Jason R. Marden. "Quality of Non-Convergent Best Response Processes in Multi-Agent Systems through Sink Equilibria." 2023 62nd IEEE Conference on Decision and Control (CDC). IEEE, 2023.
- Zhang, Runyu, et al. "Markov games with decoupled dynamics: Price of anarchy and sample complexity." 2023 62nd IEEE Conference on Decision and Control (CDC). IEEE, 2023.
- Konda, Rohit, et al. "Balancing asymptotic and transient efficiency guarantees in set covering games." 2022 American Control Conference (ACC). IEEE, 2022.
- Konda, Rohit, David Grimsman, and Jason R. Marden. "Execution order matters in greedy algorithms with limited information." 2022 American Control Conference (ACC). IEEE, 2022.
- Konda, Rohit, Rahul Chandan, and Jason R. Marden. "Mission Level Uncertainty in Multi-Agent Resource Allocation." 2021 60th IEEE Conference on Decision and Control (CDC). IEEE, 2021.

- Squires, Eric, et al. "Safety with limited range sensing constraints for fixed wing aircraft." 2021 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2021.
- Squires, Eric, et al. "Model free barrier functions via implicit evading maneuvers." arXiv preprint arXiv:2107.12871 (2021).
- Konda, Rohit, et al. "Provably-safe autonomous navigation of traffic circles." 2019 IEEE Conference on control technology and applications (CCTA). IEEE, 2019.

Abstract

Optimal Designs in Multi-Agent Systems and Industrial Refrigeration

by

Rohit Konda

The focus of this thesis is on developing control strategies for large-scale systems. We look at two distinct problem areas: the design of coordination algorithms for multi-agent systems and the optimization of industrial refrigeration systems.

In the first part of this thesis, we focus on multi-agent systems, where various decision-makers interact with each other, each with their own local objectives. To understand and design control strategies within these systems, we employ a game-theoretic viewpoint. Under this framework, the utility functions and strategic interactions are explicitly modeled to understand emergent behavior. Optimal incentive mechanisms can then be derived to align joint outcomes with the collective objective. The subsequent chapters explore different scenarios within this framework, such as collective transient behavior, coordination in limited information settings, and run-time analysis. Overall, we aim to classify multi-agent behavior and optimize outcomes under varying conditions.

In the second part of this thesis, we focus on designing control strategies for industrial refrigeration settings. Industrial refrigeration plays a significant role in various sectors and represents a major portion of total global energy usage. Despite this, there are significant control opportunities in increasing the efficiency of such systems by carefully modulating system variables, such as pressure and temperature, to operate close to optimal thermodynamic conditions. We explore the control opportunities of compressor sequencing and scheduling as a viable option to significantly reduce energy usage, by utilizing techniques from inventory control and scheduling theory. By leveraging optimization

methodologies, the manuscript seeks to enhance the efficiency of refrigeration systems, thereby reducing energy usage and environmental impact.

Contents

Curriculum Vitae	vii
Abstract	ix
1 Introduction	1
1.1 Introduction to Utility Designs in Multi-Agent Systems	2
1.2 Introduction to Industrial Refrigeration	4
Part I Utility Design in Multi-Agent Systems	6
2 Game Dynamics	7
2.1 Strategic Form Games	8
2.2 Learning Dynamics and Decision Algorithms	8
3 Guarantees in k - round walks	12
3.1 Introduction	13
3.2 Model and Preliminaries	15
3.3 One Round Walks	20
3.4 Finite Walks	46
3.5 Asymptotic Walks	51
3.6 Tradeoffs	57
3.7 Simulations	68
4 Greedy Algorithms in Limited Information Settings	72
4.1 Introduction	73
4.2 Model and Preliminaries	75
4.3 Main Theoretical Results	78
4.4 Applications in Submodular Settings	90

5	Nash equilibrium in Uncertain Settings	96
5.1	Introduction	97
5.2	Model and Preliminaries	99
5.3	Characterization of PoA	104
5.4	Set Covering Results	111
6	Nonconvergent Learning Dynamics	121
6.1	Introduction	122
6.2	Model and Preliminaries	124
6.3	Main Results	129
6.4	Better Response Sink Equilibrium	136
6.5	Appendix	138
 Part II Industrial Refrigeration		 142
7	Load Shifting in Compressor Sequencing	143
7.1	Preliminaries on Industrial Refrigeration	144
7.2	Mathematical Model	149
7.3	Results on Compressor Sequencing	151
7.4	Simulations	159
8	Inventory Control and Peak Pricing	162
8.1	Introduction to Peak Pricing	163
8.2	Mathematical Model	166
8.3	Results on Compressor Scheduling	170
8.4	Simulations	179
 Bibliography		 181

Chapter 1

Introduction

Large-scale systems are prevalent in a diverse set of domains, spanning from industrial processes [1] to transportation networks [2]. These systems typically involve numerous interconnected components that interact dynamically with each other. Managing such systems usually requires sophisticated control strategies that are capable of regulating whole-scale behavior in an effective way. In this manuscript, we use frameworks from control theory, optimization, and game theory to develop theoretical insights for efficient control strategies for large-scale systems.

In this manuscript, effective control strategies are dictated by solving an underlying optimization problem. However, in many instances, solving the optimization problem directly can be infeasible, either due to informational constraints, computational complexity, or other concerns. Thus, throughout this work, we will focus on developing online and approximate solutions that yield well performing behavior, while reducing the computational and informational burden of the underlying control algorithms. We do this in two contexts: First, we will examine control designs in general multi-agent systems, where there are different decision makers that have to interact; Second we will examine control designs in industrial processes, where we develop insights on effective control strategies

that minimize energy usage. We address both of these research directions separately in this manuscript.

1.1 Introduction to Utility Designs in Multi-Agent Systems

In the first part of the manuscript, we consider developing control strategies in multi-agent systems, where multiple autonomous decision makers interact, each with their own individual objectives. In these collective systems, agents may have different constraints, information, or preferences leading to complex interactions that traditional control methodologies may struggle to address. Thus, we look towards game theory as a powerful framework in which to address control designs in these multi-agent systems. Under a game-theoretic framework, we can formally model the strategic interactions among agents, considering their decisions as part of a game, where each agent aims to maximize their utility. By modeling the interactions and preference structure among agents, game theory allows for control designs that account for the strategic behavior, leading to distributed and efficient coordination algorithms.

Throughout our work, we shape agent's behavior through utility design (also known as mechanism design), where agent's preferences via their individual utility function is modified to promote better emergent behavior. Each agent's utility function quantifies its preferences over their decision set; under a rationality assumption, each agent employs a decision making process that maximizes its own utility. However, by carefully designing utility functions, a system designer can incentive cooperation and align individual goals with a collective global objective. From coordination in traffic networks [3] to managing distributed energy grids [4], game-theoretic approaches holds much promise as an methodology

to address increasingly complex and interconnected systems.

Under this game-theoretic viewpoint, the focus of this study is on classifying behavior under selfish processes and implementing utility design to incentivize better global behavior. The main contributions of this part are dictated as follows:

- In Chapter 3, we focus on the quality of trajectories that result from selfish decision making. We examine resource allocation scenarios and examine the performance after agents are allowed to respond k times. We provide characterizations for agent performance when $k = 1$, when k is finite, and when $k \rightarrow \infty$ asymptotically. We employ utility design to optimize the performance of the resulting game outcomes.
- In Chapter 4, we allow agents to selfishly respond in a sequential manner. In this chapter, we examine when agents have limited information and must communicate their decisions along a communication network. Along a given communication network, we examine the resulting communication time guarantees to run this sequential greedy process.
- In Chapter 5, we study scenarios where agents don't have full access to a world state. We examine the performance guarantees and optimal utility designs when agents must selfishly respond to each other in limited information scenarios.
- In Chapter 6, we extend to general scenarios where agent objectives may be misaligned for operational reasons. We then characterize the performance given a utility structure when agents may be misaligned in their preference structure and must behave rationally in response to other agent's decisions.

1.2 Introduction to Industrial Refrigeration

Industrial refrigeration is a critical component in various sectors, ensuring that perishable goods remain fresh; some examples include food and beverage, pharmaceuticals, and industrial chemicals [5–7]. By maintaining specific temperature conditions, refrigeration systems prevent spoilage and extending the shelf life of products, making it an integral part of the global supply chain. As such, industrial refrigeration represents a significant portion of commercial energy usage. However, managing refrigeration systems efficiently poses a multifaceted challenge due to the intricate interplay between different system states, such as temperature, humidity, and product specifications. By employing efficient control algorithms and data-informed methodologies, refrigeration operators can optimize the performance, potentially resulting in significant reductions in energy use and environmental impact. As industries increasingly prioritize sustainability and resource conservation, the integration of optimization methodologies into refrigeration systems represents a pivotal step toward achieving these goals.

In this manuscript, we identify novel control opportunities to increase the efficiency of refrigeration systems. We base our findings on data collected from a Butterball LLC ® located in Huntsville, AZ. We use tools from optimization and scheduling theory to characterize optimal policies in these refrigeration systems. The main contributions of this part are dictated as follows:

- In Chapter 7, we examine the problem of compressor sequencing, where the capacities of multiple compressors must be scheduled in such a way to minimize energy use while meeting the required thermal demands of the system. We characterize the optimal sequencing algorithms and introduce load shifting as a novel mechanism for aiding compressor sequencing. While load shifting has been traditionally used to respond to variable cost structures, we introduce load shifting for compressor

sequencing to allow for more efficient compressor configurations. We show the benefits of this ideology through our theoretical results and simulation models.

- In Chapter 8, we discuss optimal scheduling that addresses the necessary thermal loads while reducing operational expenses. We use methodologies from inventory management to characterize the structure of the optimal control strategies. We analyze the impact of various energy cost-rate frameworks, including time-of-use (TOU) pricing and peak pricing schemes, on the control policies.

Part I

Utility Design in Multi-Agent Systems

Chapter 2

Game Dynamics

In the advent of increasingly distributed architectures and multi-agent systems, there has been significant effort from the theoretical community to understand and design such systems. In this thesis, game theory is used as the fundamental methodology that we use to design such systems. Game theory has solidified its usefulness as a modeling strategy in economic contexts [8], computational contexts [9], and more recently, engineering contexts [10]. In essence, game theory can be encapsulated by the study of players, their decisions, and their preferences over their decisions.

However, much of the fundamental work and exploration in game theory has been on the study of fixed points and equilibrium among multi-agent decisions - this viewpoint has bled into the current framework in how game theory is mainly used. However, we take a different viewpoint in this thesis, where we explicitly model the interaction process between agents and study features of the decisions along that process. We introduce the necessary preliminaries in this section.

2.1 Strategic Form Games

In this section, we review classical models of game theory. Formally, a *strategic form* or normal form game consists a set of n players. The set of players is denoted $\mathcal{I} \doteq \{1, \dots, n\}$. Each player has a set of admissible decisions \mathcal{A}_i , in which they can select a specific decision $a_i \in \mathcal{A}_i$ from. The decision set \mathcal{A}_i can be quite abstract, but is usually assumed to be finite throughout this manuscript. Each player i can select an action $a_i \in \mathcal{A}_i$, resulting in a joint action $a \doteq (a_1, \dots, a_n)$. The space of joint actions is denoted as $\mathcal{A} \doteq \prod_i \mathcal{A}_i$.

Under the joint decision space, each player has a preference over the joint decision space \mathcal{A} . In generality, this is equivalent to \mathcal{A} being an ordered space with some ordering \preceq_i for each player. Usually, we will assume that the preference structure can be generated through a given *utility function* $U_i(a) : \mathcal{A} \rightarrow \mathbb{R}$. In this case decision a_1 is preferred over a_2 for agent i if $a_2 \preceq_i a_1$ or equivalently if $U_i(a_2) \leq U_i(a_1)$, in the usual sense in \mathbb{R} . One crucial aspect of this model is that the utility functions $U_i \neq U_j$ for $i \neq j$ may disagree in general - different agents can have different preferences over the joint decisions. In this way, game outcomes are not immediately apparent and there are different ways to model agent interactions along this model. In this manuscript, we take a dynamical viewpoint that will be discussed in the next section.

2.2 Learning Dynamics and Decision Algorithms

We take a dynamical view point to model interactions between players. Specifically, we assume that each agent is endowed an *update rule* which produces the following non-autonomous discrete process:

$$a_i^{t+1} = f_i(t, a^t; \{U_i\}_{i \in \mathcal{I}}). \quad (2.1)$$

In this way, the utility function U_i serves to parameterize the update rule f_i . These dynamics can also be extended to probabilistic formulations. We outline and discuss some classical examples of decision making dynamics below.

2.2.1 Best Response Dynamics

We first outline a fundamental learning dynamics - *best response dynamics*, that goes far back as Nash's seminal paper [11]. Informally, best response dynamics involve each player adjusting their decisions to maximize their individual utility in a continuous manner based on the current strategies chosen by the other players. The *best response* is the decision that maximizes the utility of the player, i.e. a^{br} is a best response action for agent i , if $a_i^{\text{br}} \in \arg \max_{a_i} U_i(a_i, a_{-i})$, where $a_{-i} \doteq (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$ is the joint action without the action of agent i . Note that player i must have access to the current strategy selection a_{-i} of the other agents for it to be able to compute its best response action. Before defining the dynamics as in Eq. (2.1), we first outline the *selection rule* or $s : \mathbb{N} \rightarrow 2^{\mathcal{I}}$, that defines which players are performing a best response. For example, if the best response is done simultaneously, $s(t) = \mathcal{I}$, but if done in a round robin fashion, $s(t) = t \bmod n$. Additionally, the selection rule can be defined in a probabilistic sense as well. With this, the learning dynamics can be formally described as

$$a_i^{t+1} = \begin{cases} a_i^{\text{br}} \equiv \arg \max_{a_i} U_i(a_i, a_{-i}^t) & \text{if } i \in s(t), \\ a_i^t & \text{otherwise .} \end{cases} \quad (2.2)$$

In general, we can assume that the best response set is unique, but if not, a decision can be arbitrarily selected from it. One key aspect of best response dynamics is their simplicity and intuitive appeal. Players update their strategies based on what they perceive to be the most advantageous response to the strategies of others, reflecting

a natural decision-making process observed in many real-world scenarios. This makes best response dynamics a valuable modeling tool, offering valuable insights into strategic behavior and collective outcomes in games.

2.2.2 Greedy Algorithms

We provide some connections to best response dynamics and another classical local decision making algorithm - the *greedy algorithm*. Greedy algorithms are defined by a sequential process, where agents make a series of locally optimal choices, with the hope that these local optimizations will lead to a globally optimal solution. The primary characteristic of a greedy algorithm is its short-sightedness, as agents make decisions solely on the information available at the current step. Greedy algorithms have been shown to effective algorithmic solutions for broad classes of problems including graph-theoretic, local search, and submodular settings [12]. Formally, the greedy algorithm can be defined by the dynamics in Eq. (2.2) with the selection rule $s(t) = t$ if $t \leq n$ and \emptyset if $t \geq n$. Thus we can recover classical decision making algorithms from a game-theoretic viewpoint.

2.2.3 Better Response Dynamics

We can also relax best response dynamics to consider better response dynamics, where the decision-making and rationality assumptions of agents can be relaxed. While in a best response process, agents maximize their utility, in a better response process, agents are only expected to deviate to a decision that increases their utility over their current decision. Formally, agents are able deviate to a decision in their *better response set* defined by $\{a_i \in \mathcal{A}_i : U_i(a_i, a^{t-i}) \geq U_i(a^t)\}$. Better response processes may be a more realistic decision-making model, where computing best responses may be intractable, but

it is possible for agents to make iterative improvements in their utility. This iterative adjustment process can be seen in markets, where firms continuously adjust prices and outputs in response to competitors [13], or in network settings, where users adjust their usage patterns based on congestion and other users' behaviors [14].

2.2.4 Other Learning Dynamics

There are whole host of learning dynamics that are possible to implement. For example, adaptive play is an extension of the best response process to include memory [15]. It is possible to inject noise into the learning process, as done in a noisy best response process or log-linear learning [16]. There are also iterative update algorithms, such as no-regret learning, or multiple-weights algorithm [17]. Overall, the decision-making of the agents can be modeled in a variety of ways, but intrinsically employ some component of selfish decision-making.

2.2.5 Nash Equilibrium as Fixed Points

The Nash equilibrium in game theory is a fundamental concept where each player's strategy is the best response to the strategies of others, resulting in no player having an incentive to deviate unilaterally. This equilibrium can be understood as a fixed point of best response dynamics (and other learning dynamics as well), where the strategy profile is stable because every player's strategy is optimal given the others' strategies. Formally, a joint decision a^{ne} is considered a Nash equilibrium if $a_i^{\text{ne}} \in \arg \max_{a_i} U_i(a_i, a_{-i}^{\text{ne}})$ for all players $i \in \mathcal{I}$. In certain games structures, convergence to Nash equilibrium can be guaranteed [18]. Thus, Nash equilibrium are classically taken as a model of emergent outcomes in multi-agent systems. However, in this thesis, we take more of a dynamical viewpoint, and move beyond equilibrium analysis.

Chapter 3

Guarantees in k - round walks

Developing competitive and distributed solutions for multi-agent problems is essential for many emerging application domains. Game theory has proven to be a valuable framework for designing these algorithms. However, the majority of research within this framework focuses on equilibrium behavior often neglecting transient behavior. In this chapter, we address this gap by examining the transient efficiency guarantees of *best response processes* in the context of resource allocation games, which have numerous applications in various application domains. Our primary focus in this chapter is on characterizing the optimal short-term system behavior along the best-response process. Remarkably, we find that the transient performance guarantees are relatively close to optimal long-term/asymptotic performance guarantees. Moreover, we explore the trade-offs involved in optimizing for both asymptotic and transient efficiency through various utility designs. Our analysis provides a comprehensive understanding of the span of utility designs and their joint effect on transient and asymptotic guarantees. The results and discussion in this chapter is based on the work presented in [19,20].

3.1 Introduction

Multi-agent architectures have recently gained considerable attention due to their widespread applications. The underlying goal of these distributed systems is to coordinate agents to desirable system states, as measured by some global objective, through local decision making processes. As such, *game theoretic methods* have emerged as an important design methodology in these settings. In this approach, each agent is treated as a player in a non-cooperative game with their decision-making governed by a local-objective or *utility* function.

Endowed with a utility function, each agent can update its decisions in a self-interested manner to dynamically respond to environmental changes, including those induced by other agents. In this chapter, we consider when agents use the classical update rule - the *round-robin best response* algorithm. Under this process, the decision updates are done in a sequential manner: at each iteration, a single, chosen agent optimizes its decision against its utility function with all of the other agents keeping their decisions fixed. As such, the utility structure of the agents can significantly influence the underlying dynamical behavior of the agents: this is highlighted in Figure 3.1. Understanding this relationship is important in classifying the emergent behavior, especially when a system operator would like to design or tune the utility functions in an optimal fashion.




Figure 3.1: If a given multi-agent scenario with n agents is modeled as a game, the construction of distributed algorithms can be decoupled into two domains: the design of local objectives (utilities) and the design of the learning dynamics. In this chapter, we fix the dynamics to the classical k round-robin best response and study the effects of the utility design on the efficiency bounds for the resulting decision trajectory. Moreover, we characterize the guarantees as the number of rounds k increases.

Therefore, in this chapter, we shift focus to the transient behavior of the round-robin best response algorithm. Specifically, we benchmark the iterative process to be the round-robin best response algorithm and study the performance guarantees that result from various designs of utility functions in the context of the well-studied class of *resource allocation games*.

While the many of the existing theoretical results on the Nash equilibrium behavior are positive, these performance guarantees only emerge asymptotically. In fact, arriving at Nash equilibrium may even take an exponential time [21], rendering the resulting characterizations irrelevant in many realistic multi-agent scenarios. For example, there may be an extremely large number of agents in the multi-agent scenario 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, expecting that the agents will converge to Nash equilibrium may not be a reasonable assumption.

Moreover, the work in this chapter belongs to a larger research trend that aims to study game theoretic models beyond their respective equilibrium. In contrast to the traditional game-theoretic approach, the *game dynamics* are embraced as a valuable feature of the game, where a rigorous study of actualized play can provide important insights about the game model (see, for e.g., [22–24]). Furthermore, characterizing performance guarantees along these game dynamics is valuable to understanding the transient behavior of the agents. In contrast to the study of equilibrium quality, the literature on transient guarantees is much less developed. However, we highlight an important subset of works that characterize transient performance guarantees in different game-theoretic contexts: such problem domains include affine congestion games [25–28], market sharing games [25,29], basic utility games [29], series-parallel networks, and load-balancing games [30,31].

3.2 Model and Preliminaries

We consider distributed settings that are modeled as *resource allocation games* [32]. Let $\mathcal{I} = \{1, \dots, n\}$ be the collection of agents that can possibly utilize a portion of a given finite set of resources $\mathcal{R} = \{r_1, \dots, r_d\}$. Each resource $r \in \mathcal{R}$ is associated with a *welfare rule* $w_r : \mathbb{N} \rightarrow \mathbb{R}_{>0}$ that defines the welfare accrued at each resource based on number of agents that utilize it. We use the denotation $w_r(0) = 0$. The choice of resource utilization for each agent is given by its action set $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$. As such, the quality of a joint action $a = (a_1, \dots, a_n) \in \mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ is classified through a system-level objective function $W : \mathcal{A} \rightarrow \mathbb{R}_{>0}$ of the form

$$W(a) = \sum_{r \in \mathcal{R}} w_r(|a|_r), \quad (3.1)$$

where $|a|_r = |\{i \in \mathcal{I} : r \in a_i\}|$ is the number of agents that utilize r in the joint action a . These types of objectives are commonplace in many engineering domains, with applications in information gathering, image segmentation, statistical summarization (see [33]). We include a sample of these applications below.

Example 1 (Weapon-Target Assignment). Consider a weapon-target assignment problem [34], in which a set of agents \mathcal{I} defend against a set of \mathcal{R} targets. Each agent $i \in \mathcal{I}$ has to decide which targets $r \in \mathcal{R}$ to defend against with a p_d chance of defending against each target, where its decision $a_i \subset \mathcal{R}$ is a subset of targets. Each target is also characterized by its relative value $v_r \geq 0$. As a whole, the set of agents would like to maximize the expected value of the targets defended. As such, the nonlinear objective of the agent's decisions is denoted by

$$W(a) = \sum_{r \in \mathcal{R}} v_r \cdot (1 - (1 - p_d)^{|a|_r}). \quad (3.2)$$

Example 1A (Set Covering). If the probability of defending p_d is 1, then we recover the problem of set covering [35]. In this domain, there is no benefit for more than agent to defend against a target, and as a whole, the agents would like to maximize the value of the targets defended. Thus, the welfare function simplifies to

$$W(a) = \sum_{r \in \bigcup_i a_i} v_r. \quad (3.3)$$

Example 2 (Wireless Transmission over a Network). Consider a group of communicating agents, as in [36], that send transmissions through a shared network with nodes V and edges E . Each agent would like to send a wireless transmission over the network from a given start node s_i to an end node t_i , and it must choose one out of the possible paths from s_i to t_i to transmit across. Thus, the resource set is E and the action set $\mathcal{A}_i \subset 2^E$ is the possible set of paths from s_i to t_i . An edge e may experience congestion if multiple transmissions utilize it; we assume that with each additional transmission, the rate of transmission experiences a harmonic-like decay. As a whole, the agents would like to send their transmissions with the highest rate. As such, the system welfare is

$$W(a) = \sum_{e \in \bigcup_i a_i} \sum_{j=1}^{|a|_e} \frac{1}{j}. \quad (3.4)$$

Example 3 (k -Clustering). Consider a classical dimensionality reduction problem of distilling a given data set into representative clusters, similar to [37]. Each data point d_i has a set of possible representative clusters $\mathcal{A}_i \subset \mathcal{R}$ that it can join. The objective is to compute clusterings with maximum overlap between data points. Greedy algorithms

can provide quick solutions with respect to the following welfare function.

$$W(a) = \sum_{r \in \bigcup_i a_i} (|a|_r)^2. \quad (3.5)$$

The global directive of the agents is to coordinate to a joint action that maximizes the system welfare, i.e. $a^{\text{opt}} \in \arg \max_{a \in \mathcal{A}} W(a)$. In order to coordinate the agents in a distributed fashion, we assume a game-theoretic setup, where each agent optimizes its decision with respect to a given local objective, or utility function $U_i : \mathcal{A} \rightarrow \mathbb{R}$ in a self-interested process. To establish the learning procedure of the agents, we focus attention to a class of best response processes known as *k-round walks* (or *k* round-robin best response), explicitly stated in Algorithm 1. At each step of the algorithm, an agent is selected in a round-robin fashion to perform a best response; this goes on until k rounds have been completed. For a given joint action $\alpha \in \mathcal{A}$, we say the action a_i^{br} is a *best response* for agent i if

$$a_i^{\text{br}} \in \text{Br}_i(\alpha_{-i}) = \arg \max_{a_i \in \mathcal{A}_i} U_i(a_i, \alpha_{-i}), \quad (3.6)$$

where $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$ denotes the joint action a without the action of agent i . We also assume that the round-robin walk begins with none of resources being utilized by any of the agents, denoted by the null joint action $a^\emptyset := \emptyset$.

Algorithm 1 *k*-Round Walk

Require: $a(0) \leftarrow a^\emptyset$, $i \leftarrow 1$, $\tau \leftarrow 1$, k , n

while $\tau \leq kn$ **do**

Modify action of agent i to $a_i(\tau) \leftarrow \text{Br}_i(a_{-i}(\tau - 1))$;

Fix other agent actions to $a_{-i}(\tau) \leftarrow a_{-i}(\tau - 1)$;

Increment τ by 1;

Set the next agent $i \leftarrow \tau \bmod n$ by round-robin;

end while

return $a(kn)$

Running the k -round walk algorithm induces the action trajectory $a(0) = a^\emptyset, a(1), \dots, a(kn - 1), a(kn)$ with $a(kn)$ being the end resulting joint action. The central goal of this chapter is to understand how the utility functions of the agents affect the behavior of the k -round walk. To this end, the performance of the joint action $a^k \doteq a(kn)$ after k rounds can be characterized through the following metric ¹

$$\text{Eff}(G; k) = \frac{W(a^k)}{\max_{a \in \mathcal{A}} W(a)} \in [0, 1], \quad (3.7)$$

where a ratio closer to 1 implies the efficiency of the joint action after k rounds is closer to optimal. We use tuple $G \triangleq (\mathcal{I}, \mathcal{A}, \mathcal{R}, W, \{U_i\}_{i \in \mathcal{I}})$ to define the game instance under consideration. Furthermore, we characterize the performance of the limit points ² of the k -round walk as

$$\text{Eff}(G; \infty) = \frac{W(\lim_{k \rightarrow \infty} a^k)}{\max_{a \in \mathcal{A}} W(a)}. \quad (3.8)$$

Accordingly, the form of the utility functions has a significant impact on the given metric $\text{Eff}(G; k)$. One natural utility function to consider is the *common-interest* (CI) utility, where all the agents share the same utility function $U_i(a) \equiv W(a)$ for all $a \in \mathcal{A}$ and $i \in \mathcal{I}$. In line with literature on submodular optimization [38], we demonstrate that common interest utilities exhibit constant factor efficiency guarantees $\text{Eff}(G; k)$. However, as unveiled by subsequent results in this chapter, by finely adjusting the utility functions, we can achieve superior efficiency guarantees. Following this perspective, let's consider utility functions of the form

$$U_i(a_i, a_{-i}) = \sum_{r \in a_i} u_r(|a|_r), \quad (3.9)$$

¹We note that the resulting action $a(kn)$ may not be unique if the best response $\text{Br}_i(a_{-i}(\tau))$ is not unique for some τ . In this case, we overload $W(a^k)$ to mean the minimum welfare $\min_{\alpha \in a^k} W(\alpha)$ for the k -round walk.

²Since resource allocation games are *potential games* [18], the limit points of k -round walks are necessarily a subset of the set of Nash equilibrium.

where the *utility rule* $u_r : \mathbb{N} \rightarrow \mathbb{R}_{>0}$ defines the resource-specific agent utility determined by $|a|_r$. We assume that $u_r(1) = w_r(1)$ for all $r \in \mathcal{R}$. Now, the choice of the utility rules u_r influences the resulting joint action trajectory, and in turn the resulting end joint action a^k .

In most scenarios of interest, a system designer is required to specify the utility rules u_r without specific knowledge of the resource allocation game parameters, such as the number of agent \mathcal{I} or the action set \mathcal{A} . To that end, let \mathcal{W} be the set of possible welfare rules that could be associated with any resource, i.e., $w_r \in \mathcal{W}$ for all $r \in \mathcal{R}$. Here, the system designer is tasked with associating a utility rule to each type of resource, i.e., the utility rule for any resource $r \in \mathcal{R}$ with the welfare rule w_r is of the form $u_r = \mathcal{U}(w_r)$ where we refer to the map $\mathcal{U} : \mathcal{W} \rightarrow \mathbb{R}_{\geq 0}^{\mathbb{N}}$ as the *utility design*. We refer to the common interest utility design as \mathcal{U}_{CI} . Lastly, we define the set of resource allocation games that are induced by \mathcal{W} and \mathcal{U} as $\mathcal{G}_{\mathcal{W},\mathcal{U}}$, where a game $G \in \mathcal{G}_{\mathcal{W},\mathcal{U}}$ if $w_r \in \mathcal{W}$ and $u_r = \mathcal{U}(w_r)$ for all resources $r \in \mathcal{R}$.

The central focus of this chapter is to understand how the choice of utility rules, derived from $\mathcal{U}(\cdot)$, and the number of rounds k impacts the efficacy of the emergent collective behavior in the k -round walk. Therefore, we additionally extend the efficiency measure to the set of games $\mathcal{G}_{\mathcal{W},\mathcal{U}}$ and quantify the optimal efficiency guarantees as

$$\text{Eff}(\mathcal{W}, \mathcal{U}; k) = \inf_{G \in \mathcal{G}_{\mathcal{W},\mathcal{U}}} \text{Eff}(G; k), \quad (3.10)$$

$$\text{Eff}^*(\mathcal{W}; k) = \sup_{\mathcal{U} : \mathcal{W} \rightarrow \mathbb{R}_{>0}^{\mathbb{N}}} \text{Eff}(\mathcal{W}, \mathcal{U}; k). \quad (3.11)$$

We similarly extend the definitions of $\text{Eff}(G; \infty)$. We importantly highlight that $\lim_{k \rightarrow \infty} \text{Eff}(\mathcal{W}, \mathcal{U}; k) \neq \text{Eff}(\mathcal{W}, \mathcal{U}; \infty)$ in general, as the limit $\lim_{k \rightarrow \infty}$ and the infimum $\inf_{G \in \mathcal{G}_{\mathcal{W},\mathcal{U}}}$ cannot be interchanged. We devote the next section to characterizing the introduced efficiency metrics.

Additional Notation. Given a set \mathcal{S} , $|\mathcal{S}|$ represents its cardinality. We use the denotation $w(0) = u(0) = 0$. We also assume without loss of generality that $w_r(1) = 1$ if not stated otherwise ³. We define the bent welfare rule for some $b \geq 1$ and curvature $C \in [0, 1]$ as

$$w^{b,C}(j) = (1 - C)j + C \cdot \min\{j, b\}. \quad (3.12)$$

3.3 One Round Walks

3.3.1 Main Optimization Program

Our main results address the performance of one round walks in resource allocation games. Under this distributed algorithm design, each agent runs a local optimization to decide its decision in a sequential fashion. Given a set of allowable welfare rules \mathcal{W} and a utility function design \mathcal{U} , our first main result in Theorem 1 derives the efficiency guarantees of the one-round walk through a linear program construction.

Let $w_\ell \in \mathcal{W}$ and $\mathcal{U}(w_\ell) = u_\ell$ for some index ℓ . We use the notation $\bar{w}_\ell(i) = w_\ell(i)/w_\ell(1)$ and $\bar{u}_\ell(i) = u_\ell(i)/u_\ell(1)$ to simplify the presentation of the results. Additionally, we make the mild assumption that $u_\ell(1) = w_\ell(1)$ to normalize the utility rules.

Theorem 1. *Let \mathcal{W} be the welfare set. Consider the one-round walk with a utility function design \mathcal{U} , where $\mathcal{U}(w_\ell) = u_\ell$ for each $w_\ell \in \mathcal{W}$. The resulting efficiency guarantee*

³Consider any game $G \in \mathcal{G}_{\mathcal{W}, \mathcal{U}}$. For any resource $r \in \mathcal{R}$ with a welfare rule w_r , we can define another game $G' \in \mathcal{G}_{\mathcal{W}, \mathcal{U}}$ with instead $|w_r(1)|$ copies of resource r with a welfare rule $w'_r(j) = w_r(j)/w_r(1)$. If $|w_r(1)|$ is not integer, we can scale the number of resources uniformly and round to get arbitrarily close. Thus we can assume without loss of generality that $w_r(1) = 1$.

is $\text{Eff}(\mathcal{W}, \mathcal{U}; 1) = \inf_{w_\ell \in \mathcal{W}} 1/\beta(w_\ell)$, where $\beta(w_\ell) \in [1, \infty]$ is the solution to

$$\begin{aligned} \beta(w_\ell) = \min \quad & \beta \quad \text{subject to:} \\ \beta \bar{w}_\ell(y) \geq H \cdot & \left(\sum_{i=1}^y \bar{u}_\ell(i) - z \min_{1 \leq i \leq y+1} \bar{u}_\ell(i) \right) + \bar{w}_\ell(z) \end{aligned} \quad (3.13)$$

for all $z, y \in \mathbb{N}$ s.t. $z \geq 0$ and $y \geq 1$,

and $H = \sup\{\bar{w}_\ell(i)/i : i \in \mathbb{N}, w_\ell \in \mathcal{W}\}$.

Proof. Linear Program Formulation of the One-Round Walk. We first give a linear program that computes the efficiency $\text{Eff}_n(\mathcal{W}, \mathcal{U}; 1)$ 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 the given one round walk. 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 function design \mathcal{U} . We also make the assumption that the welfare set \mathcal{W} is finite, but generalize beyond this assumption later in the proof. A comparable primal-dual approach was also explored in [39] and [28] for different settings.

First, we apply a key observation that for a game G , truncating the action set of each agent i to $\mathcal{A}_i = \{a_i^\emptyset, a_i^{\text{br}}, a_i^{\text{opt}}\}$ does not affect the efficiency metric $\text{Eff}(G; 1)$. Here, a_i^\emptyset is the null action that does not select any resources, a_i^{br} is the action that agent i takes under a best response, and a_i^{opt} is the action that agent i plays in a joint action that optimizes the welfare $a^{\text{opt}} = \arg \max_{a \in \mathcal{A}} W(a)$.⁴ Therefore, we can restrict attention to the class of games $\mathcal{G}_{\mathcal{W}, \mathcal{U}}^{n,3} \subseteq \mathcal{G}_{\mathcal{W}, \mathcal{U}}^n$, where agents only have these three actions available without loss of generality. Furthermore, scaling W uniformly does not affect the ratio $\text{Eff}(G; 1) = \frac{W(a^{\text{br}})}{W(a^{\text{opt}})}$, and we can assume that $W(a^{\text{br}}) = 1$ without loss of generality. So we

⁴Note 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.

aim to find a game that maximizes the optimal welfare $W(a^{\text{opt}})$ to provide the lowest ratio. Consolidating the previous observations results in the following optimization problem

$$\text{Eff}_n(\mathcal{W}, \mathcal{U}; 1)^{-1} = \max_{G \in \mathcal{G}_{\mathcal{W}, \mathcal{U}}^{n, 3}} W(a^{\text{opt}}) \quad \text{subject to:} \quad (3.14)$$

$$W(a^{\text{br}}) = 1, \quad (3.15)$$

$$U_i(a_{j \leq i}^{\text{br}}, a_{j > i}^{\emptyset}) \geq U_i(a_{j < i}^{\text{br}}, a_i^{\text{opt}}, a_{j > i}^{\emptyset}) \quad \forall i \in \mathcal{I}, \quad (3.16)$$

The constraint inequality in Eq. (3.16) maintains that the joint action a^{br} is indeed a best response under the reduced action set $\mathcal{A}_i = \{a_i^{\emptyset}, a_i^{\text{br}}, a_i^{\text{opt}}\}$. To reformulate the optimization problem in Eq. (3.14) as a linear program, some necessary definitions are introduced. The possible resource allocations are enumerated by the following product set

$$\mathcal{P} = \prod_{i \in \mathcal{I}} \{\emptyset, \{a_i^{\text{br}}\}, \{a_i^{\text{opt}}\}, \{a_i^{\text{br}}, a_i^{\text{opt}}\}\},$$

where each resource is classified by the agent actions that can select it. Then the respective vectors in $\{0, 1\}^n$ can be defined.

$$\begin{aligned} b_i^p &= \left\{ 1 \text{ if } a_i^{\text{br}} \in p_i, 0 \text{ otherwise} \right\}, \\ o_i^p &= \left\{ 1 \text{ if } a_i^{\text{opt}} \in p_i, 0 \text{ otherwise} \right\}, \end{aligned}$$

where $p \in \mathcal{P}$ describes a resource type. We define the norm of b^p to be $|b^p| = \sum_{i \in \mathcal{I}} b_i^p$ (similarly for $|o^p| = \sum_{i \in \mathcal{I}} o_i^p$) and denote the number of nonzero elements before index i as $|b^p|_{< i} = \sum_{1 \leq j < i} b_j^p$. With this, we describe the linear program in the following lemma.

Lemma 1. *Consider the welfare set $\mathcal{W} = \{w_1, \dots, w_m\}$. For n agents, the efficiency*

guarantee of the one-round walk with the utility function design \mathcal{U} is

$$\text{Eff}_n(\mathcal{W}, \mathcal{U}; 1)^{-1} = \min_{\{\lambda_i \geq 0\}_{i \in \mathcal{I}}, \beta} \beta \quad \text{subject to:} \quad (3.17)$$

$$\beta \bar{w}_\ell(|\mathbf{b}^p|) \geq \bar{w}_\ell(|\mathbf{o}^p|) + \sum_{i \in \mathcal{I}} \lambda_i \left[(\mathbf{b}_i^p - \mathbf{o}_i^p) \bar{u}_\ell(|\mathbf{b}^p|_{<i} + 1) \right]$$

for all $p \in \mathcal{P}$ and $1 \leq \ell \leq m$.

Proof. First we show the equivalence of the optimization program proposed in Eq. (3.14) and the primal linear program described below. We later show that the dual of this primal program is exactly the linear program in Eq. (3.17). Note that each decision variable $\eta_p^\ell \in \mathbb{R}_{\geq 0}$ is a real non-negative number.

$$\text{Eff}_n(\mathcal{W}, \mathcal{U}; 1)^{-1} = \max_{\{\eta_p^\ell \geq 0\}_{\ell, p \in \mathcal{P}}} \sum_{\substack{1 \leq \ell \leq m, \\ p \in \mathcal{P}}} \bar{w}_\ell(|\mathbf{o}^p|) \cdot \eta_p^\ell \quad \text{s.t.} \quad (3.18)$$

$$\sum_{\substack{1 \leq \ell \leq m, \\ p \in \mathcal{P}}} \bar{w}_\ell(|\mathbf{b}^p|) \cdot \eta_p^\ell = 1 \quad (3.19)$$

$$\sum_{\substack{1 \leq \ell \leq m, \\ p \in \mathcal{P}}} \left[(\mathbf{b}_i^p - \mathbf{o}_i^p) \bar{u}_\ell(|\mathbf{b}^p|_{<i} + 1) \right] \cdot \eta_p^\ell \geq 0 \quad \forall i \in \mathcal{I} \quad (3.20)$$

For the equivalence, we first 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 vector 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 $\mathcal{R}^{\ell, p} = \{r \in \mathcal{R} : \ell_r = p, w_r = w_\ell\}$. Now we show that $W(a^{\text{opt}})$ in Eq. (3.14) matches

Eq. (3.18).

$$\begin{aligned}
W(a^{\text{opt}}) &= \sum_{r \in \mathcal{R}} w_r(|a^{\text{opt}}|_r) \\
&= \sum_{\substack{1 \leq \ell \leq m, \\ p \in \mathcal{P}}} \sum_{r \in \mathcal{R}^{\ell,p}} w_\ell(|a^{\text{opt}}|_r) \\
&= \sum_{\substack{1 \leq \ell \leq m, \\ p \in \mathcal{P}}} \bar{w}_\ell(|\text{o}^p|) \cdot \eta_p^\ell,
\end{aligned}$$

where $\eta_p^\ell = |\mathcal{R}^{\ell,p}| \cdot w_\ell(1)$. 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 $|a^{\text{opt}}|_r = \sum_{j \in \mathcal{I}} \mathbb{1}_{a_j^{\text{opt}}}(r) = |\text{o}^p|$ if $r \in \mathcal{R}^{\ell,p}$; additionally, the value $w_\ell(|\text{o}^p|)$ is constant for any $r \in \mathcal{R}^{\ell,p}$. A similar argument can be made about the welfare of the best response action $W(a^{\text{br}})$, so Eq. (3.15) matches Eq. (3.19) as well.

Now we show the utility constraint in Eq. (3.16) matches the constraint in Eq. (3.20). For conciseness, let $a^1 = (a_{j < i}^{\text{br}}, a_i^{\text{br}}, a_{j > i}^{\emptyset})$ and $a^2 = (a_{j < i}^{\text{br}}, a_i^{\text{opt}}, a_{j > i}^{\emptyset})$. The utility difference can be written as

$$\begin{aligned}
U_i(a^1) - U_i(a^2) &= \sum_{r \in a_i^{\text{br}}} u_r(|a^1|_r) - \sum_{r \in a_i^{\text{opt}}} u_r(|a^2|_r) \\
&= \sum_{r \in \mathcal{R}} \left(\mathbb{1}_{a_i^{\text{br}}}(r) u_r(|a^1|_r) - \mathbb{1}_{a_i^{\text{opt}}}(r) u_r(|a^2|_r) \right)
\end{aligned}$$

$$\begin{aligned}
&= \sum_{\substack{1 \leq \ell \leq m, \\ p \in \mathcal{P}}} \sum_{r \in \mathcal{R}^{\ell, p}} \left(\mathbb{1}_{a_i^{\text{br}}}(r) u_r(|a^1|_r) - \mathbb{1}_{a_i^{\text{opt}}}(r) u_r(|a^2|_r) \right) \\
&= \sum_{\substack{1 \leq \ell \leq m, \\ p \in \mathcal{P}}} \sum_{r \in \mathcal{R}^{\ell, p}} \left[(b_i^p - o_i^p) u_\ell(|b^p|_{<i} + 1) \right] \\
&= \sum_{\substack{1 \leq \ell \leq m, \\ p \in \mathcal{P}}} \left[(b_i^p - o_i^p) \bar{u}_\ell(|b^p|_{<i} + 1) \right] \eta_p^\ell.
\end{aligned}$$

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 and using the assumption that $u_\ell(1) = w_\ell(1)$.

We assume that $\eta_p^\ell \geq 0$ to ensure a well-defined game parametrization. Observe that in the primal program in Eq. (3.18), we have relaxed $\eta_p^\ell \in \mathbb{R}_{\geq 0}$ to be any non-negative real number with η_p^ℓ denoting the relative fraction of resources with a specific resource type. We use this relaxation to normalize $W(a^{\text{br}}) = 1$ and this relaxation is done without loss of generality, since we can scale up the values $\{\eta_p^\ell\}_{\ell, p \in \mathcal{P}}$ (from the solution arguments of Eq. (3.18)) uniformly and round to derive the resource set for a corresponding valid game construction G that achieves an efficiency ratio $\text{Eff}(G; 1)$ that is arbitrarily close to the solution of the primal program.

We now verify that the dual of the primal program in Eq. (3.18) matches the linear program defined in Eq. (3.17). Note that primal program in Eq. (3.18) can be concisely

written as

$$\begin{aligned} & \max_{\eta} \quad c^T \eta \quad \text{subject to:} \\ & K\eta = 1 \\ & \begin{bmatrix} L \\ I_{m \cdot 4^n} \end{bmatrix} \eta \succeq 0, \end{aligned}$$

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 , K , L are the compactly written vectors in equations (3.18), (3.19), and (3.20) respectively. Writing the dual linear program gives

$$\begin{aligned} & \max_{\lambda \succeq 0, \xi \succeq 0, \beta} \quad -\beta \quad \text{subject to:} \\ & K_\ell^T \beta - \begin{bmatrix} L_\ell^T \\ I_{4^n} \end{bmatrix} \begin{bmatrix} \lambda \\ \xi \end{bmatrix} - c_\ell = \mathbf{0} \quad \forall 1 \leq \ell \leq m, \end{aligned}$$

where $c = (c_1^T, \dots, c_m^T)^T$ is associated with each $1 \leq \ell \leq m$ (likewise for K and L). Observe that the constraint set $K_\ell^T \beta - \begin{bmatrix} L_\ell^T \\ I_{4^n} \end{bmatrix} \begin{bmatrix} \lambda \\ \xi \end{bmatrix} - c_\ell = \mathbf{0}$ is equivalently written as $K_\ell^T \beta - L_\ell^T \lambda - c_\ell = \xi$ and as $K_\ell^T \beta - L_\ell^T \lambda - c_\ell \succeq 0$. Substituting back c_ℓ , K_ℓ , L_ℓ results in the constraint

$$\beta \bar{w}_\ell(|\mathbf{b}^p|) \geq \bar{w}_\ell(|\mathbf{o}^p|) + \sum_{i \in \mathcal{I}} \lambda_i \left[(\mathbf{b}_i^p - \mathbf{o}_i^p) \bar{u}_\ell(|\mathbf{b}^p|_{<i} + 1) \right],$$

which matches the constraint outlined in Eq. (3.17). \square

Continuing the Proof of Theorem 1. The dual program in Eq. (3.17) provides a solution for $\text{Eff}_n(\mathcal{W}, \mathcal{U}; 1)^{-1} = \beta^*$ for a fixed n and finite \mathcal{W} . However, the constraint

set is exponential in the number of agents. Thus, in this section, we remove redundant constraints to arrive at a more tractable linear program. We first show the solution is upper bounded by $\beta^* \leq \tilde{\beta}$ for any n , where $\tilde{\beta} = \max_{1 \leq \ell \leq m} \beta(w_\ell)$ and $\beta(w_\ell)$ is the solution to the program in Eq. (3.13).

Let n be the number of agents. Without loss of generality, we assume that $w_\ell(1) = u_\ell(1) = 1$ for $1 \leq \ell \leq m$. 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, $B^p : \{1, \dots, y_p\} \rightarrow \{1, \dots, n\}$ and $O^p : \{1, \dots, z_p\} \rightarrow \{1, \dots, n\}$ with

$$B^p(j) = i \text{ if } b_i^p = 1 \text{ and } |b^p|_{\leq i} = j,$$

$$O^p(j) = i \text{ if } o_i^p = 1 \text{ and } |o^p|_{\leq i} = j.$$

Considering the dual program in Eq. (3.17), we add the constraint that $\lambda_i = H = \sup\{w_\ell(i)/i : i \in \mathbb{N}, 1 \leq \ell \leq m\}$ explicitly. Since we shrink the feasible region, the optimal solution to Eq. (3.17) 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

$$0 \geq w_\ell(z_p) - \sum_{j=1}^{z_p} \lambda_{O^p(j)} u_\ell(1) \text{ for all } \ell, p.$$

Simplifying the expression gives $\sum_{j=1}^{z_p} \lambda_{O^p(j)} \geq w_\ell(z_p)$, which is always satisfied if $\lambda_i = H$ for all i . If the constraints according p are such that $b^p \neq \mathbf{0}$, then the term $\beta w_\ell(y)$ is present and strictly positive in the inequality (3.17) 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 Eq. (3.17),

for each ℓ , to

$$\beta w_\ell(y_p) \geq w_\ell(z_p) + \sum_{i=1}^{y_p} \text{Hu}_\ell(i) - \sum_{i \in \mathcal{I}} \text{Ho}_i^p u_\ell(|b^p|_{<i} + 1).$$

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

$$\beta w_\ell(y) \geq w_\ell(z) + \sum_{i=1}^y \text{Hu}_\ell(i) - \sum_{i \in \mathcal{I}} \text{Hz} \min_{1 \leq i \leq y+1} u_\ell(i),$$

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 $\beta^* \leq \tilde{\beta}$. Since $\tilde{\beta}$ is the only variable in the optimization problem, we can decouple the constraints for each ℓ and limit the number of agents $n \rightarrow \infty$ to arrive at the program in Eq. (3.13).

Now we show that the solution is lower bounded by $\beta^* \geq \tilde{\beta}$, where β^* and $\tilde{\beta}$ are defined as before. We show that when we remove dual constraints, we arrive at the set of linear programs in Eq. (3.13). Since the feasible region expands, the optimal solution potentially decreases. Let the set of agents be $\mathcal{I} = \mathbb{N}$ and $j_\ell^p = \arg \min_{1 \leq j \leq y_p+1} u_\ell(j)$. We remove all the dual constraints barring the constraints that correspond to $p \in \mathcal{P}$ with either **(a)** $y_p = 0$ and $z_p = z_\ell^* = \arg \max w_\ell(j)/j$ or **(b)** $y_p > 0$ and $B^p(j_\ell^p - 1) < O^p(1)$ and $O^p(z_p) < B^p(j_\ell^p)$. The first property refers to all resource types where a^{br} is never selected but a^{opt} is by z_ℓ^* agents. The second property refers to all resource types where the indices of the agents selecting a^{opt} are between the agents with index $B^p(j_\ell^p - 1)$ and $B^p(j_\ell^p)$.

Assume property (a). Then the corresponding dual constraint in Eq. (3.17) can be written as

$$0 \geq w_\ell(z_\ell^*) - \sum_{j=1}^{z_\ell^*} \lambda_{O^p(j)} u_\ell(1),$$

for any resource type $p \in \mathcal{P}$ that satisfies property (a) and for all ℓ . Therefore, for any $j \in \mathbb{N}$, except for at most $z^* - 1$ (with $z^* \equiv \max_{\ell} \{z_{\ell}^*\}$) values, observe that $\lambda_j \geq H$ must hold.

Now assume property (b). With respect to a resource type $p \in \mathcal{P}$ that satisfies property (b), we observe that $u_{\ell}(|b^p|_{<i} + 1) = u_{\ell}(j_{\ell}^p)$ for any agent with index $i = O^p(j)$ for some j . Therefore, under the two previous observations, we can rewrite the relaxed dual program as

$$\min_{\lambda \geq \mathbf{0}} \beta \quad \text{subject to:} \quad (3.21)$$

$$\beta w_{\ell}(y_p) \geq \sum_{j=1}^{y_p} \lambda_{B^p(j)} u_{\ell}(j) - \sum_{j=1}^{z_p} \lambda_{O^p(j)} u_{\ell}(j_{\ell}^p) + w_{\ell}(z_p)$$

for all $p \in \mathcal{P}'$ and ℓ ,

$$\lambda_i \geq H \quad \text{for all } i \in \mathbb{N} \text{ but at most } z^* - 1 \text{ values,}$$

where $\mathcal{P}' = \{p \in \mathcal{P} : p \text{ satisfies property (b)}\}$. Observe that we recover the proposed program given in Eq. (3.13) if we assume that the optimal dual variable is $\lambda_i = H$ for all $i \in \mathbb{N}$. To show this claim, we confirm that the binding constraint for β in Eq. (3.21) is larger when considering a different sequence of lambdas $\lambda \neq H\mathbf{1}$. In other words, for a given $y \geq 1$ and $z \geq 0$, we show that for the resulting dual variables,

$$\begin{aligned} \beta_{\lambda} &:= \max_{p \in \mathcal{P}'} \left\{ \frac{1}{w_{\ell}(y_p)} \left(\sum_{j=1}^{y_p} \lambda_{B^p(j)} u_{\ell}(j) - \sum_{j=1}^{z_p} \lambda_{O^p(j)} u_{\ell}(j_{\ell}^p) \right) \right\} \\ &\geq \frac{H}{w_{\ell}(y)} \left(\sum_{j=1}^y u_{\ell}(j) - \sum_{j=1}^z u_{\ell}(j^p) \right) := \beta_{y,z} \end{aligned} \quad (3.22)$$

For any $\lambda \neq H\mathbf{1}$, 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 $\lim_{j \rightarrow \infty} \lambda_j = \infty$. 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 Eq. (3.22) 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 $\lambda_{B^p(j)} > M$ for all $1 \leq j \leq y_p$. Thus $\beta_\lambda \geq \frac{1}{w_\ell(y_p)} \sum_{j=1}^{y_p} M u_\ell(j)$. 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 , for all but a finite set of values, there exists a convergent sub-sequence λ_{ss} that converges to a value $V \geq H$ 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 $|\lambda_{B^p(j)} - V| \leq \frac{\varepsilon}{2M_u^y x}$ and $|\lambda_{O^p(j)} - V| \leq \frac{\varepsilon}{2M_u^y x}$ for all j . Therefore

$$\begin{aligned} \beta_\lambda &\geq \frac{1}{w_\ell(y_p)} \left(\sum_{j=1}^{y_p} \lambda_{B^p(j)} u_\ell(i) - \sum_{j=1}^{z_p} \lambda_{O^p(j)} u_\ell(j^p) \right) \\ &\geq \frac{V}{w_\ell(y)} \left(\sum_{j=1}^y u_\ell(i) - \sum_{j=1}^z u_\ell(j^p) \right) - \frac{\varepsilon}{2} - \frac{\varepsilon}{2} \\ &\geq \beta_{y,z} - \varepsilon. \end{aligned}$$

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 for any n and we have shown the equality $\beta^* = \tilde{\beta}$ for any n .

Note that the welfare rule set \mathcal{W} was assumed to be finite for the previous arguments. Now we extend to more general sets of welfare rules. As the worst case efficiency is defined as $\text{Eff}(\mathcal{W}, \mathcal{U}; 1) = \inf_{G \in \mathcal{G}_{\mathcal{W}, \mathcal{U}}} \text{Eff}(G; 1)$, for a given sequence $\varepsilon_j \rightarrow 0$, there always exists a game $G_j \in \mathcal{G}_{\mathcal{W}, \mathcal{U}}$ such that $\text{Eff}(G_j; 1) \leq \text{Eff}(\mathcal{W}, \mathcal{U}; 1) + \varepsilon_j$. Take a sequence of $\varepsilon_j \rightarrow 0$. Define a finite welfare set for each step as $\mathcal{W}_j := \bigcup_{1 \leq k \leq j} \{w_r : r \in \mathcal{R}_k\}$ described as the

union of the welfare rules for each game in the sequence. For each step, we have that $\text{Eff}(\mathcal{W}, \mathcal{U}; 1) \leq \tilde{\beta} \leq \text{Eff}(G_j; 1) \leq \text{Eff}(\mathcal{W}, \mathcal{U}; 1) + \varepsilon_j$, where $\tilde{\beta}$ is the solution derived from the set of linear programs in Eq. (3.13) for the welfare set \mathcal{W}_j . Taking $j \rightarrow \infty$ gives the result. \square

Remark 1. Observe that the value $\beta(\tilde{w}_\ell) = \beta(w_\ell)$ is equal for the welfare rule $\tilde{w}_\ell = a \cdot w_\ell$ for any $a > 0$ if the corresponding utility rule also satisfies $\tilde{u}_\ell = a \cdot u_\ell$. For a finite collection of welfare rules, say $\{w_1, \dots, w_m\}$, consider the set $\mathcal{W} = \{w : w = \sum_{j=1}^m a_j w_j : a_j \geq 0 \text{ for all } j\}$ that can be defined by the possible non-negative linear combinations. In this instance, the quantity $\inf_{w_\ell \in \mathcal{W}} 1/\beta(w_\ell)$ is then equal to $\min\{1/\beta(w_1), \dots, 1/\beta(w_m)\}$, which can be computed from a finite set of linear programs.

Remark 2. Note that if the number of agents n is known, the linear program in Eq. (3.13) provides a non-trivial lower bound for $1/\beta(w_\ell)$ when only including the constraints for $z, y \leq n$. Thus, it is possible to derive lower bounds on the efficiency guarantees through a set of tractable optimization programs.

The above theorem sets forth a prescriptive process by which to characterize the efficiency guarantees of the one round walk through a linear program construction ⁵. This is done through a novel parametrization of the set of resource allocation games and careful elimination of the redundant constraints in the dual of resulting program. While directly solving the optimization in Eq. (3.13) requires keeping track of a countable number of constraints, this linear program construction provides valuable insights into the achievable efficiency guarantees. Under certain sub-classes of welfare rules, the optimal utility function design that optimizes $\text{Eff}^*(\mathcal{W}; 1)$ can actually be derived in closed form. This is done in the next two subsections with regards to *submodular* and *supermodular* welfare rules.

⁵We note that the LP in Eq. (3.13) is decoupled for each welfare rule w_ℓ .

3.3.2 One Round Walks in Submodular Settings




Figure 3.2: In the top chart, following the results in Theorem 3, we visually compare the efficiency guarantees of the best one-round walk with the guarantees of the common interest version and the best achievable polynomial-time guarantees. We note that the best one-round walk and common-interest are at worst $\sim 80\%$ within the best polynomial-time guarantees. In the bottom chart, we display the fractional performance gains of the best one-round walk over the common-interest one (up to $\sim 13\%$ better).

In this subsection, 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. Many well-studied applications including viral marketing, information gathering, image segmentation, statistical summarization involve welfare objectives that are submodular (see [33] for a survey on application domains). Example (2) involves welfare rules that are submodular. We

formally define submodular welfare rules below.

Definition 1 (Submodularity). A welfare rule w is submodular if w is non-decreasing and concave in j , or equivalently that $w(j+1) - w(j)$ is non-negative and non-increasing in j .

Under the assumption of submodularity, we can simplify the linear program in Theorem 1.

Corollary 1. *Let \mathcal{W} be a set of submodular welfare rules. Consider the one-round walk with a utility function design \mathcal{U} . The resulting efficiency guarantee is $\text{Eff}(\mathcal{W}, \mathcal{U}; 1) = \inf_{w_\ell \in \mathcal{W}} 1/\beta(w_\ell)$, where $\beta(w_\ell)$ is given by Eq. (3.13) with $H = 1$.*

Proof. We directly apply Theorem 1 to derive the efficiency guarantees for submodular welfare rules. Note that if w_ℓ is submodular, $\bar{w}_\ell(i) \leq i$ for all i . Thus, we can substitute $H = \sup_{i,\ell} \bar{w}_\ell(i)/i = \bar{w}_\ell(1)/1 = 1$. \square

Furthermore, given a set of submodular welfare rules, we can derive the optimal one-round walk, as well as its respective efficiency guarantees, through a corresponding linear program. The construction of the program is derived from the characterization result in Corollary 1 with the non-trivial fact that the optimal utility rules are non-increasing in this domain.

Theorem 2. *Consider the set \mathcal{W} where each $w_\ell \in \mathcal{W}$ is a submodular welfare rule. The utility rules u_ℓ^1 of the optimal utility design for the one round walk are given by the*

solutions to

$$\begin{aligned}
(u_\ell^1, \beta(w_\ell)) &\in \arg \min_{\beta, u \in \mathbb{R}_{\geq 0}^{\mathbb{N}}} \beta \quad \text{subject to:} & (3.23) \\
\beta w_\ell(y) &\geq \sum_{i=1}^y u(i) - zu(y+1) + w_\ell(z) \quad \forall y, z \geq 1, \\
u(1) &= w_\ell(1),
\end{aligned}$$

with a corresponding efficiency guarantee of $\text{Eff}^*(\mathcal{W}; 1) = \inf_{w_\ell \in \mathcal{W}} 1/\beta(w_\ell)$.

Proof. We simply refer to \bar{w}_ℓ as w in the following discussion. If the utility rule u is assumed to be non-increasing, we will show that we recover the linear program in Eq. (3.23). If u is 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 Eq. (3.13) to produce the linear program in Eq. (3.23), rewritten below.

$$\begin{aligned}
(\beta^*, u^*) &\in \arg \min_{\beta, \{u(i)\}_{i \in \mathcal{I}}} \beta \quad \text{subject to:} & (3.24) \\
\beta w(y) &\geq \sum_{i=1}^y u(i) - zu(y+1) + w(z) \quad \forall y, z \geq 1 \\
u(1) &= 1,
\end{aligned}$$

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 for this simplified program. First, rearranging the terms in the constraint in Eq. (3.24) gives that for any $y \geq 1$,

$$u^*(y+1) \geq \sup_{z \geq 1} \left(\frac{1}{z} \left(\sum_{i=1}^y u^*(i) + w(z) - \beta^* w(y) \right) \right). \quad (3.25)$$

We verify $u^*(y+1)$ is well-defined. Note that since u^* is optimal, the efficiency bound $\beta^* < \infty$ is nontrivial (as the common interest design guarantees an efficiency guarantee greater than $1/2$ [40]). 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 Eq. (3.25) 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 $z \rightarrow \infty$ in Eq. (3.25) 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 $z_{y+1} \in \arg \max_{z \geq 1} w(z) - zu(y+1)$ be the number that achieves the maximum.

We verify that z_{y+1} is well-defined. Suppose for contradiction that $w(z) - zu^*(y+1)$ is always increasing in z , so z_{y+1} is not well defined. Since $\beta^* < \infty$, the limit $\lim_{z \rightarrow \infty} w(z) - zu^*(y+1)$ must converge and therefore $u^*(y+1)$ must be equal to $Q = \lim_{z \rightarrow \infty} \Delta w(z)$, 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 Eq. (3.24), with $y-1$ and $z \rightarrow \infty$ gives $\beta w(y-1) \geq \lim_{z \rightarrow \infty} w(z) - zu^*(y) \geq \infty$, which is a contradiction.

Now, substituting z_{y+1} into Eq. (3.25) for y and $y+1$ produces the following expressions

$$\begin{aligned} u^*(y+1) &= \frac{1}{z_{y+1}} \left(\sum_{i=1}^y u^*(i) + w(z_{y+1}) - \beta^* w(y) \right) \\ u^*(y) &\geq \frac{1}{z_{y+1}} \left(\sum_{i=1}^{y-1} u^*(i) + w(z_{y+1}) - \beta^* w(y-1) \right). \end{aligned}$$

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 Eq. (3.24) for $y+j$

and $y + j + 1$ gives

$$u^*(y + j + 1) \geq \frac{1}{z_{y+j}} \left(\sum_{i=1}^{y+j} u^*(i) + w(z_{y+j}) - \beta^* w(y + j) \right)$$

$$u^*(y + j) = \frac{1}{z_{y+j}} \left(\sum_{i=1}^{y+j-1} u^*(i) + w(z_{y+j}) - \beta^* w(y + j - 1) \right).$$

Thus by substituting the second expression into first, the following inequality holds

$$u^*(y + j + 1) \geq u^*(y + j) + \frac{u^*(y + j) - \beta^* \Delta w(y + j)}{z_{y+j}}. \quad (3.26)$$

We show, by induction, that the following expression holds for any $j \geq 1$,

$$\frac{u^*(y + j) - \beta^* \Delta w(y + j)}{z_{y+j}} \geq \frac{u^*(y + 1) - \beta^* \Delta w(y + 1)}{z_{y+1}} > 0. \quad (3.27)$$

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 Eq. (3.26) 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 Eq. (3.27) 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 Eq. (3.27) 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 Eq. (3.24) where $y \rightarrow \infty$ and $z = 0$. Since $w(y) \leq y$,

$$\beta^* \geq \lim_{y \rightarrow \infty} \frac{1}{y} \sum_{i=1}^y u^*(i) \geq \infty, \quad (3.28)$$

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 the linear program. \square

Remark 3. Note that the optimization problem in Eq. (3.23) is intractable to solve directly. However, if we fix the number of agents to n , and only consider the variables $u \in \mathbb{R}_{\geq 0}^n$ that are non-increasing and constraints for $1 \leq y, z \leq n$, we can derive lower bounds on the optimal efficiency guarantees.

Thus, in the submodular setting, it is possible to derive characterizations of the optimal one-round walk. While computing these characterizations is intractable in general, it is possible to compute the optimal one-round walk guarantees in closed form for certain classes of welfare rules. In fact, this is possible if the submodular welfare rules are parametrized through their *curvature*. Curvature is a classical parametrization used widely in submodular optimization problems (see [41, 42]) that characterizes the rate of diminishing returns associated with a submodular welfare function. We note that any submodular welfare function has a curvature $C \in [0, 1]$. In our setting, curvature can be defined as follows.

Definition 2 (Curvature). A submodular welfare rule w has a curvature of $C \in [0, 1]$ if $C = 1 - \lim_{n \rightarrow \infty} (w(n + 1) - w(n)) / w(1)$.

With this, we can arrive at a tight, closed-form characterization of the optimal performance guarantees, as shown below.

Theorem 3. *Let the set \mathcal{W} comprise of all submodular welfare rules w that have curvature of at most $C \in [0, 1]$. The efficiency guarantees of the one round walk for the optimal utility design as well as common interest utility design are defined by*

$$\text{Eff}^*(\mathcal{W}; 1) = 1 - \frac{C}{2}, \quad (3.29)$$

$$\text{Eff}(\mathcal{W}; \mathcal{U}_{\text{CI}}, 1) = (1 + C)^{-1}. \quad (3.30)$$

Furthermore, the utility rules $u_\ell^1 = \mathcal{U}_1(w_\ell)$ for the optimal utility design can be compactly expressed as

$$u_\ell^1(j) = \sum_{b \in \mathbb{N}} a_b u_b(j), \quad (3.31)$$

where $a_b \in \mathbb{R}_{\geq 0}$ and u_b are defined in Eq. and Eq.

Proof. Proof of Efficiency for the Common Interest Utility Design. We verify the equality in Eq. (3.30). The common interest utility design can be defined through the utility rules $u_\ell^{\text{mc}} = \mathcal{U}_{\text{CI}}(w_\ell)$ that satisfy $u_\ell^{\text{mc}}(j) = w_\ell(j) - w_\ell(j - 1)$ for all j . Now, we can use Corollary 1 to characterize the efficiency guarantee of the common interest utility design.

Consider the set \mathcal{W} of submodular welfare rules that have curvature of at most C and assume $w_\ell(1) = 1$ without loss of generality. The utility rule u_ℓ^{mc} must be non-increasing, and the constraints in Eq. (3.13) can be rewritten as

$$\beta w_\ell(y) \geq \sum_{i=1}^y u_\ell^{\text{mc}}(i) - z u_\ell^{\text{mc}}(y + 1) + w_\ell(z), \quad (3.32)$$

for any $y \geq 1$ and $z \geq 0$. We claim the binding constraint is when $z = y$. Fixing y , the

only terms that depend on z is $-zu_\ell^{\text{mc}}(y+1) + w_\ell(z)$. Examining the difference between terms from $z+1$ against z gives

$$\begin{aligned} & w_\ell(z+1) - (z+1)u_\ell^{\text{mc}}(y+1) - w_\ell(z) + zu_\ell^{\text{mc}}(y+1) \\ &= w_\ell(z+1) - w_\ell(z) - u_\ell^{\text{mc}}(y+1) \\ &= u_\ell^{\text{mc}}(z+1) - u_\ell^{\text{mc}}(y+1). \end{aligned}$$

Since u_ℓ^{mc} is non-increasing, note that $u_\ell^{\text{mc}}(z+1) - u_\ell^{\text{mc}}(y+1)$ is greater than 0 if $z < y$ and less than 0 if $z > y$. Therefore the tightest constraint is when $z = y$. Now we simplify the solution for $\beta(w_\ell)$ in Eq. (3.13) under the assumption that $z = y$ as

$$\begin{aligned} \beta(w_\ell) &= \max_{y \geq 1} \left\{ \frac{1}{w_\ell(y)} \left(\sum_{j=1}^y u_\ell^{\text{mc}}(j) - yu_\ell^{\text{mc}}(y+1) + w_\ell(y) \right) \right\} \\ &= \max_{y \geq 1} \left\{ 2 - \frac{y}{w_\ell(y)} u_\ell^{\text{mc}}(y+1) \right\}, \end{aligned}$$

in which we have used the identity $\sum_{j=1}^y u_\ell^{\text{mc}}(j) = \sum_{j=1}^y w_\ell(j) - w_\ell(j-1) = w_\ell(y)$. Since w_ℓ is submodular, $j/w_\ell(j) \geq 1$ for any $j \in \mathbb{N}$, and because w_ℓ has at most curvature of C , $u_\ell^{\text{mc}}(j) \geq 1 - C$ for any $j \in \mathbb{N}$ as well. Therefore, the solution is upper bounded by $\beta(w_\ell) \leq 1 + C$ and since w was chosen arbitrarily from \mathcal{W} , the resulting efficiency guarantee is $\text{Eff}(\mathcal{G}_{\mathcal{W}, \mathcal{U}_{\text{mc}}}; 1) = \inf_\ell 1/\beta(w_\ell) \geq (1 + C)^{-1}$. This efficiency guarantee is actually tight if we consider the b -covering welfare rule $w^{b,C}$ with curvature C , as in Eq. (3.34). Observe that under the b -covering welfare, the maximum is $\max_{y \geq 1} \frac{y}{w^{b,C}(y)} u^{\text{mc}}(y+1) = 1 - C$ at $y = b$. Therefore, $\text{Eff}(\mathcal{G}_{\mathcal{W}, \mathcal{U}_{\text{mc}}}; 1) = (1 + C)^{-1}$ with equality and we show the claim. \square

Proof of Efficiency of the Optimal Utility Design. We verify the equality in Eq. (3.29) and structure of the optimal utility rules. Given a curvature C , let \mathcal{W} be the set of welfare rules that have curvature of at most C . From [42, Lemma 2], 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 $w = \sum_{b \in \mathbb{N}} \alpha^b w^{b,C}$, with

$$\alpha^b = (2w(b) - w(b-1) - w(b+1))/C \text{ and} \quad (3.33)$$

$$w^{b,C}(j) = \begin{cases} j, & \text{if } 0 \leq j \leq b \\ b + (1 - C) \cdot (j - b) & \text{if } j > b. \end{cases} \quad (3.34)$$

We refer to these welfare rules as *b-covering* welfare rules. We note that for any $b \in \mathbb{N}$, the welfare rule $w^{b,C}$ has a curvature of C . We consider a linear utility function design $\mathcal{U}_{\text{lin}}(w_\ell = \sum_{b \in \mathbb{N}} \alpha^b w^{b,C}) = \sum_{b \in \mathbb{N}} \alpha^b u^b$. Note that the constraint in Eq. (3.13) is satisfied for any linear combination of $w^{b,C}$ and u^b , so only we only need to confirm optimality of u^b for each b . For each welfare rule $w^{b,C}$, we claim that the corresponding optimal utility rule from running the program in Eq. (3.23) is

$$u^b(j) = \begin{cases} (1 - \beta^b) \left(\frac{b+1}{b}\right)^{j-1} + \beta^b & \text{if } j \leq b + 1 \\ (1 - C)\beta^b & \text{if } j \geq b + 1, \end{cases} \quad (3.35)$$

where $\beta^b = \frac{(\frac{b+1}{b})^b}{(\frac{b+1}{b})^b - C}$ is the resulting optimal efficiency. Taking the minimum across b , we have that $\min_{b \in \mathbb{N}} \frac{1}{\beta^b} = 1 - C/2$ for $b = 1$. Therefore, using Theorem 2, the optimal efficiency guarantee is $\text{Eff}^*(\mathcal{W}; 1) = 1 - C/2$.

Now we verify that u^b and β^b are indeed the optimal solutions. We first remove all constraints in Eq. (3.23) 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 Eq. (3.24) gives that for any $y \geq 1$, the

optimal solution satisfies

$$u^*(y+1) = \sup_{z \geq 1} \left(\frac{1}{z} \left(\sum_{i=1}^y u^*(i) + w(z) - \beta^* w(y) \right) \right). \quad (3.36)$$

Substituting in for w and the binding constraint $z = b$, the recursive equation for u^b is then

$$\begin{aligned} u^b(1) &= 1 \\ u^b(j+1) &= \frac{1}{b} \sum_{i=1}^j u^b(i) + 1 - \frac{1}{b} \beta^* w^{b,C}(j), \end{aligned}$$

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.

$$\begin{aligned} x_1(t+1) &= x_1(t) + x_2(t) \\ x_2(t+1) &= \frac{1}{b}(x_1(t) + x_2(t)) + s(t) \\ s(t) &= 1 - \frac{1}{b} \beta^* w^{b,C}(t). \end{aligned}$$

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

$$\begin{aligned} x_2(1) &= 1 \\ x_2(t) &= \frac{1}{b} B^{t-2} + \sum_{\tau=1}^{t-2} \frac{1}{b} B^{t-2-\tau} (1 - \beta^* w^{b,C}(\tau)) \\ &\quad + (1 - \beta^* w^{b,C}(t-1)) \quad t > 1, \end{aligned} \quad (3.37)$$

where $B = \frac{b+1}{b}$. Simplifying the expression for $x_2(t)$ for $t-1 > b$ and substituting

$w^{b,C}(t) = (1 - C)t + C \min(t, b)$ results in the following

$$\begin{aligned} x_2(t) &= \frac{1}{b} B^{t-2} \left(1 + \sum_{\tau=1}^b B^{-\tau} (1 - \beta^* \tau) \right. \\ &\quad \left. + \sum_{\tau=b+1}^{t-2} B^{-\tau} (1 - \beta^* ((1 - C)\tau + Cb)) \right) \\ &\quad + (1 - \beta^*(t - 1 - C(t - 1) + Cb)). \end{aligned}$$

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

$$x_2(t) = B^{t-2}(\beta^*(CB^{1-b} - B) + B) + (1 - C)\beta^*.$$

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 2. This is only possible when $\beta^* \geq \frac{B^b}{B^b - C}$. Therefore the optimal solution must be $\beta^* = \beta^b = \frac{B^b}{B^b - C}$. Substituting for β^* in the expression in Eq. (3.37) and simplifying results in the closed form expression in Eq. (3.35) for u^b . It can be seen that u^b defined in Eq. (3.35) 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 Eq. (3.23), we examine the terms $w^{b,C}(z) - zu^b(y + 1)$ for any $y \geq 1$. Note that $1 = w^{b,C}(z) - w^{b,C}(z - 1) \geq u^b(y + 1)$ when $z \leq b$ and $(1 - C) = w^{b,C}(z) - w^{b,C}(z - 1) \leq u^b(y + 1)$ when $z \geq b$ for any y . Thus the maximum $\max_z w^{b,C}(z) - zu^b(y + 1)$ occurs when $z = b$, and we have shown the claim. \square

Remark 4. We note that the efficiency guarantees, in Eq. (3.30), of the common interest algorithm exactly matches the bound given for general submodular set functions [41].

In Theorem 3, we have characterized the efficiency guarantees of the optimal utility design and common interest design in closed form. A visual comparison of the guarantees is depicted in Figure 3.2. We also compare the efficiency guarantees to the best approximation guarantee $1 - C/e$ that is achievable by any polynomial time algorithm [43, Theorem 2] in this setting. By only carefully designing the objectives that agents greedily optimize against, we see that there can be significant gains in the performance guarantees in submodular resource allocation games.

3.3.3 One Round Walks in Supermodular Settings

In this subsection, we now consider welfare rules that are *supermodular*. Under this welfare structure, cooperative resource utilization results in a surplus of system welfare. Applications of supermodular games include clustering (see Example (3) for more details) and power allocation in networks [43]. A formal definition of supermodular welfare rules is as follows.

Definition 3 (Supermodularity). A welfare rule w is supermodular if w is non-decreasing and convex in j , or that $w(j+1) - w(j)$ is non-negative and non-decreasing in j .

Unlike in the submodular setting, the efficiency guarantees of the optimal utility design and the common interest utility design can be characterized in closed form for supermodular welfare rules. This is done in the following theorem.

Theorem 4. Consider the set \mathcal{W} where each $w_\ell \in \mathcal{W}$ is a supermodular welfare rule. The efficiency guarantees of the optimal utility design and the common interest utility design are

$$\text{Eff}(\mathcal{W}, \mathcal{U}_{\text{CI}}, 1) = \text{Eff}^*(\mathcal{W}; 1) = \inf_{w_\ell \in \mathcal{W}} \lim_{n \rightarrow \infty} \frac{n}{\bar{w}_\ell(n)} \quad (3.38)$$

Furthermore, the utility rules $u_\ell^1 = \mathcal{U}_1(w_\ell)$ for the optimal design is any rule that is

non-decreasing and satisfies

$$\sum_{i=1}^j u_\ell^1(i)/w_\ell(j) \leq 1 \text{ for all } j \geq 1.$$

Proof. We first show that $\text{Eff}^*(\mathcal{W}; 1) \leq \min_{1 \leq \ell \leq m} \lim_{n \rightarrow \infty} \frac{n}{\bar{w}_\ell(n)}$ for a finite supermodular welfare rule set ⁶. We do this through a game construction, depicted in Figure 3.3. Let $w^* = \arg \min_{1 \leq \ell \leq m} \frac{n}{\bar{w}_\ell(n)}$ be the welfare rule rule that attains the minimum. Let the game G have n agents with agent i having the action set $\mathcal{A}_i = \{a_i^\emptyset, a_i^{\text{br}}, a_i^{\text{opt}}\}$. 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 $j \neq i$ have selected r_1 through a_j^{opt} . Thus a^{br} is a possible solution with a welfare of $W(a^{\text{br}}) = n \cdot w^*(1)$. The welfare of the optimal allocation a^{opt} is $W(a^{\text{opt}}) = w^*(n)$. Therefore, we have that $\text{Eff}^*(\mathcal{W}; 1) \leq \text{Eff}(G; 1) = \min_{1 \leq \ell \leq m} \frac{n}{\bar{w}_\ell(n)}$ for any n and this is increasing in n so we have the claim.

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 $\sum_{i=1}^j u_\ell(i)/w_\ell(j) \leq 1$ for every j and ℓ , the one-round efficiency is lower bounded by $\text{Eff}_n(\mathcal{W}, \mathcal{U}; 1) \geq \min_{1 \leq \ell \leq m} \frac{n}{w_\ell(n)}$ for all n . To do this, we can use a modified version of the linear program in Eq. (3.13) for n agents, in which $\text{Eff}_n(\mathcal{W}, \mathcal{U}; 1) \geq \min_{1 \leq \ell \leq m} \frac{1}{\beta_\ell}$, where $\beta_\ell \in \mathbb{R}$ is the solution to

$\beta_\ell = \min \beta$ subject to:

$$\beta \bar{w}_\ell(y) \geq H \left(\sum_{i=1}^y \bar{u}_\ell(i) - z \min_{1 \leq i \leq y+1} \bar{u}_\ell(i) \right) + \bar{w}_\ell(z)$$

for all $0 \leq z \leq n$ and $1 \leq y \leq n$,

⁶We assume that the welfare set \mathcal{W} is finite for ease of presentation, and it is straightforward to extend to more general sets of welfare rules.

where the linear program is a lower bound since we consider tighter constraints that allow y and z to range from 1 to n . Since w_ℓ is supermodular, $H = \max_\ell \bar{w}_\ell(n)/n$ and assuming u_ℓ is non-decreasing, $\min_{1 \leq i \leq y+1} \bar{u}_\ell(i) = \bar{u}_\ell(1) = 1$. Thus, we can simplify the constraint as

$$\beta \geq \left(H \sum_{i=1}^y \bar{u}_\ell(i) - Hz + \bar{w}_\ell(z) \right) / \bar{w}_\ell(y) \tag{3.39}$$

With this, we observe that $\bar{w}_\ell(z) - Hz$ is convex in z . So the binding constraint for z occurs at either the end point $z = 0$ or $z = n$. Observe that $\max\{\bar{w}_\ell(n) - Hn, \bar{w}_\ell(0) - H0\} = 0$ and the terms can be cancelled out. Additionally, $\max_y \sum_{i=1}^y \bar{u}_\ell(i) / \bar{w}_\ell(y) = 1$ occurs at the binding constraint $y = 1$, by assumption that $\sum_{i=1}^j u_\ell(i) / w_\ell(j) \leq 1$ for all $1 \leq j \leq n$. Therefore, $\beta_\ell = H = \max_\ell \bar{w}_\ell(n)/n$ for all ℓ under the binding constraint of $y = 1$ and $z = 0$ and we indeed have that $\text{Eff}^*(\mathcal{W}; 1) \geq \min_{1 \leq \ell \leq m} \lim_{n \rightarrow \infty} \frac{n}{\bar{w}_\ell(n)}$.

Note that the marginal contribution utility rule (see proof of Theorem 3) satisfies the assumptions of optimality in Theorem 4 as $\sum_{i=1}^j u_\ell^{\text{mc}}(i) = w_\ell(j)$ for any j and u^{mc} is non-decreasing for supermodular welfare rules. Thus, the common interest utility design inherits the same efficiency guarantee $\text{Eff}(\mathcal{W}, \mathcal{U}_{\text{CI}}; 1) = \text{Eff}^*(\mathcal{W}; 1)$. □




Figure 3.3: We depict the game construction in the proof of Theorem 4. If every previous agent does not utilize r_1 in the execution of the non-oblivious algorithm, then the subsequent agents are also indifferent to their decisions and may also choose not to utilize r_1 resulting in a poor joint decision.

□

Remark 5. We can similarly define curvature for supermodular welfare rules, where $C = 1 - \lim_{n \rightarrow \infty} (w(n+1) - w(n))/w(1)$. Under this definition, note that the efficiency guarantees in Theorem 4 can be equally stated as $\text{Eff}^*(\mathcal{W}; 1) = (1 - C)^{-1}$.

Remark 6. We remark the efficiency guarantees $\text{Eff}^*(\mathcal{W}; 1)$ of the optimal non-oblivious algorithm exactly matches the guarantees of Nash-seeking algorithms with optimal sharing rules [43, Theorem 4]. Thus, greedy algorithms have similar guarantees to more complex algorithm designs.

We can consider the non-oblivious algorithm with a *Shapley* (or equal-shares) [18] utility function design for supermodular settings. Shapley utility functions are desirable due to their well known budget balance property [32], where $\sum_i U_i(a) = W(a)$ for all joint actions $a \in \mathcal{A}$. We observe that the Shapley utility rules (defined in this setting as $\mathcal{U}(w) = u_{\text{shap}}$ with $u_{\text{shap}}(j) = w(j)/j$ for all $j \in \mathbb{N}$) satisfies the assumptions of Theorem 4 and thus maximizes the possible efficiency guarantees. However, we note that the optimal utility rules are not unique, as the constant utility function design (defined as $\mathcal{U}(w) = u_1$ with $u_1(j) = w(1)$ for all $j \in \mathbb{N}$) also satisfies the assumptions of Theorem 4. Additionally, the standard greedy algorithm also has equivalent guarantees for supermodular welfare rules. However, the average case guarantees of different non-oblivious algorithms may be different; we leave it to future work to classify these algorithms based on their average behavior.

3.4 Finite Walks

We restrict attention to submodular resource allocation games in this chapter. We remark that the optimal efficiency guarantees after one round are relatively close to the

optimal polynomial time guarantees (see Figure 3.2) ⁷. However, while small, there is still a gap between the optimal efficiency guarantees for $k = 1$ and the best polynomial time guarantee. We expect this efficiency gap to decay as we run the k -round walk for more rounds. However, we show surprisingly that further rounds do not increase the relative efficiency guarantees. Specifically, running the k -round walk *can not* improve the resulting efficiency guarantee for any given k over the guarantee of the one-round walk. This is made formal in the upper bound characterization stated in the next theorem.

Theorem 5. *Let the set \mathcal{W} comprise of all submodular welfare rules w that have curvature of at most $C \in [0, 1]$. Then the efficiency guarantees of the k -round walk, for any $k \geq 1$, is upper bounded by the expressions*

$$\text{Eff}^*(\mathcal{W}; k) \leq \text{Eff}^*(\mathcal{W}; 1) = 1 - C/2, \quad (3.40)$$

$$\text{Eff}(\mathcal{W}, \mathcal{U}_{\text{CI}}; k) = \text{Eff}(\mathcal{W}, \mathcal{U}_{\text{CI}}; 1) = (1 + C)^{-1}, \quad (3.41)$$

respectively for the the optimal utility design and the common interest utility.

Proof. We first provide upper bounds on the efficiency metric $\text{Eff}^*(\mathcal{W}; k)$. To do this, we construct a game G such that for any utility design \mathcal{U} , rounds $k \geq 1$, and curvature C , we have that $\text{Eff}(\mathcal{W}, \mathcal{U}; k) \leq \text{Eff}(G; k) \leq 1 - C/2$. Let C be the curvature and consider the bent welfare rule $w^{b,C}$ with $b = 1$ as in Eq. (3.12) with $w^{b,C}(2) = 2 - C$. Additionally, let $u = \mathcal{U}(w^{b,C})$ be the corresponding utility rules 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

⁷Furthermore, when utilizing the common interest utility design, the efficiency guarantees are identical.

construction the agents will be determined by u according to the following three cases:

(a) $0 \leq u(2) \leq (1 - C)$, (b) $(1 - C) \leq u(2) \leq 1$, and (c) $u(2) \geq 1$.

For case (a), Agent 1's actions are $\mathcal{A}_1 = \{a_1^\emptyset, a_1^1 = \mathcal{R}_1, a_1^2 = \mathcal{R}_2\}$. Agent 2's actions are $\mathcal{A}_2 = \{a_2^\emptyset, a_2^1 = \mathcal{R}_3, a_2^2 = \mathcal{R}_1\}$. The optimal allocation is $a^{\text{opt}} = \{a_1^2, a_2^2\}$ resulting in a welfare of $2x$. An allocation that can occur after a one round walk is $a^{\text{br}} = \{a_1^1, a_2^1\}$ resulting in a welfare of $(1 + u(2))x$. Therefore, $\text{Eff}(\mathcal{G}; 1) \leq \frac{(1+u(2))x}{2x} \leq 1 - \frac{C}{2}$ by assumption of $u \leq 1 - C$. Additionally, observe that a^{br} is a Nash equilibrium and therefore is still the resulting allocation after any number of additional rounds $k \geq 1$. Therefore $\text{Eff}(\mathcal{W}, \mathcal{U}; k) \leq \text{Eff}(\mathcal{G}; k) \leq 1 - \frac{C}{2}$ for this case of utility design.

For case (b), Agent 1's actions are $\mathcal{A}_1 = \{a_1^\emptyset, a_1^1 = \mathcal{R}_1, a_1^2 = \mathcal{R}_2\}$. Agent 2's actions are $\mathcal{A}_2 = \{a_2^\emptyset, a_2^1 = \mathcal{R}_3, a_2^2 = \mathcal{R}_1\}$. The optimal allocation is $a^{\text{opt}} = \{a_1^2, a_2^2\}$ resulting in a welfare of $2x$. An allocation that can occur after a one-round walk is $a^{\text{br}} = \{a_1^1, a_2^2\}$ resulting in a welfare of $w^{b,C}(2) \cdot x$. Therefore, $\text{Eff}(\mathcal{G}; 1) \leq \frac{w^{b,C}(2) \cdot x}{2x} = 1 - \frac{C}{2}$. For $k \geq 2$, 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 remain at a' for $k - 1$ rounds and in the last round, switch to a^{br} . Therefore $\text{Eff}(\mathcal{W}, \mathcal{U}; k) \leq \text{Eff}(\mathcal{G}; k) \leq 1 - \frac{C}{2}$ for this case.

For case (c), Agent 1's actions are $\mathcal{A}_1 = \{a_1^\emptyset, a_1^1 = \mathcal{R}_1, a_1^2 = \mathcal{R}_2\}$. Agent 2's actions are $\mathcal{A}_2 = \{a_2^\emptyset, a_2^1 = \mathcal{R}_1, a_2^2 = \mathcal{R}_3\}$. The optimal allocation is $a^{\text{opt}} = \{a_1^2, a_2^2\}$ resulting in a welfare of $(1 + u(2))x$. An allocation that can occur after a one round walk is $a^{\text{br}} = \{a_1^1, a_2^1\}$ resulting in a welfare of $w^{b,C}(2) \cdot x$. Therefore, $\text{Eff}(\mathcal{G}; 1) = \frac{w^{b,C}(2) \cdot x}{(1+u(2))x} \leq 1 - \frac{C}{2}$ 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 $\text{Eff}(\mathcal{W}, \mathcal{U}; k) \leq \text{Eff}(\mathcal{G}; k) \leq 1 - \frac{C}{2}$ for this case.

Since $u = \mathcal{U}(w^{b,C})$ was chosen arbitrarily, we have that the upper bound holds for any utility design and we have shown that $\text{Eff}^*(\mathcal{W}; k) \leq 1 - C/2$. 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 .




Figure 3.4: The worst case game construction achieving the k -round walk guarantee dictated in Equation (3.40). P_1 and P_2 represent the two agents and $|x|$ and $|x \cdot u(2)|$ represent the size of the resources. The black lines represent the selections in the different joint actions by the two agents.

Common Interest Utility First, we note that since any best response improves the welfare under the common interest utility, we must have that $\text{Eff}(\mathcal{W}, \mathcal{U}_{\text{CI}}; k) \geq \text{Eff}(\mathcal{W}, \mathcal{U}_{\text{CI}}; 1) = (1+C)^{-1}$. Now we show that the upper bound $\text{Eff}(\mathcal{W}, \mathcal{U}_{\text{CI}}; k) \leq (1+C)^{-1}$ to complete the equality in Eq. (3.41). As before, a game G is constructed such that under the common interest design \mathcal{U}_{CI} , $k \geq 1$, and curvature C , we have that $\text{Eff}(\mathcal{W}, \mathcal{U}_{\text{CI}}; k) \leq \text{Eff}(G; k) \leq (1+C)^{-1}$. Let G have n players with a resource set $\mathcal{R} = \mathcal{R}^{\text{opt}} \cup \mathcal{R}^{\text{both}} \cup \{r^n\}$ with $|\mathcal{R}^{\text{opt}}| = n$ and $|\mathcal{R}^{\text{both}}| = n - 1$. Each agent i has three actions in its action set $\mathcal{A}_i = \{a_i^{\emptyset}, a_i^{\text{br}}, a_i^{\text{opt}}\}$. The resources are selected by the agents in the following manner: each resource $r_j^{\text{opt}} \in \mathcal{R}^{\text{opt}}$ is selected by agent j in action

$a_j^{\text{opt}} \ni r_j^{\text{opt}}$ for all $1 \leq j \leq n$; each resource $r_j^{\text{both}} \in \mathcal{R}^{\text{both}}$ is selected by agent $j + 1$ in action $a_{j+1}^{\text{opt}} \ni r_j^{\text{both}}$ and by agent j in action $a_j^{\text{br}} \ni r_j^{\text{both}}$ for all $1 \leq j \leq n - 1$; agent n selects the resource r^n in action a_n^{br} . See Figure 3.5 for a visual representation of this game.

Given a curvature C , consider two bent welfare rules $w_1, w_2 \in \mathcal{W}$ with curvature C such that $w_1 = w^{b,C}$ with $b = 1$ and $w_2 = C \cdot w_1$. For any $r \in \mathcal{R}^{\text{both}} \cup \{r^n\}$, let the corresponding welfare rule be $w_r = w_1$ and for any $r \in \mathcal{R}^{\text{opt}}$, let the corresponding welfare rule be $w_r = w_2$. Under this game construction it can be seen that under a^{br} , each resource $r \in \mathcal{R}^{\text{both}} \cup \{r^n\}$ is selected by exactly one agent, resulting in a welfare of $W(a^{\text{br}}) = n$; also, under a^{opt} , each resource $r \in \mathcal{R}^{\text{both}} \cup \mathcal{R}^{\text{opt}}$ is selected by exactly one agent, resulting in a welfare of $W(a^{\text{opt}}) = (n - 1)(1 + C) + C$. Assuming that a^{br} is the joint action that results after k rounds, we have that $\text{Eff}(G, k) \leq \frac{n}{(n-1)(1+C)+C}$. Limiting the number of agents $n \rightarrow \infty$ to infinity gives the result. To verify that a^{br} can result after k 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 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} is a Nash equilibrium, so after any number of rounds k , the joint action a^{br} is still the result of a k -round walk.

□

Remark 7. We remark that the results in Theorem 5 are not endemic to the k -round walk algorithm. Allowing for variable turn order or only allowing strict best responses does not affect the resulting upper bounds.

Notably, for any curvature $C \in [0, 1]$, the upper bound in Eq. (3.40) exactly matches the characterization in Eq. (3.29) and likewise for the upper bound in Eq. (3.41) and the characterization in Eq. (3.30). Therefore, in regards to the efficiency guarantees,




Figure 3.5: The worst case game construction achieving the k -round walk guarantee dictated in Equation (3.41). In this figure, rows represent players and columns represent resources. Red circles represent selections in a^{br} and blue circles represent selections in a^{opt} .

running the k -round walk algorithm for more than one round does not lead to gains in performance.

3.5 Asymptotic Walks

With this, we can characterize the optimal efficiency guarantees for the k -round walk when $k = \infty$ as well as the efficiency guarantees for the common-interest utility below. We still restrict attention to submodular resource allocation games. We see that the asymptotic guarantees match the best polynomial time guarantee in Figure 3.2.

Proposition 1. *Let the set \mathcal{W} comprise of all submodular welfare rules w that have curvature of at most $C \in [0, 1]$. The optimal efficiency guarantees are*

$$\text{Eff}^*(\mathcal{W}; \infty) = 1 - C/e. \tag{3.42}$$

and the guarantees associated with the common interest are

$$\text{Eff}(\mathcal{W}, \mathcal{U}_{\text{CI}}; \infty) = (1 + C)^{-1}. \quad (3.43)$$

Proof. Note that resource allocation games are isomorphic to *potential games*, and, as such, the limit points from any best response process must necessarily be in the set of Nash equilibrium $\text{NE} \subseteq \mathcal{A}$ of the game. Thus the limit points $\lim_{k \rightarrow \infty} a^k$ of the k -round walk are Nash equilibrium as well. We consider $a^{\text{ne}} \in \text{NE}$ to be a Nash equilibrium if any unilateral deviations are not preferable by any agent, or

$$a_i^{\text{ne}} \in \arg \max_{a_i \in \mathcal{A}_i} U_i(a_i, a_{-i}^{\text{ne}}) \quad \text{for all } i \in \mathcal{I}. \quad (3.44)$$

Measuring the quality of Nash equilibrium is done through the classical metric of *price of anarchy* as follows.

$$\text{PoA}(G) = \frac{\min_{a \in \text{NE}} W(a)}{\max_{a \in \mathcal{A}} W(a)}. \quad (3.45)$$

We similarly define $\text{PoA}(\mathcal{W}, \mathcal{U}) = \inf_{G \in \mathcal{G}_{\mathcal{W}, \mathcal{U}}} \text{PoA}(G)$ mirroring Eq. (3.10). The price of anarchy is a well understood metric, with a host of results on its characterization, complexity, and design [44]. As solutions of the ∞ -round walk must also be Nash equilibrium, we have that $\lim_{k \rightarrow \infty} a^k \subseteq \text{NE}$. However, this inclusion may be strict, as not every Nash equilibrium may be reachable from the k -round walk considered in Algorithm 1. But in the next theorem, we show equivalence of price of anarchy and the efficiency of the ∞ -round walk. Thus in the subsequent sections, we can use previous results in the literature on price of anarchy to quantify the efficiency of the ∞ -round walk.

Lemma 2. *Let \mathcal{W} be a set of welfare rules. The efficiency of the ∞ -round walk is*

equivalent to the price of anarchy

$$\text{Eff}(\mathcal{W}, \mathcal{U}; \infty) = \text{PoA}(\mathcal{W}, \mathcal{U}), \quad (3.46)$$

if the utility rules $u_\ell = \mathcal{U}(w_\ell)$ for $1 \leq \ell \leq m$ are non-increasing, and $\text{Eff}(\mathcal{W}, \mathcal{U}; \infty) \geq \text{PoA}(\mathcal{W}, \mathcal{U})$ otherwise.

Proof. For ease of notation, we remove the subscript of u_ℓ^{mc} as u_ℓ . Since $\lim_{k \rightarrow \infty} a^k \subseteq \text{NE}$, by definition, the efficiency guarantee of $\text{Eff}(\mathcal{W}, \mathcal{U}; \infty) \geq \text{PoA}(\mathcal{W}, \mathcal{U})$ must be higher than the guarantee for the total set of Nash equilibrium.

Now we show the $\text{Eff}(\mathcal{W}, \mathcal{U}; \infty) \leq \text{PoA}(\mathcal{W}, \mathcal{U})$ by a game construction G , in which a Nash equilibrium with the efficiency arbitrarily close to $\text{PoA}(\mathcal{W}, \mathcal{U})$ can result from a one-round walk. Let $\varepsilon_1 > 0$ and $\text{PoA}^n(\mathcal{W}, \mathcal{U})$ refer to the price of anarchy for the set of games in $\mathcal{G}_{\mathcal{W}, \mathcal{U}}$ that have only n number of agents. Note that $\text{PoA}^n(\mathcal{W}, \mathcal{U})$ is non-increasing in n and lower bounded by 0. Therefore $\text{PoA}^n(\mathcal{W}, \mathcal{U})$ is a convergent sequence in n and for any ε_1 , there exists an $N_1 \in \mathbb{N}$ such that $\text{PoA}^{N_1}(\mathcal{W}, \mathcal{U}) - \text{PoA}(\mathcal{W}, \mathcal{U}) \leq \varepsilon_1$.

Generalizing [39, Theorem 2] to a set of welfare rules provides a characterization of the price of anarchy for N_1 agents as $\text{PoA}^{N_1}(\mathcal{W}, \mathcal{U}) = Q^{-1}$ with

$$\begin{aligned} Q &= \max_{\theta(y, x, z, \ell)} \sum_{\substack{1 \leq \ell \leq m, \\ y, x, z}} w_\ell(z + x) \theta(y, x, z, \ell) & (3.47) \\ \text{s.t.} \quad & \sum_{\substack{1 \leq \ell \leq m, \\ y, x, z}} [y u_\ell(y + x) - z u_\ell(y + x + 1)] \theta(y, x, z, \ell) \geq 0 \\ & \sum_{\substack{1 \leq \ell \leq m, \\ y, x, z}} w_\ell(y + x) \theta(y, x, z, \ell) = 1 \\ & \theta(y, x, z, \ell) \geq 0, \end{aligned}$$

where $y, x, z \in \mathbb{N}$ with $1 \leq y + x + z \leq N_1$. We refer to $\Theta(y, x, z, \ell)$ to denote the corresponding optimal variables for $\theta(y, x, z, \ell)$ 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, ℓ pair and $1 \leq k \leq D$, we construct a set of resources $\mathcal{R}_{y,x,z,\ell}^k$ with $|\mathcal{R}_{y,x,z,\ell}^k| = \Theta(y, x, z, \ell)/D$ ⁸. Each agent i has three actions in its action set $\mathcal{A}_i = \{a_i^\emptyset, a_i^{\text{ne}}, a_i^{\text{opt}}\}$. Each agent i selects $\{\mathcal{R}_{y,x,z,\ell}^k\}_{i \leq k \leq y+x+i-1}$ in a_i^{ne} for each pair y, z, x, ℓ . If $y + z + x \leq i \leq N_2$, agent i selects $\{\mathcal{R}_{y,x,z,\ell}^k\}_{i-z \leq k \leq x+i-1}$ in a_i^{opt} for each pair y, z, x, ℓ . Otherwise for $1 \leq i \leq y + z + x - 1$, $a_i^{\text{opt}} = a_i^\emptyset$ and agent i doesn't select any resources in a_i^{opt} . This is shown in Figure 3.6 for one y, x, z, ℓ . pair




Figure 3.6: In this figure, rows represent players and columns represent resources. Red circles represent selections in a^{ne} and blue circles represent selections in a^{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

⁸While $\Theta(y, x, z, \ell)/D$ might not be an integer, we can scale $|\mathcal{R}_{y,x,z,\ell}^k|$ uniformly and round to arrive at a game construction with the arbitrarily close efficiency guarantees.

rest of the agents, the utility difference of a unilateral deviation to a_i^{opt} from a_i^{ne} is

$$\begin{aligned}
& U_i(a^{\text{ne}}) - U_i(a_i^{\text{opt}}, a_{-i}^{\text{ne}}) \\
& \geq \sum_{r \in a_i^{\text{ne}}} u_r(|a^{\text{ne}}|_r) - \sum_{r \in a_i^{\text{opt}}} u_r(|(a_i^{\text{opt}}, a_{-i}^{\text{ne}})|_r) \\
& \geq \sum_{y,x,z,\ell} [(y+x)u_\ell(a+x) - \\
& \quad xu_\ell(y+x) - zu_\ell(y+x+1)] \cdot |\mathcal{R}_{y,x,z,\ell}^k| \\
& \geq \frac{1}{D} \sum_{y,x,z,\ell} [yu_\ell(y+x) - zu_\ell(y+x+1)] \Theta(y,x,z,\ell) \\
& \geq 0.
\end{aligned}$$

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 $\mathcal{R}_{y,x,z,\ell}^k$. The third inequality comes from simplifying. The fourth inequality comes from the fact that since $\Theta(y,x,z,\ell)$ has to satisfy the inequality constraint in Eq. (3.47) to be feasible. Similarly, in a one-round walk, the best response for the first $y+x+z-1$ agents is a_i^{ne} . The best response for the other agents during the one-round walk is also a_i^{ne} , since

$$\begin{aligned}
& U_i(a_{j<i}^{\text{ne}}, a_i^{\text{ne}}, a_{j>i}^{\emptyset}) - U_i(a_{j<i}^{\text{ne}}, a_i^{\text{opt}}, a_{j>i}^{\emptyset}) \\
& = \sum_{y,x,z,\ell} [\sum_{j=1}^{y+x} u_\ell(i) - xu_\ell(y+x) - zu_\ell(y+x+1)] |\mathcal{R}_{y,x,z,\ell}^k| \\
& \geq \frac{1}{D} \sum_{y,x,z,\ell} [yu_\ell(y+x) - zu_\ell(y+x+1)] \Theta(y,x,z,\ell) \\
& \geq 0.
\end{aligned}$$

We use similar arguments as before, where the second inequality comes from the

fact that u_ℓ is non-increasing. 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 any k rounds can also be itself a^{ne} . Therefore in this game, $\text{Eff}(\mathcal{G}; \infty) \leq \text{PoA}(\mathcal{G})$. Now we calculate the efficiency of the Nash equilibrium $W(a^{\text{ne}})$ with respect to $W(a^{\text{opt}})$. We have that

$$\begin{aligned} W(a^{\text{ne}}) &= \sum_{y,x,z,\ell} w_\ell(y+x) \cdot \Theta(y,x,z,\ell) \frac{N_2 - 2(y+x-1)}{N_2} \\ &\quad + 2 \sum_{\substack{1 \leq \ell \leq m \\ 1 \leq i \leq y+x-1}} w_\ell(i) \frac{\Theta(y,x,z,\ell)}{N_2} = 1 + O\left(\frac{1}{N_2}\right), \end{aligned}$$

where, since $\Theta(y,x,z,\ell)$ is feasible, then it satisfies the equality constraint that $\sum_{y,x,z,\ell} w_\ell(y+x)\Theta(y,x,z,\ell) = 1$. $O(\frac{1}{N_2})$ reflects that the rest of the terms are on order of $1/N_2$.

Similarly,

$$\begin{aligned} W(a^{\text{opt}}) &= \sum_{y,x,z,\ell} w_\ell(z+x) \cdot \Theta(y,x,z,\ell) \frac{N_2 - 3(z+x-1)}{N_2} \\ &\quad + 2 \sum_{\substack{1 \leq \ell \leq m \\ 1 \leq i \leq z+x-1}} w_\ell(i) \frac{\Theta(y,x,z,\ell)}{N_2} = Q + O\left(\frac{1}{N_2}\right), \end{aligned}$$

where, since $\Theta(y,x,z,\ell)$ is optimal, then $\sum_{y,x,z,\ell} w_\ell(z+x)\Theta(y,x,z,\ell) = Q = \text{PoA}^{N_1}(\mathcal{W}, \mathcal{U})^{-1}$.

For any ε_2 , we can choose N_2 , such that $O(\frac{1}{N_2}) \leq \varepsilon_2$, so $\text{PoA}(\mathcal{G}) \leq \text{PoA}^{N_1}(\mathcal{W}, \mathcal{U})^{-1} + \varepsilon_2$.

To put everything together, we have that

$$\begin{aligned} \text{Eff}(\mathcal{W}, \mathcal{U}; \infty) &\leq \text{Eff}(\mathcal{G}; \infty) = \text{PoA}(\mathcal{G}) \\ &\leq \text{PoA}^{N_1}(\mathcal{W}, \mathcal{U})^{-1} + \varepsilon_2 \leq \text{PoA}(\mathcal{W}, \mathcal{U}) + \varepsilon_1 + \varepsilon_2, \end{aligned}$$

and since ε_1 and ε_2 are arbitrary, we have the result. \square

Remark 8. A simple corollary from the above proof also shows that $\text{Eff}(\mathcal{W}, \mathcal{U}; k) \leq \text{PoA}(\mathcal{W}, \mathcal{U})$ for any $k \geq 1$ as well. Thus the efficiency guarantees of k -round walks are upper bounded by the price of anarchy.

The equality $\text{Eff}^*(\mathcal{W}; \infty) = 1 - C/e$ in Eq. (3.42) comes from the fact that $\text{Eff}^*(\mathcal{W}; \infty) = \sup_{\mathcal{U}} \text{PoA}(\mathcal{W}, \mathcal{U})$ by Lemma 2 for the set of non-increasing utility rules and that $\sup_{\mathcal{U}} \text{PoA}(\mathcal{W}, \mathcal{U}) = 1 - C/e$ comes from [42, Theorem 1].

Now we show $\text{Eff}(\mathcal{W}, \mathcal{U}_{\text{CI}}; \infty) = (1 + C)^{-1}$. Since any best response with a common interest utility must increase the welfare W , the limiting efficiency $\text{Eff}(\mathcal{W}, \mathcal{U}_{\text{CI}}; \infty) \geq \text{Eff}(\mathcal{W}, \mathcal{U}_{\text{CI}}; 1) = 1/(1 + C)$ is greater than the efficiency of the one-round walk. Since we consider welfare rules w_r that are submodular, then the utility rules u_r^{mc} are non-increasing, and we can apply Lemma 2 to have that $\text{Eff}(\mathcal{W}, \mathcal{U}_{\text{CI}}; \infty) = \text{PoA}(\mathcal{W}, \mathcal{U}_{\text{CI}})$. From applying [43, Corollary 1] with the bent welfare rule $w^{b,C}$ in Eq. (3.12) gives

$$\begin{aligned} \text{PoA}(\mathcal{W}, \mathcal{U}_{\text{CI}})^{-1} &\geq \text{PoA}(w^{b,C}, \mathcal{U}_{\text{CI}})^{-1} = \\ &1 + \max_{j \geq 1} \left\{ \frac{j}{w^{b,C}(j)} [2w^{b,C}(j) - w^{b,C}(j-1) - w^{b,C}(j+1)] \right\} \end{aligned}$$

Simplifying the inequality for $b = 1$ gives $\text{PoA}(\mathcal{W}, \mathcal{U}_{\text{CI}}) \leq (1 + C)^{-1}$. Since $\text{Eff}(\mathcal{W}, \mathcal{U}_{\text{CI}}; \infty)$ is both upper bounded and lower bounded by $(1 + C)^{-1}$, we have the result. \square

3.6 Tradeoffs

In Proposition 1 and Theorem 3, we describe the optimal performance guarantees for $k = 1$ and $k = \infty$. Accompanying these results are characterizations of the utility designs that achieve said guarantees. As these utility designs are not equivalent, this prompts the natural question: Does optimizing the efficiency guarantees for $k = 1$ have downstream effects on the efficiency for $k = \infty$, and vice versa? This question is precisely

addressed in the next theorem. Explicitly, we identify the reciprocal guarantees for both the utility design that optimizes the transient guarantees for $k = 1$ and the utility design that optimizes the asymptotic guarantees for $k = \infty$.

Theorem 6. *Let the set \mathcal{W} comprise of all submodular welfare rules w that have curvature of at most $C \in [0, 1]$. Consider \mathcal{U}_∞^* to be the utility design that achieves the optimal $\text{Eff}^*(\mathcal{W}; \infty)$ and \mathcal{U}_1^* to be the utility design that achieves the optimal $\text{Eff}^*(\mathcal{W}; 1)$. Then we have that*

$$\text{Eff}(\mathcal{W}, \mathcal{U}_1^*; \infty) = \text{Eff}^*(\mathcal{W}; 1) = 1 - C/2 \quad (3.48)$$

$$\text{Eff}(\mathcal{W}, \mathcal{U}_\infty^*; 1) \leq 1 + \frac{(C - 3)C}{(2 - C)e + C} \quad (3.49)$$

with $\text{Eff}(\mathcal{W}, \mathcal{U}_\infty^*; 1) < \text{Eff}^*(\mathcal{W}; 1)$ holding strictly for $C > 0$ and $\text{Eff}(\mathcal{W}, \mathcal{U}_\infty^*; 1) = 0$ for curvature $C = 1$.

Proof. We show the trade-offs in Theorem 6 that result from considering utility designs that maximize the one-round walk efficiency versus the ∞ -round walk. We first show the equality in Eq. (3.48). From Lemma 2, we have

$$\text{Eff}(\mathcal{W}, \mathcal{U}_1^*; \infty) = \text{PoA}(\mathcal{W}, \mathcal{U}_1^*) \geq \text{Eff}(\mathcal{W}, \mathcal{U}_1^*; 1) = 1 - C/2,$$

since \mathcal{U}_1^* is a non-increasing utility design, which is shown in Theorem 2. We now show $\text{PoA}(\mathcal{W}, \mathcal{U}_1^*) \leq 1 - C/2$. Consider the bent rule $w^{b,C}$ in Eq. (3.12) with $b = 1$ and a utility rule u_1 with $u_1(1) = 1$ and $u_1(2) = (2 - 2C)/(2 - C)$. In Theorem 3, it was shown that u_1 is the utility rule that maximizes the one-round efficiency for $\mathcal{W} = \{w^{b,C}\}$. It can be easily verified that u_1 satisfies the assumptions of [43, Theorem 2] and we can derive

the price of anarchy as

$$\text{PoA}(w^{b,C}, u_1)^{-1} = \max_{1 \leq l \leq j} \left\{ \frac{w^{b,C(l)+ju_1(j)-lu_1(j+1)}}{w^{b,C(j)}} \right\}.$$

Under $j = 1$ and $l = 1$, we have that $\text{PoA}(w^{b,C}, u_1) \leq 1 - C/2$. Using the fact that $\text{PoA}(\mathcal{W}, \mathcal{U}_1^*) \leq \text{PoA}(w^{b,C}, u_1)$, we have the upper bound as well.

We show that $\text{Eff}(\mathcal{W}, \mathcal{U}_\infty; 1) \leq 1 + \frac{(C-3)C}{(2-C)e+C}$ in Lemma 3 and $\text{Eff}(\mathcal{W}, \mathcal{U}_\infty; 1) = 0$ for $C = 1$ in Lemma 4. These lemmas are stated below.

Lemma 3. *Let the set \mathcal{W} comprise of all submodular welfare rules w that have curvature of at most $C \in [0, 1]$. Consider \mathcal{U}_∞^* to be the utility design that achieves the optimal $\text{Eff}^*(\mathcal{W}; \infty)$. Then we have that*

$$\text{Eff}(\mathcal{W}, \mathcal{U}_\infty^*; 1) \leq 1 + \frac{(C-3)C}{(2-C)e+C}. \quad (3.50)$$

Proof. From results in [42, Lemma 1 iii], the utility design that optimizes the price of anarchy $\mathcal{U}_{\text{poa}} \doteq \arg \max_{\mathcal{U}} \text{PoA}(\mathcal{W}, \mathcal{U})$ have non-increasing utility rules. Since these utility rules are non-increasing, under Lemma 1, the resulting asymptotic efficiency $\text{Eff}(\mathcal{W}, \mathcal{U}_{\text{poa}}, \infty) = \text{PoA}(\mathcal{W}, \mathcal{U}_{\text{poa}})$ and thus $\mathcal{U}_{\text{poa}} \equiv \mathcal{U}_\infty^*$ is the utility design that maximizes the asymptotic efficiency.

To construct the upper bound in Eq. (3.50), we characterize the one-round efficiency of the asymptotically optimal utility design against the bent welfare rule $w^{b,C}$ for $b = 1$. We have that $\text{Eff}(\mathcal{W}, \mathcal{U}_\infty^*; 1) \leq \text{Eff}(w^{1,C}, u^\infty; 1)$, where $u^\infty \equiv \mathcal{U}_\infty^*(w^{1,C})$. We can use the linear program given in Thm 1 to characterize the one-round efficiency. Under only the bent welfare rule, the constraint for a given z, y simplifies to

$$\beta \cdot w^{1,C}(y) \geq \sum_{i=1}^y u^\infty(i) - z \min_{1 \leq i \leq y+1} u^\infty(i) + w^{1,C}(z).$$

According to [42, Lemma 1], the asymptotically optimal utility rule u^∞ for the bent welfare rule is given by the following recursive equation

$$\begin{aligned} u^\infty(1) &= 1, \\ u^\infty(j+1) &= \max\{ju^\infty(j) - \rho w^{1,C}(j) + 1, 1 - C\}, \end{aligned}$$

with $\rho = (1 - C/e)^{-1}$. Now we characterize the binding constraints for z . Since $u^\infty(j) \leq 1$ for all j and $u^\infty(j) \geq w^{1,C}(j) - w^{1,C}(j-1)$ for $j \geq 2$, the binding constraint is when $z = 1$. Coupled with the fact that u^∞ is non-increasing, we can simplify the one-round characterization as

$$\begin{aligned} \text{Eff}(w^{1,C}, u^\infty; 1)^{-1} &= \max_{y \geq 1} \left\{ \frac{1 + \sum_{i=1}^y u^\infty(i) - u^\infty(y+1)}{w^{1,C}(y)} \right\} \\ &\geq \frac{(2-C)e + C}{(2-C)(e-C)}, \end{aligned}$$

where the second inequality comes from only considering $y = 2$ and solving for $u^\infty(j)$ explicitly for $j = 2$ and $j = 3$. Taking the reciprocal on both sides and simplifying gives the expression in Eq. (3.50). \square

Lemma 4. *Let the set \mathcal{W} comprise of all submodular welfare rules w that have curvature of at most $C = 1$. Consider \mathcal{U}_∞^* to be the utility design that achieves the optimal $\text{Eff}^*(\mathcal{W}; \infty)$. Then $\text{Eff}(\mathcal{W}, \mathcal{U}_\infty^*; 1) = 0$.*

Proof. By definition, $\text{Eff}(\mathcal{W}, \mathcal{U}_\infty^*; 1) \geq 0$ must be greater than zero. For the upper bound, we construct a game $G \in \mathcal{G}_{\mathcal{W}, \mathcal{U}_\infty^*}$ such that one-round walk efficiency is $\text{Eff}(G; 1) = 0$. Consider a game G 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 $w_r = w^{b,C}$ as the bent welfare rule with curvature of $C = 1$ for some fixed $b \geq 1$, as defined in Eq.

The corresponding utility rule is $u^\infty = \mathcal{U}_\infty^*(w^{b,C})$ is the following recursive expression from [42, Lemma 1]⁹,

$$\begin{aligned} u^\infty(1) &= 1 \\ u^\infty(j+1) &= \frac{1}{b}[ju^\infty(j) - \rho^b \min\{j, b\}] + 1, \end{aligned}$$

with $\rho^b = (1 - \frac{b^b e^{-b}}{b!})^{-1}$. The number of resources in each set is $|\mathcal{R}_1| = v$ and $|\mathcal{R}_{j+1}| \sim v \cdot u^\infty(j)$ for $1 \leq j \leq n$ and for some $v \geq 0$. If $u^\infty(j)$ 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 $\mathcal{R}_1 = a_i^{\text{br}}$ and $\mathcal{R}_{i+1} = a_i^{\text{opt}}$ 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

$$\text{Eff}(\text{G}; 1) \leq \frac{W(a^{\text{br}})}{W(a^{\text{opt}})} = \frac{vb}{v \sum_{1 \leq i \leq n} u^\infty(i)}.$$

Now we show that as we increase n , the series $\sum_{1 \leq i \leq n} u^\infty(i)$ diverges, and the efficiency can get arbitrarily bad as the number of agents increase. To construct the closed form expression of $u^\infty(j)$, we construct the following LTV state space system with $u^\infty(j) := x(t)$

$$\begin{aligned} x(t+1) &= A(t)x(t) + s(t) & A(t) &= \frac{t}{b} \\ s(t) &= 1 - \frac{\rho^b}{b} \min(t, b) \end{aligned}$$

Solving for the solution $x(t)$ using the state transition matrix with the initial condition

⁹The recursive expression found in [42] is for the utility rule that maximizes the price of anarchy. Since this utility rule is non-increasing, we can apply Lemma 2 to translate the results to ∞ -round walks.

$x(1) = 1$ results in the following expression

$$\begin{aligned} x(t) &= \prod_{\tau=1}^t \frac{\tau}{b} + \sum_{T=1}^{t-1} \left[\left(1 - \frac{\rho^b}{b} \min(t, b)\right) \prod_{\tau=T+1}^{t-1} \frac{\tau}{b} \right] \\ &= \frac{t!}{b^t} \left(1 + \sum_{T=1}^t \frac{b^T}{T!} \left(1 - \frac{\rho^b}{b} \min(t, b)\right) \right) \end{aligned}$$

If $t \geq b$, then

$$\begin{aligned} x(t) &= \frac{t!}{b^t} \left(1 - (e^b - 1)(\rho - 1) + \sum_{T=t+1}^{\infty} \frac{b^T}{T!} (\rho^b - 1) + \right. \\ &\quad \left. \sum_{T=1}^b \frac{b^T}{T!} \frac{\rho^b(b-T)}{b} \right) \\ &= \frac{t!}{b^t} \sum_{T=t+1}^{\infty} \frac{b^T}{T!} (\rho^b - 1) \\ &\geq (\rho^b - 1) \frac{b}{t+1} \\ &\sim O\left(\frac{1}{t}\right) \end{aligned}$$

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 $\sum_{i=1}^N u^\infty(i)$ diverges and the claim is shown. Now we verify the equality

$$\begin{aligned} \sum_{T=1}^b \frac{b^T}{T!} \frac{\rho^b(b-T)}{b} &= (e^b - 1)(\rho^b - 1) - 1 \\ \sum_{T=1}^b \frac{b^T(b-T)}{bT!} &= \frac{1}{\rho^b} (e^b \rho^b - e^b - \rho^b) \\ \sum_{T=1}^b \frac{b^T}{T!} - \sum_{T=1}^b \frac{b^{T-1}}{(T-1)!} &= (e^b - 1 - e^b(1 - \frac{b^b e^{-b}}{b!})) \\ \frac{b^b}{b!} - 1 &= \frac{b^b}{b!} - 1 \end{aligned}$$

The last equality results from recognizing the terms on the left hand side as a telescoping sum. \square

From the given Lemmas, we have shown the sum claims in Theorem 6. \square

We first observe that there are no gains in the asymptotic guarantees of the utility design \mathcal{U}_1^* over the respective optimal guarantees for $k = 1$. Again, we see that there are diminishing returns of running the k -round algorithm for more than one round. Additionally, the transient guarantees of \mathcal{U}_∞^* are strictly less than $\text{Eff}^*(\mathcal{W}; 1)$ as expected. However, if the curvature $C = 1$ is maximal, we note that transient guarantees of \mathcal{U}_∞^* unexpectedly degrade to 0. Interestingly, optimizing for asymptotic performance does not necessarily translate to good transient performance in our setting. Moreover, it may even result in highly undesirable behavior in the transient in certain settings.

To clarify this stark trade-off between the transient and asymptotic guarantees, we restrict attention to the class of *set covering games* [35] (see Example 1.A) and characterize the exact Pareto optimal frontier. Set covering games are natural generalizations of covering problems [35], and are characterized by the following welfare rule (with curvature $C = 1$).

$$w_{\text{sc}}(j) = \begin{cases} 1, & \text{for } j \geq 1 \\ 0, & \text{for } j = 0 \end{cases}. \quad (3.51)$$

With this, we arrive at the following Pareto frontier characterization, depicted in Figure 3.8. Note that the end points of the trade-off curve matches the ones dictated in Theorem 6 for curvature $C = 1$ exactly.

Theorem 7. *Let $\mathcal{W} = \{w_{\text{sc}}\}$, where w_{sc} , defined in Eq. (3.51), is the set covering welfare rule and $\mathcal{U}(w_{\text{sc}}) = u$ is the corresponding distribution rule. If the efficiency of the limit point of the k -round walk is $\text{Eff}(w_{\text{sc}}, u; \infty) = Q \in [\frac{1}{2}, 1 - \frac{1}{e}]$ ¹⁰, the maximum efficiency*

¹⁰We will sometimes use abuse of notation $\text{Eff}(w_r, u; k)$ to mean $\text{Eff}(\mathcal{W} = \{w_r\}, \mathcal{U}; k)$ with $\mathcal{U}(w_r) = u$

$\max_u \text{Eff}(w_{\text{sc}}, u; 1)$ achievable after $k = 1$ rounds is

$$\left[\sum_{j=0}^{\infty} \max \left\{ j! \left(1 - \frac{1-Q}{Q} \sum_{\tau=1}^j \frac{1}{\tau!} \right), 0 \right\} + 1 \right]^{-1}. \quad (3.52)$$

Proof. To characterize the Pareto optimal frontier in Eq. (3.52), we first derive the closed form solution for the one-round walk efficiency $\text{Eff}(w_{\text{sc}}, u; 1)$ specifically for the set covering welfare rule w_{sc} .

Lemma 5. *Let $\mathcal{W} = \{w_{\text{sc}}\}$, where w_{sc} is the set covering welfare rule defined in Eq. (3.51), and $\mathcal{U}(w_{\text{sc}}) = u$ be the corresponding utility rule. Then the one-round walk efficiency guarantee is*

$$\text{Eff}(w_{\text{sc}}, u; 1) = \left[\sum_{i \in \mathbb{N}} u(i) - \min_{i \in \mathbb{N}} u(i) + 1 \right]^{-1}. \quad (3.53)$$

Proof. Examine the linear program in Corollary 1 with substituting the set covering welfare defined in Eq. (3.51). Under the substitution, the 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 $w_{\text{sc}}(j) = \min(1, j) = 1$ when $j \geq 1$. Observe that the binding constraint occurs when we limit $y \rightarrow \infty$ and set $z = 1$ (and not $z = 0$ since $u(1) = 1$, the term $1 - \min_j u(j) \geq 0$). Under the binding constraint, $\text{Eff}(w_{\text{sc}}, u; 1) = \beta^{-1}$, where β^{-1} matches the given expression in Eq. (3.53). \square

To describe the trade-off, we now provide an explicit expression of Pareto optimal utility rules, i.e., the utility rules u that satisfy either $\text{Eff}(w_{\text{sc}}, u; 1) \geq \text{Eff}(w_{\text{sc}}, u'; 1)$ or $\text{Eff}(w_{\text{sc}}, u; \infty) \geq \text{Eff}(w_{\text{sc}}, u'; \infty)$ for all $u' \neq u$.

for a specific w_r and u_r .




Figure 3.7: The worst case game construction achieving the one-round walk guarantee dictated by Lemma 5. The agents are represented by circles, each oval represents the number of resources with the set covering welfare, and the black lines represent the agent action selections. In this game, all the agents can either stack on the first resource set or spread out.

Lemma 6. For a given $\mathcal{X} \geq 0$, a utility rule $u^{\mathcal{X}}$ that satisfies $\text{Eff}(w_{\text{sc}}, u; \infty) \geq 1/(1 + \mathcal{X})$ while maximizing $\text{Eff}(w_{\text{sc}}, u; 1)$ is defined as in the following recursive formula:

$$\begin{aligned} u^{\mathcal{X}}(1) &= 1 \\ u^{\mathcal{X}}(j+1) &= \max\{ju^{\mathcal{X}}(j) - \mathcal{X}, 0\}. \end{aligned} \tag{3.54}$$

Proof. We only consider utility rules that are non-increasing. Under this assumption, from Lemma 2, we have that $\text{Eff}(w_{\text{sc}}, u; \infty) = \text{PoA}(w_{\text{sc}}, u)$. According to Corollary 2 in [39], the price of anarchy for n agents can be written as

$$\frac{1}{\text{PoA}^n(w_{\text{sc}}, u)} = 1 + \max_{1 \leq j \leq n-1} \{ju(j) - u(j+1), (n-1)u(n)\}.$$

We define the following constant for each possible utility rule.

$$\mathcal{X}_u = \max_{1 \leq j \leq n-1} \{ju(j) - u(j+1), (n-1)u(n)\}. \tag{3.55}$$

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 $\text{PoA}^n(w_{\text{sc}}, u) = \text{PoA}^n(w_{\text{sc}}, u') = 1/(1 + \mathcal{X}_u)$. By induction, we show that $u(j) \leq u'(j)$ for all j . The base case is satisfied,

as $1 = u(1) \leq u'(1) = 1$. Under the assumption $u(j) \leq u'(j)$, we also have that

$$ju(j) - \mathcal{X}_u = u(j+1) \leq u'(j+1) = ju'(j) - \mathcal{X}_u^j, \quad (3.56)$$

where $\mathcal{X}_u^j = ju'(j) - u'(j+1) \leq \mathcal{X}_u$ by definition in Eq. (3.55), 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 Eq. (3.53) is diminished when $u(i) \leq u'(i)$ for all i and $\text{Eff}(w_{\text{sc}}, u; 1) \geq \text{Eff}(w_{\text{sc}}, u'; 1)$, 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 $u^{\mathcal{X}}$ in Eq. (3.54) when we limit $n \rightarrow \infty$. Finally, we note that for infinite n , $\mathcal{X} < \frac{1}{e-1}$ is not achievable, as shown in [35]. \square

With the two previous lemmas, we can move to proving Theorem 7. We first characterize a closed form expression of the maximal utility rule $u^{\mathcal{X}}$, which is given in Lemma 6. We fix \mathcal{X} so that $\text{Eff}(w_{\text{sc}}, u^{\mathcal{X}}; \infty) = \frac{1}{\mathcal{X}+1} = Q$. To calculate the expression for $u^{\mathcal{X}}$ for a given \mathcal{X} , a corresponding time varying, discrete time system to Eq. (3.54) 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

$$y(1) = 1$$

$$y(t) = \max \left[\prod_{\ell=1}^{t-1} \ell - \mathcal{X} \left(\sum_{\tau=1}^{t-2} \prod_{\ell=\tau+1}^{t-1} \ell \right) - \mathcal{X}, 0 \right] \quad t > 1.$$

Simplifying the expression and substituting for $u^{\mathcal{X}}(j)$ gives

$$u^{\mathcal{X}}(j) = \max \left[(j-1)! \left(1 - \mathcal{X} \sum_{\tau=1}^{j-1} \frac{1}{\tau!} \right), 0 \right] \quad j \geq 1.$$

Substituting the expression for the maximal $u^{\mathcal{X}}$ into Eq. (3.53) gives the one round efficiency. Notice that for $\mathcal{X} \geq \frac{1}{e-1}$, $\lim_{j \rightarrow \infty} u^{\mathcal{X}}(j) = 0$, and therefore $\min_j u^{\mathcal{X}}(j) = 0$. Shifting the variables $j' = j + 1$, we get the statement in Eq. (3.52). \square




Figure 3.8: We depict the Pareto-optimal frontier of the one-round efficiency $\text{Eff}(w_{\text{sc}}, u; 1)$ versus the asymptotic efficiency guarantees $\text{Eff}(w_{\text{sc}}, u; \infty)$ 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.

Notably in Figure 3.8, we see a stark drop-off in transient efficiency when the asymptotic efficiency is close to the optimal guarantee of $1 - \frac{1}{e}$. This extreme trade-off should prompt a more careful interpretation of asymptotic results, especially in the setting of resource allocation games.

3.7 Simulations

3.7.1 Comparison of One-Round Guarantees

In this section, we compare the efficiency guarantees of optimal utility design and the common interest design in Figure 3.9. We do this by utilizing the results from Theorem 1 for a distribution of various welfare rules. Specifically, given a welfare rule w and respective utility rule u defined only on a finite number of entries, an efficiency guarantee $\text{Eff}(\mathcal{G}_{\{w\},\{u\}})$ can be derived from Theorem 1 through considering only a finite number of constraints.

First, we randomly generate 5000 welfare rules $w \in \mathbb{R}^{15}$ defined for 15 entries. We do this by fixing $w(1) = 1$ and sampling the difference $w(j) - w(j - 1)$ uniformly from $[0, 1]$ for all $2 \leq j \leq 15$ to recursively generate a normalized and monotone welfare rule. For each welfare rule, we derive the efficiency guarantee of the common interest design by running the linear program in Eq. (3.13), with $H = \max_i \bar{w}(i)/i$ and the marginal contribution utility rule. Additionally, for each welfare rule, 200 normalized utility rules are randomly generated, where $u(1) = 1$ and $u(j)$ is uniformly sampled from $[0, 1]$ for $2 \leq j \leq 15$, and the efficiency guarantees for each utility rule are derived. The maximum calculated efficiency from this set is taken as a lower bound for the guarantees of the optimal utility design. Accordingly, we generate a histogram of efficiency guarantees for the common interest design and the optimal utility design, shown in Figure 3.9.

While the calculated efficiency results may not be tight for both the common interest and optimal utility design, this simulation highlights the potential gains from considering a non-oblivious design. As the mean guarantee of the common interest design is .28 and the mean guarantee of the optimal utility design is .54, we see an improvement of 93% in the mean efficiency guarantee. We remark that since the optimal utility design has equivalent run-time as the common interest design, these gains are realizable just through manipulating the local objectives for agents.

Non-Oblivious vs Standard Guarantees




Figure 3.9: In this figure, we compare a histogram of efficiency results of the common interest and optimal utility designs for a distribution of welfare rules. Each sample corresponds to computing a lower bound on the efficiency guarantees for a given welfare rule set $\mathcal{W} = \{w\}$ through utilizing the results in Theorem 1. We observe a noticeable improvement in the efficiency guarantees from the common interest to the optimal utility design.

3.7.2 Multiple Rounds

To illustrate the theoretical results, we examine the average performance over 5 rounds of the k -round algorithm of three utility designs: the common interest utility design, the utility design that optimizes the efficiency for $k = 1$, and the utility design that optimizes the efficiency for $k = \infty$. The average performance is measured across 100 random instances of *weapon-target assignment problems* (see Example 1) with 20 agents with a defense rate of $p_d = .5$. In each simulated instance, we set the number of targets that the agents can possibly defend to 30. The values v_r for each target r are uniformly selected from the unit interval $[0, 1]$ and subsequently normalized by dividing by $\sum_{r \in \mathcal{R}} v_r$. Each agent has 2 actions available, in addition to the empty allocation a^\emptyset . Each action a is a consecutive selection of 2 resources chosen uniformly randomly from the resource set \mathcal{R} .

The resulting system welfare across 5 rounds for each utility design is highlighted in Figure 3.10, where the distributions of the system welfare across the randomized instances are depicted with a box and whisker plot. Note that the optimal allocation may also not achieve a 100% detection rate. In Figure 3.10, we see that worst instance of the optimal one-round performs better than the greedy and asymptotically optimal utility designs when $k = 1$. This is supported in the worst-case analysis presented in this chapter. Additionally, we note that the resulting efficiency plateaus quickly, with almost no differences in efficiency after two rounds of best response - confirming that successive rounds give diminishing returns in system performance. Interestingly, on average, the differences in performance across utility designs is much more subtle.




Figure 3.10: We plot the average rate of defense in a randomly generated set of weapon-target assignment problems with respect to three utility designs: the one-round optimal, the common interest, and the asymptotically optimal utility design. We see that in the short term, the one-round optimal design performs better in the worst case than the common interest and the asymptotically optimal utility designs.

Chapter 4

Greedy Algorithms in Limited Information Settings

In this chapter, we investigate a multi-agent decision problem where agents aim to coordinate and optimize a given system-level objective. While finding the globally optimal solution may be intractable, the *greedy algorithm* is a classical and efficient approach for obtaining good approximate solutions, particularly for submodular optimization domains. The execution of the greedy algorithm requires agents to be ordered such that each agent performs a local optimization based on the solutions provided by preceding agents.

In settings with limited information, passing solutions from one agent to another can be challenging, as direct communication may not be possible. Thus, the communication time required for executing the greedy algorithm is heavily dependent on the order of which the agents are arranged in. In this study, we characterize the relationship between the communication complexity of running the greedy algorithm and agent ordering. We see that the complexity is $O(n)$ with the optimal ordering, but increases significantly to $O(n^2)$ in the worst possible ordering. Inspired by these findings, we propose an

algorithm that identifies an efficient ordering and executes the greedy algorithm in a computationally efficient fashion. Our proposed method shows significant advantages over existing approaches for distributed submodular maximization, offering both theoretical insights and practical benefits. The results and discussion in this chapter is based on the work presented in [45].

4.1 Introduction

Many real-world problems are well-modeled as multiagent decision problems. In these scenarios, the set of n decision makers, or agents, coordinate to a joint decision that maximizes some objective function. In general, finding the optimal decision set is computationally intractable, even for a centralized authority. Therefore, there exist a multitude of techniques for arriving at a joint decision, which may be an approximation of the optimal. For instance, consensus algorithms offer a way for agents to converge as a group toward a unified decision [46–48]. In other settings, a game-theoretic approach is advantageous, where agents arrive at a joint decision which is some form of equilibrium (e.g., Nash equilibrium [11], Wardrop equilibrium [49], Stackelberg equilibrium [50], etc.).

Another common approach to multiagent decision problems is a greedy algorithm [51, 52]. A common theme among greedy algorithms is that at each iteration of the algorithm, a myopic choice is made: simply pick the best immediate option, ignoring the effect on future iterations of the algorithm. As with the other algorithms mentioned above, greedy algorithms in general are not always guaranteed to find an optimal solution to a given problem, however, they are often easy to implement, execute quickly, and in some cases provide some degree of optimality.

This chapter focuses on the scenario where a greedy algorithm is used to solve a multiagent decision problem. In this setting, a greedy algorithm is implemented by first

ordering the agents. Then each agent sequentially makes its decision by choosing the action that maximizes the objective function, based solely on the decisions of previous agents in the sequence. An underlying element of the greedy algorithm is that the agents are able to coordinate with each other via some network. In the best case, such a network would allow for each agent to communicate with all other agents directly. In many applications, however, this is not realistic; communication between agents i and j must pass through other agents in the network. If i and j are on opposite ends of the network, or if the network has highly-limited bandwidth, this communication may be delayed. In light of this, two questions arise:

1. Given the structure of the communication network, how does the ordering of the agents affect the time it takes to complete the greedy algorithm?
2. Can the agents coordinate among themselves to find the ordering that will cause the greedy algorithm to complete as fast as possible?

We address the first question by showing, given a network structure, that the greedy algorithm finishes in $O(n^2)$ time steps for the worst ordering and $O(n)$ time steps for the best ordering. We then address the second question by presenting a fully-distributed algorithm whereby agents can find a near-optimal ordering while simultaneously running the greedy algorithm.

Of particular import in this chapter are submodular maximization problems, which are prevalent in modeling many applications, such as sensor placement [12], data summarization [53, 54], robot path planning [55, 56], task allocation [57], inferring influence in a social network [58], image segmentation [59], outbreak detection in networks [60], and leader selection in multiagent systems [61]. A key feature that is shared among the objective functions in these various domains is a property of *diminishing returns*. For example, in outbreak detection, the added benefit of placing an outbreak sensor on a node in a network

is valuable when there are few other sensors in the network, and less valuable when there are already many other sensors present. Objectives that exhibit such properties are *submodular*.

While such problems are NP-Hard in general, the property of submodularity can be exploited to show that certain algorithms can achieve near-optimal results. The seminal work in [38] shows that a centralized greedy algorithm can, in fact, provide a solution that is guaranteed to be within 1/2 of the optimal solution. More sophisticated algorithms have pushed this guarantee from 1/2 to $1 - 1/e \approx 0.63$ [62, 63]. Progress beyond this approximation frontier is not possible for polynomial time algorithms as it was also shown that no such algorithm can achieve higher guarantees than $1 - 1/e$, unless $P = NP$ [64]. It will be shown that the greedy algorithm will complete in fewer time steps than existing methods, while still maintaining 1/2-optimality in the resulting decision set.

4.2 Model and Preliminaries

Consider a distributed optimization problem with n agents $\mathcal{I} = \{1, \dots, n\}$, where each agent is endowed with a decision or action set \mathcal{A}_i . We denote an action as $a_i \in \mathcal{A}_i$, and a joint action profile as $a \in \mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$. We assume that each agent i has the ability to “opt out” of participating in the decision process. This is modeled by having an action $a_i^\emptyset \in \mathcal{A}_i$, so that when agent i chooses action a_i^\emptyset , agent i is opting out. The quality of each joint action profile is evaluated with a global objective function $W(a) : \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}$ that a system designer seeks to maximize. In other words, the goal of the system designer is to coordinate the agents to a joint action profile that satisfies

$$a^{\text{opt}} \in \arg \max_{a \in \mathcal{A}} W(a). \quad (4.1)$$

In general, solving the multi-agent decision problem in Eq. (4.1) is infeasible, due to

computational, informational, communication constraints etc. Therefore, fast, distributed algorithms are employed to compute good approximate solutions. The *greedy algorithm* has cemented its place as a universal approach to arrive at approximate solutions in many application domains. In this algorithm, the set of agents is ordered (for instance, according to its index i) and then each agent sequentially solves the reduced optimization problem

$$a_i^{\text{gr}} \in \arg \max_{a_i \in \mathcal{A}_i} W(a_1^{\text{gr}}, \dots, a_{i-1}^{\text{gr}}, a_i, a_{i+1}^{\emptyset}, \dots, a_n^{\emptyset}), \quad (4.2)$$

where each agent i chooses the best action a_i^{gr} given that the previous agents in the sequence have also played their best action and the successive agents in the sequence have opted out. After each agent chooses according to Eq. (4.2), then the algorithm is complete and the resulting set of decisions $(a_1^{\text{gr}}, \dots, a_n^{\text{gr}})$ comprises the joint decision set a^{gr} . The process completes in n time steps, where a time step is comprised of an agent making a decision and communicating that decision to future agents in the sequence.

However, the greedy algorithm makes a key assumption that agents have access to the decisions of the previous agents. In purely distributed systems, this assumption may be infeasible. There have been prior works that study the performance of the greedy algorithm with relaxed informational assumptions, in that agent i only knows the decisions of some strict subset of the previous agents $S \subset \{1, \dots, i-1\}$ [52, 65]. However, this work takes a different approach, where we assume that agents can make up for their informational deficiencies through a communication infrastructure. We model the communication constraints through an underlying graph structure $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where each vertex in \mathcal{V} corresponds to an agent in \mathcal{I} and each edge $(i, j) \in \mathcal{E} = \mathcal{V} \times \mathcal{V}$ implies that agents i and j can communicate with one another. The graph \mathcal{G} is assumed to be connected and undirected throughout this chapter, unless explicitly stated. The set of agents that agent i can communicate with is agent i 's neighborhood \mathcal{N}_i .

The primary focus of this work is to examine the interplay between the communication graph \mathcal{G} and the order in which the greedy algorithm in Eq. (4.2) is solved under. For a given graph \mathcal{G} , the order $\pi : \mathcal{V} \rightarrow \mathcal{I}$ is defined by which label i is given to each vertex v . Therefore, given \mathcal{G} , we would like to characterize the communication time guarantees of the worst order and the best order. To analyze the spectrum of possible guarantees with respect to different ordering methods, we define the following two quantities

$$T_{\min}(\mathcal{G}) = \min_{\pi} T(\mathcal{G}, \pi), \quad (4.3)$$

$$T_{\max}(\mathcal{G}) = \max_{\pi} T(\mathcal{G}, \pi), \quad (4.4)$$

where $T(\mathcal{G}, \pi)$ refers to the time it takes for the communication process to finish for a given graph \mathcal{G} and ordering π . We will use π_{best} and π_{worst} to refer to the orderings that are the solutions of Eq. (4.3) and Eq. (4.4) respectively. We remark that only in the full information setting, where \mathcal{G}_c is the complete graph, is the run-time for any order the same, with $T_{\max}(\mathcal{G}_c) = T_{\min}(\mathcal{G}_c) = n - 1$.

We can describe the k -hop communication, in which an agent i 's greedy action a_i^{gr} is passed along to agents outside of its neighborhood \mathcal{N}_i , using the following graph-theoretic notation. A *walk* on the graph \mathcal{G} is a sequence of vertices $\gamma = (v_1, \dots, v_m)$, in which each successive pair $(v_j, v_{j+1}) \in \mathcal{E}$ for all $1 \leq j < m$. We denote the length of the walk as $|\gamma|$ being the number of vertices in the sequence. A *spanning walk* is a walk in which all vertices in the graph are visited and a *minimum spanning walk* is a spanning walk with shortest length. A *path* p is a walk in which all the vertices $\{v_j\}_{j \leq m}$ are all distinct. The expression of $T(\mathcal{G}, \pi)$ is given as

$$T(\mathcal{G}, \pi) = \sum_{i=1}^{n-1} \left(\min_{p_{i \rightarrow i+1}} |p_{i \rightarrow i+1}| - 1 \right), \quad (4.5)$$

where $p_{i \rightarrow i+1}$ is a path on the graph from the vertex (labeled with) i to the vertex $i + 1$. This expression is motivated by a natural communication process, where initially agent 1 computes its greedy action a_1^{gr} at time 0. Then agent 1 communicates a_1^{gr} to agent 2 through a k -hop walk through the graph, where each hop is assumed to take 1 time step. Then agent 2 computes a_2^{gr} given a_1^{gr} and passes both actions to agent 3 through another k -hop walk. Continuing this process, agent $n - 1$ passes $\{a_j^{\text{gr}}\}_{j < n}$ to agent n and agent n computes a_n^{gr} finishing the process. To isolate the run-time analysis with respect to only the communication time, we also assume that agents can solve for their greedy action a_i^{gr} arbitrarily fast.

4.3 Main Theoretical Results

4.3.1 Motivating Example

To make the communication process concrete, consider when the given communication graph is a line graph as shown in Figure 4.1. In this graph scenario, agent 1 initially computes its greedy response a_1^{gr} and passes the action it has played to agent 2 at $t = 0$. Then at $t = 1$, agent 2 (knowing a_1^{gr}) can compute a_2^{gr} and passes both a_1^{gr} and a_2^{gr} along to agent 3. Continuing this to $t = 5$, agent 6 will have been passed the greedy actions of all previous agents 1 through 5 and play its greedy action a_6^{gr} , completing the greedy algorithm in Eq. (4.2). This will complete in $T(\mathcal{G}, \pi) = 5$ time steps, which is the best that one can hope for when implementing the greedy algorithm in a limited information setting.

However, consider the following order in Figure 4.2. This situation can occur if the order π is improperly picked by the system operator. Under this ordering, agent 2 can only receive the greedy action of agent 1 through a 3-hop path through agents 6 and 4,



Figure 4.1: Example of a line graph, where we have labeled the vertices according to the best ordering π_{best} . In this example, the agents compute their greedy action and pass it down the line.

since there is not a direct communication link between agent 1 and agent 2. Following this logic, the greedy algorithm will complete at time $T(\mathcal{G}, \pi) = 3 + 5 + 4 + 3 + 2 = 17$, which can be seen to be significantly higher than the previous well chosen order.




Figure 4.2: Example of a line graph, but instead we consider the adversarial ordering π_{worst} in which the vertices are labeled intermittently. In this instance, the k -hop communication path must bounce back and forth between agents to complete the greedy algorithm.

Extending this argument to n agents, under a line graph, the greedy algorithm under the best ordering π_{best} will complete in $T(\mathcal{G}, \pi_{\text{best}}) = n - 1$ steps and under the worst ordering π_{worst} will complete in $T(\mathcal{G}, \pi_{\text{worst}}) = \lfloor n^2/2 \rfloor - 1$ steps, where $\lfloor a \rfloor$ is the floor function. Therefore, there may be a significant gap in the communication complexity that results from choosing different orderings. We analyze the possible gap by characterizing the quantities $T_{\max}(\mathcal{G})$ and $T_{\min}(\mathcal{G})$ in this chapter.

4.3.2 Communication Run-time Characterizations

We outline the main theorem of the chapter below, where the worst case communication time over any graph structure is given for the best and worst orderings. The corresponding graph structures and orders that attain the worst-case communication time are displayed in Figure 4.2 and Figure 4.3.




Figure 4.3: A 7 node stargraph with agent 2 in the center. We note that the communication time $T(\mathcal{G}, \pi)$ on this graph using any ordering must be greater than $2 \cdot 7 - 4 = 10$. For $n \geq 3$ agents, the n -node star graph is the worst case graph for the best ordering π_{best} .

Theorem 8. *Let $n \geq 3$ be the number of agents. The maximum communication time required to complete the greedy algorithm in Eq. (4.2) for any undirected, connected communication graph \mathcal{G} with the best and worst orderings is equal to*

$$\max_{\mathcal{G}} T_{\min}(\mathcal{G}) = 2n - 4 \quad (4.6)$$

$$\max_{\mathcal{G}} T_{\max}(\mathcal{G}) = \lfloor n^2/2 \rfloor - 1, \quad (4.7)$$

where $T_{\min}(\mathcal{G})$ is defined in Eq. (4.3) and $T_{\max}(\mathcal{G})$ is defined in Eq. (4.4) and $\lfloor a \rfloor$ is the largest integer that is below a .

Proof. We first show the equality in Eq. (4.6). For a given graph \mathcal{G} and the optimal ordering π_{best} , we claim that the communication time is equal to

$$T_{\min}(\mathcal{G}) = |\gamma_{\min}|, \quad (4.8)$$

where γ_{\min} is a minimum spanning walk of \mathcal{G} . Note that for any given ordering π , the

communication time for $T(\mathcal{G}, \pi)$ is given in Eq. (4.5). Let $p_{i \rightarrow i+1}^*$ be the shortest path from i to $i+1$. The walk $\gamma_{\text{cat}} = p_{1 \rightarrow 2}^* p_{2 \rightarrow 3}^* \cdots p_{n-1 \rightarrow n}^*$ is defined as the concatenation of the shortest paths from 1 to n with the duplicate vertices from $p_{i \rightarrow i+1}^*$ and $p_{i+1 \rightarrow i+2}^*$ removed. Then according to Eq. (4.5), we have that $T(\mathcal{G}, \pi) = |\gamma_{\text{cat}}|$. We note that since $\pi^{-1}(i)$ and $\pi^{-1}(i+1)$ are in $p_{i \rightarrow i+1}^*$, then γ_{cat} is a spanning walk and thus $|\gamma_{\text{cat}}| \geq |\gamma_{\min}|$. Since the ordering π was arbitrary and $T(\mathcal{G}, \pi) = |\gamma_{\text{cat}}|$, we have that $\min_{\pi} T(\mathcal{G}, \pi) \geq |\gamma_{\min}|$. This expression actually holds with equality, as π_{best} can be taken as the order that the vertices first appear in γ_{\min} , matching Eq. (4.8) and the claim is shown.

Notice that a spanning walk for the graph $\mathcal{G}_1 = (\mathcal{V}_1, \mathcal{E}_1)$ is also a spanning walk for the graph $\mathcal{G}_2 = (\mathcal{V}_1, \mathcal{E}_1 \cup \{e\})$ with the added edge e . Then the length of the minimal spanning walk $|\gamma_{\min}^1|$ for \mathcal{G}_1 must be at least $|\gamma_{\min}^2|$ for \mathcal{G}_2 . Thus $T_{\min}(\mathcal{G}_1) \geq T_{\min}(\mathcal{G}_2)$, and to calculate $\max_{\mathcal{G}} T_{\min}(\mathcal{G})$, it is sufficient to restrict to the class of tree graphs, which is the class of connected graphs with the least number of edges.

Consider any spanning walk γ on the tree graph $\mathcal{G}_{\mathcal{T}} = (\mathcal{V}_{\mathcal{T}}, \mathcal{E}_{\mathcal{T}})$ starting at the vertex v_1 and ending at the vertex v_n . Let $\mathcal{E}_{\mathcal{T}}^p \subset \mathcal{E}_{\mathcal{T}}$ be the set of edges that belong to the unique path p between v_1 and v_n . We claim that γ must visit each edge $e \in \mathcal{E}_{\mathcal{T}}$ at least once and every edge $e \in \mathcal{E}_{\mathcal{T}} \setminus \mathcal{E}_{\mathcal{T}}^p$ at least twice. If there exist an edge $\hat{e} \in \mathcal{E}_{\mathcal{T}}$ such that $\hat{e} \notin \gamma$, then since γ is a spanning walk, then the graph $\hat{\mathcal{G}} = (\mathcal{V}_{\mathcal{T}}, \mathcal{E}_{\mathcal{T}} \setminus \{\hat{e}\})$ must also be a connected graph. But this is a contradiction, since $\mathcal{G}_{\mathcal{T}}$ is assumed to be a tree graph. If an edge $\bar{e} \in \mathcal{E}_{\mathcal{T}} \setminus \mathcal{E}_{\mathcal{T}}^p$ is removed from $\mathcal{G}_{\mathcal{T}}$, there must be two nonempty components, one containing both v_1 and v_n and another containing neither. Thus if \bar{e} is only traversed once in the walk, it must hop from v_1 from the first component to the other component once. However, the spanning walk γ cannot come back to the first component again, contradicting our definition of v_n and the claim is shown.

Therefore for a tree graph $\mathcal{G}_{\mathcal{T}}$ and any spanning walk γ , we have that $|\gamma| \geq 2|\mathcal{E}_{\mathcal{T}}| - |\mathcal{E}_{\mathcal{T}}^p|$. Moreover there exists a spanning walk that has the length equal to $|\gamma| = 2|\mathcal{E}_{\mathcal{T}}| - |\mathcal{E}_{\mathcal{T}}^p|$ in

which the vertices that are not along the path p from v_1 to v_n are reached through a cycle that visits each edge not part of p twice. Thus, the length of the minimum spanning walk can be written as the following optimization problem. Here, $\text{diam}(\mathcal{G}_{\mathcal{T}})$ is the diameter of the graph $\mathcal{G}_{\mathcal{T}}$ and $\Gamma_{\mathcal{G}_{\mathcal{T}}}$ is the set of spanning walks.

$$\begin{aligned} |\gamma_{\min}| &= \min_{\gamma \in \Gamma_{\mathcal{G}_{\mathcal{T}}}} |\gamma| \\ &= 2|\mathcal{E}_{\mathcal{T}}| - \max_{v_1, v_n} |\mathcal{E}_{\mathcal{T}}^p| = 2(n-1) - \text{diam}(\mathcal{G}_{\mathcal{T}}). \end{aligned}$$

For a given tree graph $\mathcal{G}_{\mathcal{T}}$ with more than $n \geq 3$ vertices, the diameter must be greater than $\text{diam}(\mathcal{G}_{\mathcal{T}}) \geq 2$. Therefore, the length of the minimum spanning walk must be less than $2n - 4$ for any tree graph $\mathcal{G}_{\mathcal{T}}$. Moreover, the star graph is the tree graph with a graph diameter of 2, so $\max_{\mathcal{G}} T_{\min}(\mathcal{G}) = 2n - 4$.

Now, we show the equality in Eq. (4.7). We claim that the connected, undirected graph that attains $\max_{\mathcal{G}} T_{\max}(\mathcal{G})$ is the line graph. We observe, similarly as before, that if a path $p_{i \rightarrow i+1}$ exists from i to $i+1$ for the graph $\mathcal{G}_1 = (\mathcal{V}_1, \mathcal{E}_1)$, then it must also exist for the graph $\mathcal{G}_2 = (\mathcal{V}_1, \mathcal{E}_1 \cup \{e\})$ with the added edge e for any $1 \leq i \leq n-1$. As the run-time in Eq. (4.5) is defined by the shortest path from i to $i+1$, the run-time for \mathcal{G}_2 is lower bounded by $T(\mathcal{G}_1, \pi) \geq T(\mathcal{G}_2, \pi)$. Thus, we can assume that worst-case graph is a tree graph without loss of generality. If $\mathcal{G}_{\mathcal{T}}$ is a tree graph, then the path $p_{i \rightarrow i+1}$ from $\pi^{-1}(i)$ to $\pi^{-1}(i+1)$ is unique.

We now claim that for any tree graph $\mathcal{G}_{\mathcal{T}}$ for some ordering π , there exists an ordering π_L with the line graph that achieves at least the same run-time. If $\mathcal{G}_{\mathcal{T}}$ is a tree graph that is not the line graph, there exists at least one vertex v_c of $\mathcal{G}_{\mathcal{T}}$ with degree $d \geq 3$. Let $\{v_j\}_{1 \leq j \leq d}$ be the vertices in the neighborhood of v_c . Also, let $\mathcal{T}_j \subset \mathcal{G}_{\mathcal{T}}$ be the corresponding tree component containing v_j that results from removing the edge (v_j, v_c) . Let $v_j^d \in \mathcal{T}_j$ be the vertex farthest away from v_j and d_j be the distance from v_j to v_j^d .

Additionally let $H = \{(\pi^{-1}(i), \pi^{-1}(i+1))\}_{i < n}$ and

$$H^j = \{(v, v') \in H : v \in \mathcal{T}_1, v' \in \mathcal{T}_j \text{ or } v \in \mathcal{T}_j, v' \in \mathcal{T}_1\}.$$

We also denote that (v, v_c) and (v_c, v) are included in H^1 . Consider the modified graph $\mathcal{G}_{\mathcal{T}}^J$, where the edge $(v_{j=1}, v_c)$ is replaced with $(v_{j=1}, v_{j=J}^d)$ for some $1 < J \leq d$. Then the communication time for the graph $\mathcal{G}_{\mathcal{T}}^J$ is

$$\begin{aligned} T(\mathcal{G}_{\mathcal{T}}^J, \pi) &= \sum_{(v, v') \in H} p_{v, v'}^{\mathcal{G}_{\mathcal{T}}^J} = \sum_{j \geq 1} \sum_{(v, v') \in H^j} p_{v, v'}^{\mathcal{G}_{\mathcal{T}}^J} \\ &\geq \sum_{(v, v') \in H} p_{v, v'}^{\mathcal{G}_{\mathcal{T}}} + \sum_{j > 1} \sum_{(v, v') \in H^j} d_J - 2 \sum_{(v, v') \in H^j} d_J \\ &\geq T(\mathcal{G}, \pi) \text{ for some } J. \end{aligned}$$

Here, $p_{v, v'}^{\mathcal{G}_{\mathcal{T}}^J}$ refers to the path from v to v' in $\mathcal{G}_{\mathcal{T}}^J$. The first equality comes from rewriting Eq. (4.5) using H^j . The last inequality comes from the fact that the degree $d \geq 3$ and that

$$\max_{J > 1} \left\{ \sum_{j > 1} \sum_{(v, v') \in H^j} d_J - 2 \sum_{(v, v') \in H^j} d_J \right\} \geq 0.$$

By applying a similar argument inductively for every tree \mathcal{T}_j for $j \leq d-2$ and for every vertex v_c with degree more than 3, we have the claim.

The worst ordering for the line graph and the corresponding guarantee of $\lfloor n^2/2 \rfloor - 1$ is given by the work in [66, Theorem 8]. We give a sketch of the proof for completeness. Let v_i be the i 'th vertex in the line graph. Consider any ordering π . If π has any of the following properties, then there must exist another ordering π' that has a larger communication time that results from swapping positions of some $\pi(v)$ and $\pi(v')$.

1. For some $1 < i < n$, either $(\pi(v_1), \pi(v_i), \pi(v_{i+1}))$ or $(\pi(v_i), \pi(v_{i+1}), \pi(v_n))$ is monotonic.

2. For some $i, j < n$, there are two pairs $(\pi(v_i), \pi(v_{i+1}))$ and $(\pi(v_j), \pi(v_{j+1}))$ that are separated by some threshold $1 \leq \ell \leq n$.
3. There is a triple $(\pi(v_i), \pi(v_{i+1}), \pi(v_{i+2}))$ that is monotonic.

If an ordering π does not have any of the previous three properties, then it must have all even indexed vertices $\{v_i\}_{i \text{ even}}$ below the threshold $\lfloor n/2 \rfloor$ and all odd indexed vertices $\{v_i\}_{i \text{ odd}}$ above the threshold with the middle vertices corresponding to the endpoints 1 and n . An example of this configuration is shown in Figure 4.2. It can be seen that the line graph with this configuration has a $T(\mathcal{G}, \pi) = \lfloor n^2/2 \rfloor - 1$. \square

4.3.3 Distributed Orderings that are Near-Optimal

According to Theorem 8, there is a significant complexity gap between using the best ordering π_{best} and worst ordering π_{worst} for the communication time. However, finding the best order π_{best} in general is not practical either due to computational restrictions or lack of information about the graph. So we would like to be able compute orderings that get as close to the run-time with π_{best} as possible in a feasible manner. Therefore in this section, we construct an algorithm that can quickly compute a good ordering in conjunction with executing the greedy algorithm. An outline of the proposed algorithm is in Algorithm 2. The proposed design in essence computes a spanning walk on the graph that is close to the length of the minimum spanning walk through a variant of a depth-first search algorithm.

The distributed implementation of Algorithm 2 to compute an approximate solution to Eq. (4.1) offers significant benefits over other distributed approaches. The communication scheme is simple, which allows for linear-time guarantees. This also means that the message complexity is low, where the bulk of the message is comprised of the previous agent's actions and the communication is robust to time delays. Lastly, since the base

of Algorithm 2 is the greedy algorithm, we also inherit the corresponding performance guarantees. To be able to run Algorithm 2, we assume that each agent (vertex) can store and access the following variables.

- $v.\text{actions} = \emptyset$ is the set of greedy actions that v knows.
- $v.\text{order} = \emptyset$ is the index in \mathcal{I} that v is labeled with.
- $v.\text{parent} = \emptyset$ is v 's parent in the depth first search.
- $v.\text{neighborhood}$ is the neighborhood set of v .

We also assume that a seed v_{seed} is given as the starting point of the Algorithm 2. The communication time of Algorithm 2 is equivalent to the total number of calls to MESSAGE, where the vertex v messages either a vertex that hasn't been visited or its parent $v.\text{parent}$. We keep track of the communication time through the variable t . The communication time guarantees of Algorithm 2 is given below.

Proposition 2. *Let n be the number of agents and $T_{\text{alg}}(\mathcal{G}, v_{\text{seed}})$ be the output of Algorithm 2 given a communication graph \mathcal{G} and a seed vertex $v_{\text{seed}} \in \mathcal{V}$. The maximum communication time for any undirected, connected \mathcal{G} and seed v_{seed} is*

$$\max_{\mathcal{G}, v_{\text{seed}}} T_{\text{alg}}(\mathcal{G}, v_{\text{seed}}) = 2n - 2 \quad (4.9)$$

Proof. Consider an arbitrary graph \mathcal{G} with n agents and a seed vertex v_{seed} . Let ℓ be the number of calls to MESSAGE where a vertex v messages another unvisited vertex and m be the number of calls to MESSAGE where a vertex v messages its parent. It can be seen that Algorithm 2 will eventually visit all the vertices in the graph, so ℓ must equal $n - 1$. Additionally, since v_{seed} does not have a parent and Algorithm 2 terminates if v_{seed} does not send a message to an unvisited neighbor, $m \leq \ell = n - 1$. Therefore for any \mathcal{G}

and v_{seed} the communication time is bounded above by $T_{\text{alg}}(\mathcal{G}, v_{\text{seed}}) = m + \ell = 2n - 2$. Furthermore, it can be seen that for the star graph with n vertices, $m = \ell = n - 1$, and the equality in Eq. (4.9) is shown. \square

Thus the communication guarantees of Algorithm 2 is only off by a constant of 2 from the optimal communication guarantee of $2n - 4$ from the best ordering π_{best} . We remark that this difference can be further reduced if the termination condition is changed from ‘ v .parent is not empty’ to ‘ v .order = n ’, where n is the number of agents.

Algorithm 2 Distributed Near-Optimal Ordering

Require: graph \mathcal{G} and a vertex $v_{\text{seed}} \in \mathcal{V}$

Output: time t

initialize the time $t \leftarrow 0$

INIT($v_{\text{seed}}, \emptyset, \emptyset, 1$)

return t

procedure INIT($v, v_{\text{par}}, \alpha, i$)

label v as done

update v .order $\leftarrow i$ and v .parent $\leftarrow v_{\text{par}}$

let v compute a_i^{gr} from Eq. (4.2) given actions α

update v .actions $\leftarrow \alpha \cup \{a_i^{\text{gr}}\}$

MESSAGE(v, v .actions)

end procedure

procedure MESSAGE(v, α)

update v .actions $\leftarrow \alpha$

if exists w in v .neighborhood not labeled done **then**

increment $t \leftarrow t + 1$

INIT($w, v, \alpha, |\alpha|+1$)

else if v .parent is not empty **then**

increment $t \leftarrow t + 1$

MESSAGE(v .parent, α)

end if

end procedure

4.3.4 Directed Communication Graphs

In this section, we consider the communication time guarantees with respect to the more general class of connected, directed graphs using different orderings. Under the class of undirected graphs, there is a significant gap in the communication time guarantees for the best π_{best} and the worst π_{worst} orderings. Not surprisingly, when we relax to the optimization problem $\max_{\mathcal{G}_{\text{dir}}} T_{\text{max}}(\mathcal{G}_{\text{dir}})$ over the class of directed graphs, the worst case guarantees also increase. However, when considering the optimization problem for the best ordering over directed graphs $\max_{\mathcal{G}_{\text{dir}}} T_{\text{min}}(\mathcal{G}_{\text{dir}})$, we have that the worst case guarantees are also of quadratic order. Therefore in directed graphs, the gap between the performance guarantees under different orderings is relatively small.




Figure 4.4: Example of a directed cycle graph, where the vertices are labeled adversarially with π_{worst} . Here, the k -hop communication must cycle back to get to the next agent.

For the graph example in Figure 4.4 using the worst ordering π_{worst} , notice that to get from i to $i + 1$, every edge but one in the directed graph must be traversed, resulting in a communication time of 5×5 for 6 agents. For the graph example in Figure 4.5 using any ordering π , the vertices labeled with 2, 3, and 4 must be traversed every time to reach the vertices labeled with 5, 6, 7, 8 in order, starting from the vertex labeled 1. Thus the communication time for the graph under the best ordering is 4×4 for 8 agents. Extending these constructions to n agents, we arrive at the following lemma.

Proposition 3. *Let $n \geq 3$ be the number of agents. The maximum communication time required to complete the greedy algorithm in Eq. (4.2) for any directed, connected*




Figure 4.5: Example of a directed graph that has the worst communication complexity for the best ordering π_{best} . Here, the vertices in the latter half must cycle back to get to the next labeled vertex.

communication graph \mathcal{G}_{dir} with the best and worst ordering is

$$\max_{\mathcal{G}_{\text{dir}}} T_{\min}(\mathcal{G}_{\text{dir}}) = \left\lfloor \frac{n}{2} \right\rfloor \cdot \left\lceil \frac{n}{2} \right\rceil, \quad (4.10)$$

$$\max_{\mathcal{G}_{\text{dir}}} T_{\max}(\mathcal{G}_{\text{dir}}) = (n-1)^2, \quad (4.11)$$

where $\lfloor a \rfloor$ is largest integer smaller than a and $\lceil a \rceil$ is the smallest integer larger than a .

Proof. First, extending the directed cycle graph \mathcal{C}_n in Figure 4.4 to n agents with the given worst case labeling π_{worst} produces a lower bound

$$\max_{\mathcal{G}_{\text{dir}}} T_{\max}(\mathcal{G}_{\text{dir}}) \geq T(\mathcal{C}_n, \pi_{\text{worst}}) = (n-1)^2.$$

The upper bound $\max_{\mathcal{G}_{\text{dir}}} T_{\max}(\mathcal{G}_{\text{dir}}) \leq (n-1)^2$ also holds, since for any graph \mathcal{G} and ordering π , the communication time must satisfy $T(\mathcal{G}, \pi) \leq \sum_{i=1}^{n-1} (n-1)$ as the length of any path in the graph cannot be greater than $n-1$.

Now we show that $\max_{\mathcal{G}_{\text{dir}}} T_{\min}(\mathcal{G}_{\text{dir}}) \geq \lfloor \frac{n}{2} \rfloor \cdot \lceil \frac{n}{2} \rceil$ by extending the graph construction, denoted \mathcal{D}_n , in Figure 4.5 to n agents. Formally, the edge set of \mathcal{D}_n includes (v_j, v_{j+1}) for

every $1 \leq j \leq \lceil n/2 \rceil - 2$ as well as the edges $(v_{\lceil n/2 \rceil - 1}, v_j)$ and (v_j, v_1) for every $\lceil n/2 \rceil \leq j \leq n$. We confirm that $T(\mathcal{D}_n, \pi_{\text{best}}) = \lfloor \frac{n}{2} \rfloor \cdot \lceil \frac{n}{2} \rceil$ for the best order. Consider any ordering π and without loss of generality, assume that $\pi(v_j) < \pi(v_{j+1})$ for all $\lceil n/2 \rceil \leq j \leq n$. Thus the communication time has to be lower bounded by

$$T(\mathcal{D}_n, \pi) \geq \sum_{j \geq \lceil n/2 \rceil}^{n-1} \min_{p_{v_j \rightarrow v_{j+1}}} |p_{v_j \rightarrow v_{j+1}}| - 1 \geq \lfloor \frac{n}{2} \rfloor \cdot \lceil \frac{n}{2} \rceil,$$

as $(v_j, v_1, v_2, \dots, v_{\lceil n/2 \rceil - 1}, v_{j+1})$ is the unique path $p_{v_j \rightarrow v_{j+1}}$ from v_j to v_{j+1} with a length of $\lceil n/2 \rceil + 1$. Observe that the ordering $\pi(v_j) = j + 1$ for $1 \leq j \leq n - 1$ and $\pi(v_n) = 1$ achieves this communication time and so the lower bound is indeed tight.

Now we show that for any given directed graph \mathcal{G} , there exists an ordering $\hat{\pi}$ in which $T(\mathcal{G}, \hat{\pi}) \leq \lfloor \frac{n}{2} \rfloor \cdot \lceil \frac{n}{2} \rceil$. Let $p_{\text{long}} = (v_1, \dots, v_l)$ be the longest path of the graph, where $l = |p_{\text{long}}|$. If $l = n$, then p_{long} is a spanning walk on the graph and the ordering $\pi(v_j) = j$ along the longest path results in a communication time of $T(\mathcal{G}, \hat{\pi}) = n - 1 \leq \lfloor \frac{n}{2} \rfloor \cdot \lceil \frac{n}{2} \rceil$. Otherwise, since \mathcal{G} is assumed to be strongly connected there exists a vertex v_J in p_{long} that is adjacent to another vertex \bar{v} that is not in the path p_{long} . We construct the ordering $\hat{\pi}$ as follows. The vertices in p_{long} are labeled as $\hat{\pi}(v_j) = j$ for $j \leq J$ and $\hat{\pi}(v_j) = n - l + j$ for $j > J$. The vertex \bar{v} is labeled with $\hat{\pi}(\bar{v}) = J + 1$ and the labels for the rest of the vertices in \mathcal{G} can be arbitrarily selected from $\{J + 2, \dots, n - l + J\}$. The resulting communication time along this order, using Eq. (4.9), is

$$\begin{aligned} T(\mathcal{G}, \hat{\pi}) &= \sum_{i=1}^{n-1} \left(\min_{p_{i \rightarrow i+1}} |p_{i \rightarrow i+1}| - 1 \right) \\ &= (l - 1) + \sum_{i=J+1}^{n-l+J} \left(\min_{p_{i \rightarrow i+1}} |p_{i \rightarrow i+1}| - 1 \right), \end{aligned}$$

as the vertices $\hat{\pi}^{-1}(i)$ and $\hat{\pi}^{-1}(i + 1)$ are adjacent to each other if $1 \leq i \leq J$ or $n - l +$

$J + 1 \leq i \leq n - 1$ according to the prescribed order $\hat{\pi}$. For any $i \in \mathcal{I}$, we observe that $\min_{p_{i \rightarrow i+1}} |p_{i \rightarrow i+1}| \leq l$ by definition of p_{long} . Now we have the upper bound

$$T(\mathcal{G}, \hat{\pi}) \leq (l - 1) + (n - l)(l - 1) \leq \left\lfloor \frac{n}{2} \right\rfloor \cdot \left\lceil \frac{n}{2} \right\rceil,$$

where the expression achieves the maximum at $l = \lceil n/2 \rceil + 1$. Thus, since \mathcal{G} was arbitrary chosen,

$$\max_{\mathcal{G}_{\text{dir}}} T_{\min}(\mathcal{G}_{\text{dir}}) = \max_{\mathcal{G}_{\text{dir}}} T(\mathcal{G}, \hat{\pi}) \leq \left\lfloor \frac{n}{2} \right\rfloor \cdot \left\lceil \frac{n}{2} \right\rceil,$$

and thus we have shown equality. □

4.4 Applications in Submodular Settings

4.4.1 Simulations

We analyze our theoretical results for the communication time guarantees empirically through a simulation, presented in Figure 4.6 and Figure 4.7. We use the model of Erdos-Renyi networks [67], where each possible undirected pair of edges (i, j) has a probability $P \leq 1$ of existing, to generate a sample set of possible graph structures. For Figure 4.6, we sample 200 instances of Erdos-Renyi networks with 6 nodes and a probability parameter of $P = .3$. For each graph, we calculate the communication time for the greedy algorithm using the best ordering π_{best} , the ordering given by Algorithm 2, and a randomly assigned ordering. For Figure 4.7, we sample 300 instances of Erdos-Renyi networks with 40 nodes and a probability parameter of $P = .05$. In Figure 4.7, we calculate the communication time for only the ordering given by Algorithm 2 and a randomly assigned ordering.

We observe in Figure 4.6 that indeed the best ordering π_{best} achieves the lowest distribution of communication times, centered closely to $n = 6$. The distribution of

communication times of the ordering given in Algorithm 2 is noticeably close to the one of π_{best} , with indeed no run-times over $2n = 12$. The distribution of the communications using random orderings does perform the worst with the largest spread. In Figure 4.7, we compare communication times from the ordering from Algorithm 2 directly with the random ordering, as computing the best ordering π_{best} is infeasible for large n . We see the same trends in Figure 4.6 reflected in a more extreme fashion. The distribution of communication times using the ordering of Algorithm 2 is still upper bounded by $2n = 80$. However, the distribution times of communication times using the random ordering is now centered much higher with a larger spread as well. Therefore, we see significant benefits from using the ordering from Algorithm 2 rather than the naive approach of using random ordering.

4.4.2 Submodular Maximization

In this section, we discuss submodular maximization problems, which can be modeled as multiagent decision problems. Consider a base set of elements E , and let $a_i \subseteq E$, $\mathcal{A}_i \subseteq 2^E$, and $a_i^\emptyset = \emptyset$. The objective function takes the form $W(a) = f(\cup_{a_i \in a} a_i)$, where $f : 2^E \rightarrow \mathbb{R}$ has the following properties for any $A \subseteq B \subseteq E$:

1. *Submodular*: $f(A \cup \{x\}) - f(A) \geq f(B \cup \{x\}) - f(B)$ for all $x \in E \setminus B$
2. *Monotone*: $f(A) \leq f(B)$
3. *Normalized*: $f(\emptyset) = 0$

In this setting, it has been shown that the greedy algorithm that is implemented in Algorithm 2 guarantees that $W(\tilde{a}) \geq (1/2)W(a^{\text{opt}})$, where \tilde{a} is defined in Eq. (4.2).

4.4.3 Comparison with Other Distributed Algorithms

As mentioned previously, much work has been done to develop other algorithms to solve submodular maximization. For instance, [68] presents a similar distributed algorithm, using a multilinear extension, and a distributed pipage rounding technique. At each time step, each agent performs a calculation for each action based on a sample of K actions drawn from a probability distribution. After T time steps, the performance guarantee is $(1 - 1/e)(1 - (2d(\mathcal{G})n + n/2 + 1)(n/T))$ with probability at least $1 - 2nTe^{-K/(8T^2)}$. Thus, for high T and $K = O(T^2)$, there is a high probability that the algorithm gives the $1 - 1/e$ guarantee. Using this information, the algorithm could provide a $1/2$ guarantee only for $T \geq 4.78(2d(\mathcal{G})n^2 + n^2/2 + 1)$, and only with high probability when $K = O(T^2)$.

The paper [69] describes a Jacobi-style algorithm, where at each time step agent i creates a strategy profile, i.e., a probability distribution across each of its actions. Then, it chooses K of those values to share with its neighbors to propagate through the network. It was shown that the resulting decision set approaches being within $1/2$ the optimal as the number of iterations increases. It is only shown in the paper that the probability of achieving the $1/2$ guarantee is $1 - O(1/T)$ rather than an explicit time expectation. However, the examples in the paper suggest that it may take $T \geq n^2$ or more time steps to realize this.

In another example, [70] presents the Constraint-Distributed Continuous Greedy (CDCG), a consensus-style algorithm, in which agent i shares an m -vector with all its neighbors at each time step, where m is the number of actions available to i . It is shown that the resulting decision set approaches being within $1 - 1/e$ of the optimal as the number of iterations T increases. The error in the performance guarantee vanishes at a rate of $O(n^{5/2}/T)$, and therefore, it may require $T \geq n^{5/2}$ time steps in order to reach an

acceptable error.

In each of the three methods listed above, each time step requires each agent do perform some calculation for each of its actions. The time requirement for each to reach an acceptable solution is expected to be greater than $2n - 2$, which is the number of time steps it takes to complete Algorithm 2. This suggests that there is a tradeoff between performance guarantees and time complexity: Algorithm 2 achieves the $1/2$ guarantee quickly, but other algorithms converge to a solution within $1 - 1/e$, but more slowly ¹.

¹Although we do not present a rigorous analysis here, we assert that Algorithm 2 also requires less information exchange at each time step. This will be a topic of future work.

Histogram of Communication versus Different Orderings




Figure 4.6: We show the distribution over of communication times needed to complete the greedy algorithm for 200 instances of random graphs generated by a Erdos-Renyi process with respect to the random ordering, the best ordering π_{best} , and the ordering from Algorithm 2. For the graph parameters $n = 6$ number of agents and $P = .3$ probability of an edge existing, we see that π_{best} gives slightly lower average communication times than the ordering from Algorithm 2, but both offer significant improvements over the random ordering.

Histogram of Communication versus Different Orderings




Figure 4.7: We compare the communication times under the random ordering to the times under the ordering algorithm proposed in Algorithm 2. For the graph parameters $n = 40$ number of agents and $P = .05$ probability of an edge existing, we see a marked decrease in the communication time from the random ordering to using the proposed algorithm over a set of 300 randomly generated graphs.

Chapter 5

Nash equilibrium in Uncertain Settings

In recent years, considerable research has focused on the development of distributed protocols for controlling multi-agent systems. The scale and limited communication bandwidth typical of these systems make centralized control infeasible. Due to strict operating conditions, it is unlikely that each agent in a multi-agent system will possess local information that accurately reflects the true state of the system. Despite this, much of the existing literature assumes that agents have perfect knowledge of their environment.

This chapter aims to understand the impact of inconsistencies in agents' local information on the performance of multi-agent systems. Specifically, we approach the design of multi-agent operations from a game-theoretic perspective, where individual agents are assigned utilities that drive their local decision-making processes. We present a practical method for designing utilities that optimize the efficiency of the collective behavior (i.e., the price of anarchy) for certain classes of set covering games, assuming the extent of information inconsistencies is known. In scenarios where the extent of these inconsistencies is unknown, we find, perhaps counterintuitively, that underestimating the level of uncertainty leads

to a better price of anarchy compared to overestimating it. This chapter is based on the work in [71].

5.1 Introduction

The fundamental challenge in the control of multi-agent systems arises from the stringent requirements placed on their scalability, communication and privacy. As these requirements cannot be satisfied by a centralized approach, we must use distributed protocols where the agents act independently according to their local information. A common approach for distributed design is to cast the global system objective as an optimization problem where the aforementioned requirements are embedded as constraints [72]. Then, through careful consideration of the problem's structure, a distributed algorithm is designed that gives good approximate solutions.

A commonly held assumption in the design of distributed protocols is that all the agents either have perfect knowledge on the underlying problem setting or quickly obtain it through communication (see, e.g., [73]). However, in practice, an agent's knowledge on the *ground truth* may be limited, especially in scenarios where accurate estimation of the true state of the environment is difficult. This prompts the following line of questioning:

- *How robust is the performance of a given distributed control algorithm to inconsistencies in agents' knowledge?*
- *Can distributed control algorithms be explicitly designed to be robust against inconsistencies in agents' knowledge?*

In this chapter, we investigate these two questions from a game theoretic perspective. Our main contribution is a general framework for evaluating the robustness of distributed control algorithms in which the agents' decision making is based on their own (possibly

inaccurate) knowledge of the problem parameters. Each agent acts in its own self interest, maximizing its utility function according to its local knowledge of the underlying problem setting. The system performance guarantees are measured under the well studied notion of *price of anarchy* [74], defined as the ratio between the system welfare at the worst system outcome and the system optimum. Here, the worst system outcome is defined as the worst emergent allocation of agents (i.e., pure Nash equilibrium) with respect to the worst possible disposition of the agents' knowledge. We then apply our framework to the class of set covering games [35] in the setting where each agent's estimate of the problem parameters lies within a bounded interval centered around the true system state.

The study of uncertainty among agents is not new in the field of game theory. In fact, John Harsanyi's seminal work on *incomplete information games* [75] was one of the first significant contributions to this field. The incomplete information model has been used to study multi-agent systems in a variety of contexts, including network games [76], team decision making [77] and pursuit/evasion [78]. While this body of important work is relevant to our problem setting, our analysis departs from the incomplete information framework as, in our setting, agents do not possess probabilistic models of the system state and have only limited knowledge on the other agents' beliefs. Thus, the system operator must account for the worst possible scenario when designing the distributed protocols. The differences between the incomplete information framework and the framework proposed in this chapter are roughly comparable to those between stochastic and robust optimization.

The literature on "robust" formulations is much more restricted. In [79], the authors consider how to generate a set of possible utility functions that are consistent with a limited amount of information. A distribution-free analysis of incomplete information games is considered in [80] through a proposed equilibrium concept. To the author's knowledge, the closest work is [81], where price of anarchy results are considered in

scenarios where the agents have biases in their perceived utilities. While similar in flavor, this chapter studies utility design under a difference class of utility deviations that result from limited information scenarios.

In this chapter, we propose a novel game theoretic framework for studying multi-agent systems that does not impose the assumption that agents possess perfect knowledge of the underlying problem setting. We apply our framework to the class of set covering games [35] and study the performance of distributed control algorithms designed without the perfect knowledge assumption in this setting.

5.2 Model and Preliminaries

For any $p, q \in \mathbb{N}$ with $p < q$, let $[p] = \{1, \dots, p\}$ and $[p, q] = \{p, \dots, q\}$. Given a set S , $|S|$ represents its cardinality. We expand on the class of resource allocation games [82] to model a general distributed scenario. A set of agents $\mathcal{I} = [n]$ must be allocated to a set of resources $\mathcal{R} = \{r_1, \dots, r_m\}$. Each agent $i \in \mathcal{I}$ is associated with a set of permissible actions $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$ and we denote the set of admissible joint allocations by the tuple $a = (a_1, \dots, a_n) \in \mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$. Each resource $r \in \mathcal{R}$ is associated with a state y_r and a welfare function $W_r : [n] \times Y_r \rightarrow \mathbb{R}$ that measures the system performance at that resource as a function of its aggregate utilization. In other words, $W_r(k; y_r)$ is the system performance at resource r when there are $k \in [n]$ users selecting r and the local state of resource r is y_r . Lastly, the system-level welfare is captured by the function $W : \mathcal{A} \times Y \rightarrow \mathbb{R}$, where $Y = \prod_{r \in \mathcal{R}} Y_r$ defines the set of possible system states. In general, for a given allocation $a \in \mathcal{A}$ and system state $y \in Y$ the system-level welfare is of the form

$$W(a; y) = \sum_{r \in \mathcal{R}} W_r(|a|_r; y_r), \quad (5.1)$$

where $|a|_r = |\{i \in \mathcal{I} \text{ s.t. } r \in a_i\}|$ represents the number of agents selecting r in the allocation a . Given a system state $y \in Y$, the goal of the multi-agent system is to coordinate to the allocation that optimizes the system-level welfare

$$a^{\text{opt}} \in \arg \max_{a \in \mathcal{A}} W(a; y). \quad (5.2)$$

However, deriving the optimal allocation requires an infeasible amount of computational resources and coordination. Therefore, we focus on arriving at approximate solutions through a designed utility function $U_i : \mathcal{A} \times Y \rightarrow \mathbb{R}$ and model the emergent collective behavior by a pure Nash equilibrium, which we will henceforth refer to simply as an equilibrium. Given the local knowledge $y_i \in Y$ of each agent $i \in \mathcal{I}$, an equilibrium is defined as an allocation $a^{\text{ne}} \in \mathcal{A}$ such that for any agent $i \in \mathcal{I}$

$$U_i(a_i^{\text{ne}}, a_{-i}^{\text{ne}}; y_i) \geq U_i(a_i, a_{-i}^{\text{ne}}; y_i), \quad \forall a_i \in \mathcal{A}_i. \quad (5.3)$$

It is important to highlight that an equilibrium may or may not exist for such situations, particularly in cases where the agents do not evaluate their utility functions for the same state, i.e., $y_i \neq y_j$. Throughout this chapter, we will assume that equilibria exist within the games under consideration. Nevertheless, our results extend to other solution concepts that are guaranteed to exist such as coarse correlated equilibrium [83].

Throughout this chapter, we will focus on the scenario where there is a true system state $y_{\text{true}} \in Y$ and each agent $i \in \mathcal{I}$ has its own knowledge of this state y_i which may or may not reflect the true state, i.e., y_i need not equal y_{true} . While the agents' knowledge $y = (y_1, \dots, y_n) \in Y_n$ will invariably influence their local behavior and resulting equilibria, we will measure the performance of the resulting equilibrium a^{ne} according to the true state, i.e., $W(a^{\text{ne}}; y_{\text{true}})$. Accordingly, our goal is to assess how discrepancies in the agents' knowledge impacts the quality of the resulting equilibria.

Furthermore, we investigate the optimal design of the utility functions in scenarios where such discrepancies may exist.

To ground these questions moving forward, we consider an extension of the utility functions considered in the framework of Distributed Welfare Games [82], where each resource is associated with a utility generating function of the form $U_r : [n] \times Y_r \rightarrow \mathbb{R}$. Here, the utility generating function defines the benefit associated with each agent selecting resource r , and can depend on both (i) the number of agents selecting resource r and (ii) the state of resource r . Given these utility generating functions $\{U_r\}_{r \in \mathcal{R}}$, the utility of an agent $i \in \mathcal{I}$ in an allocation $a \in \mathcal{A}$ is separable and of the form

$$U_i(a; y_i) = \sum_{r \in a_i} U_r(|a|_r; y_{i,r}). \quad (5.4)$$

Note that each agent $i \in \mathcal{I}$ uses its own state values $y_i = \{y_{i,r}\}_{r \in \mathcal{R}}$ to compute its utility at each resource.

We measure the efficiency of the resulting equilibria through the well-studied price of anarchy metric [74]. We begin by formally expressing a game by the tuple

$$\mathcal{G} = (\mathcal{I}, \mathcal{R}, \mathcal{A}, \{W_r, Y_r, U_r\}_{r \in \mathcal{R}}, y_{\text{true}}, \{y_i\}_{i \in \mathcal{I}}).$$

Note that this tuple includes all relevant information to define the game. We define the price of anarchy of the game G by

$$\text{PoA}(G) := \frac{\min_{a \in \text{NE}(G)} W(a; y_{\text{true}})}{\max_{a \in \mathcal{A}} W(a; y_{\text{true}})} \leq 1, \quad (5.5)$$

where $\text{NE}(G) \subseteq \mathcal{A}$ denotes the set of equilibrium of the game G . We will often be concerned with characterizing the price of anarchy for broader classes of games where resources share common characteristics. To that end, let $Z_r = \{W_r, Y_r, U_r\}$ define the characteristics of a given resource r . Further, let \mathcal{Z} denote a family of possible resource

characteristics. We define the family of games $\mathcal{G}_{\mathcal{Z}}$ as all games of the above form where $\{W_r, Y_r, U_r\} \in \mathcal{Z}$ for each resource $r \in \mathcal{R}$. The price of anarchy of the family of games $\mathcal{G}_{\mathcal{Z}}$ is defined as

$$\text{PoA}(\mathcal{G}_{\mathcal{Z}}) := \inf_{G \in \mathcal{G}_{\mathcal{Z}}} \text{PoA}(G) \leq 1. \quad (5.6)$$

For brevity we do not explicitly highlight the number of agents in a class of games as that is always assumed to be less than n . In order to express the informational inconsistencies between the agents' knowledge and the true state, we define the metric $\rho_d : \mathcal{G}_{\mathcal{Z}} \rightarrow \mathbb{R}_{\geq 0}$ as

$$\rho_d(G) = \max_{i \in \mathcal{I}} d(y_i^G; y_{\text{true}}^G), \quad (5.7)$$

where $d : Y \times Y \rightarrow \mathbb{R}_{\geq 0}$ is some distance measure such that $d(y, y') = 0$ if and only if $y = y'$. Observe that, under this notation, a perfect information scenario where all agents know the true state corresponds with $d(y_i^G; y_{\text{true}}^G) = 0$ for all $i \in \mathcal{I}$ and $\rho_d(G) = 0$. Conversely, when the agents have limited knowledge on y_{true}^G , $\rho_d(G)$ measures the extent of the uncertainty where a higher $\rho_d(G)$ indicates that the agent's evaluation of the state y_i^G is "further" from the true state y_{true}^G . Consolidating these limitations, the set of games in which $\rho_d(G) \leq \delta$ is denoted by $\mathcal{G}_{\mathcal{Z}}^{\delta}$.

In particular, we use the following distance measure for the rest of the chapter:

$$d(y; y_{\text{true}}) = \max_{r \in \mathcal{R}, k \in [n]} \frac{|W_r(k; y_r) - W_r(k; y_{\text{true},r})|}{W_r(k; y_{\text{true},r})}. \quad (5.8)$$

Note that, for a given instance, this measure allows us to capture the level uncertainty that agents have on the system welfare independently of the individual resources. For context, we introduce the following application domains.

Example 4 (Forest Fire Detection). Consider the scenario detailed in [84] where a set

of unmanned aerial vehicles (UAVs) coordinate to cover a forest region to maximize the detection of a forest fire - modeled as a covering game [35]. The UAVs are the agents in the game and the resource set \mathcal{R} correspond to a finite partition of the forest region that the UAVs are tasked to cover. Each UAV carries a sensor with a limited sensing range and must select a position to survey (with a corresponding sensing range) - this choice is modeled by an action set \mathcal{A}_i . The state y_r of each resource r corresponds to the risk that a forest fire might emerge in that resource. We wish to allocate the UAVs to maximize

$$W(a; \{y_r\}_{r \in \mathcal{R}}) = \sum_{r \in \cup a_i} y_r, \quad (5.9)$$

which must balance focusing on the high risk areas and covering as much of the forest region as possible.

Example 5 (Weapon-Target Assignment). Consider the weapon-target assignment problem described in [85] where a set of weapons $\mathcal{I} = \{1, \dots, n\}$ are assigned to a set of targets \mathcal{T} with the objective of maximizing the expected value of targets engaged. When $k \in \{1, \dots, n\}$ weapons engage a target $t \in \mathcal{T}$, its expected value is $v_t \cdot [1 - (1 - p_t)^k]$ where $v_t \geq 0$ is t 's associated value and $p_t \in [0, 1]$ is t 's probability of successful engagement. Based on its range and specifications, each weapon $i \in \mathcal{I}$ can only engage particular subsets of the targets corresponding with the actions $a_i \in \mathcal{A}_i \subseteq 2^{\mathcal{T}}$. Accordingly, under an allocation of weapons $a = (a_1, \dots, a_n)$, the operator's welfare is measured as

$$W(a; \{(v_t, p_t)\}_{t \in \mathcal{T}}) = \sum_{t \in \mathcal{T}} v_t \cdot [1 - (1 - p_t)^{|a_t|}]. \quad (5.10)$$

Observe that this scenario can be modeled as a resource allocation game where each weapon is an agent, each target is a resource and each target t has state $y_t = (v_t, p_t)$.

Though a resource characteristic is a triplet $Z_r = \{W_r, Y_r, U_r\}$, in many cases only

W_r and Y_r are inherited from the problem setting, while the utility generating rule U_r is designed. Accordingly, we will often think of the utility generating function at each resource as being derived from $\{W_r, Y_r\}$, i.e., $U_r = \Pi(W_r, Y_r)$ where Π is the *utility mechanism*. Let the set $\mathcal{Z}(\Pi) = \{W_r, Y_r, \Pi(W_r, Y_r)\}_{r \in \mathcal{R}}$.

The main focus of this chapter is on determining the utility mechanism that maximizes the price of anarchy, i.e.,

$$\Pi^{\text{opt}} = \arg \max_{\Pi} \text{PoA}(\mathcal{G}_{\mathcal{Z}(\Pi)}^{\delta}). \quad (5.11)$$

Accordingly, one may wish to understand how the achievable performance guarantees are affected by the amount of uncertainty $\delta \geq 0$. In scenarios where the system designer does not know the true value of δ , one may additionally seek to characterize the degradation in performance guarantees for estimates on δ of varying levels of accuracy. In the forthcoming sections, we provide preliminary results along these lines of questioning.

5.3 Characterization of PoA

Having defined our general model for limited information scenarios, in this section, we concentrate on a specific class of resource characteristics. Doing so allows us to formulate the optimization problem in (5.11) as a linear program by leveraging recent results in [86]. Let $\mathcal{Z}_{w,u}$ correspond to a set of resource characteristics,¹

$$Y_r = \mathbb{R}_{\geq 0} \quad (5.12)$$

$$W_r(|a|_r; y_r) = y_r \cdot w(|a|_r) \quad (5.13)$$

$$U_r(|a|_r; y'_r) = y'_r \cdot u(|a|_r) \quad (5.14)$$

¹The results in this section can be extended to settings where W_r and U_r are linear combinations over a set of basis functions pair $\{w^j, u^j\}$, $j = 1, \dots, L$, following the results in [83]. We state our results for only one basis function pair $\{w, u\}$ (i.e., $L = 1$) for ease of presentation.

respectively, where $y_r, y'_r \in Y_r$ and $w : [n] \rightarrow \mathbb{R}_{>0}$ and $u : \{1, \dots, n\} \rightarrow \mathbb{R}$ are fixed across all resources $r \in \mathcal{R}$. We assume that $w(0) = u(0) = 0$ and $w(1) = u(1) = 1$ to normalize the welfare W and utility U_i functions. With abuse of notation, we use the denotation $\mathcal{G}_{w,u}$ to refer to the family of games $\mathcal{G}_{\mathcal{Z}_{w,u}}$. In this model, y_r corresponds to a measure of value of the resource r . Additionally, when a set of agents cover a certain resource r , w and u correspond to the resource agnostic measure of the added system welfare and agent utility, respectively. In this setting, the distance measure in (5.8) can be rewritten as

$$\begin{aligned} d(y; y_{\text{true}}) &= \max_{r \in \mathcal{R}, k \in [n]} \frac{|y_r \cdot w(k) - y_{\text{true},r} \cdot w(k)|}{y_{\text{true},r} \cdot w(k)} \\ &= \max_{r \in \mathcal{R}} \frac{|y_r - y_{\text{true},r}|}{y_{\text{true},r}}, \end{aligned}$$

directly encoding the relative uncertainty of y from y_{true} . In other words, given a maximum uncertainty $0 \leq \delta \leq 1$, the state y_r must be in the continuous interval $[(1 - \delta)y_{\text{true},r}, (1 + \delta)y_{\text{true},r}]$ for all r .² In the forthcoming results, we will also use the parameter

$$B_\delta = \frac{1 + \delta}{1 - \delta}$$

to state certain equations more concisely. The following theorem presents a tractable linear program for computing the price of anarchy:

Theorem 9. *Consider a class of resource allocation games with $\mathcal{Z}_{w,u}$ for a given w and u . Additionally, let $\delta \in [0, 1)$ denote the limitations of the agents' knowledge. It holds*

²Note that we only consider the domain $\delta \in [0, 1)$, since if $\delta \geq 1$, the player valuation y_r^i can be arbitrarily close to 0 for any $y_{\text{true},r}$

that $\text{PoA}(\mathcal{G}_{w,u}^\delta) = 1/V^*$ where V^* is the optimal value of the following linear program:

$$\begin{aligned}
V^* = & \max_{\{\theta(a,x,b)\}_{a,x,b}} \sum_{a,x,b} w(b+x)\theta(a,x,b) \text{ s.t.} \\
& \sum_{a,x,b} \left[B_\delta a u(a+x) - b u(a+x+1) \right] \theta(a,x,b) \geq 0 \\
& \sum_{a,x,b} w(a+x)\theta(a,x,b) = 1 \\
& \theta(a,x,b) \geq 0 \quad \forall a,x,b
\end{aligned} \tag{5.15}$$

where $a, x, b \in \mathbb{N}$ such that $1 \leq a+x+b \leq n$.

Proof. First, we show that $\text{PoA}(\mathcal{G}_{w,u}^\delta)$ is lower bounded by $1/V^*$. Consider the reduced family of games $\mathcal{G}_{w,u}^{\delta,2} \subset \mathcal{G}_{w,u}^\delta$ where all the agents have only two actions; $\mathcal{A}_i = \{a_i^{\text{ne}}, a_i^{\text{opt}}\}$. In this reduced game, a^{ne} corresponds to a Nash equilibrium and a^{opt} corresponds to the action that maximizes the welfare. Note that $\text{PoA}(\mathcal{G}_{w,u}^\delta) = \text{PoA}(\mathcal{G}_{w,u}^{\delta,2})$ and that for any game $G \in \mathcal{G}_{w,u}^\delta$, uniformly scaling the values y_r such that $W(a^{\text{ne}}) = \sum_{r \in \mathcal{R}} y_r w(|a_r^{\text{ne}}|) = 1$ does not affect the price of anarchy. It follows that $\text{PoA}(\mathcal{G}_{w,u}^\delta) = 1/W^*$ where

$$\begin{aligned}
W^* = & \max_{G \in \mathcal{G}_{w,u}^{\delta,2}} W(a^{\text{opt}}) \\
\text{s.t. } & U_i(a^{\text{ne}}; \delta) \geq U_i(a_i^{\text{opt}}, a_{-i}^{\text{ne}}; \delta), \quad \forall i \in \mathcal{I} \\
& W(a^{\text{ne}}) = 1.
\end{aligned} \tag{5.16}$$

Observe that, as written, the above linear program is intractable as there are infinitely many games in $\mathcal{G}_{w,u}^{\delta,2}$. To reduce the complexity, we define a game parameterization based

on n partitions of the set of resources, defined as follows for each agent $i \in \mathcal{I}$:

$$\begin{aligned}\mathcal{R}_{a_i^{\text{ne}}} &= \{r \in \mathcal{R} : r \in a_i^{\text{ne}} \setminus a_i^{\text{opt}}\}, \\ \mathcal{R}_{a_i^{\text{opt}}} &= \{r \in \mathcal{R} : r \in a_i^{\text{opt}} \setminus a_i^{\text{ne}}\}, \\ \mathcal{R}_{a_i^{\text{opt}} \cap a_i^{\text{ne}}} &= \{r \in \mathcal{R} : r \in a_i^{\text{opt}} \cap a_i^{\text{ne}}\}, \\ \mathcal{R}_{a_i^{\emptyset}} &= \{r \in \mathcal{R} : r \notin a_i^{\text{opt}} \cup a_i^{\text{ne}}\}.\end{aligned}$$

Now consider an arbitrary game $G \in \mathcal{G}_{w,u}^{\delta,2}$ with resources \mathcal{R} , agent valuations y_i for each agent i and true resource values y_{true} . We can rewrite the Nash constraint in (5.16) for each agent $i \in \mathcal{I}$ as

$$\sum_{r \in a_i^{\text{ne}}} y_r^i \cdot u(|a_r^{\text{ne}}|) \geq \sum_{r \in a_i^{\text{opt}}} y_r^i \cdot u(|(a_i^{\text{opt}}, a_{-i}^{\text{ne}})_r|).$$

Under the partition defined for each agent i , we observe that the Nash condition can be rewritten as

$$\begin{aligned}& \sum_{r \in \mathcal{R}_{a_i^{\text{ne}}}} y_r^i \cdot u(|a_r^{\text{ne}}|) + \sum_{r \in \mathcal{R}_{a_i^{\text{opt}} \cap a_i^{\text{ne}}}} y_r^i \cdot u(|a_r^{\text{ne}}|) \\ & \geq \sum_{r \in \mathcal{R}_{a_i^{\text{opt}}}} y_r^i \cdot u(|(a_i^{\text{opt}}, a_{-i}^{\text{ne}})_r|) \\ & \quad + \sum_{r \in \mathcal{R}_{a_i^{\text{opt}} \cap a_i^{\text{ne}}}} y_r^i \cdot u(|(a_i^{\text{opt}}, a_{-i}^{\text{ne}})_r|).\end{aligned}$$

Canceling the terms in $\mathcal{R}_{a_i^{\text{opt}} \cap a_i^{\text{ne}}}$, we get

$$\sum_{r \in \mathcal{R}_{a_i^{\text{ne}}}} y_r^i \cdot u(|a_r^{\text{ne}}|) \geq \sum_{r \in \mathcal{R}_{a_i^{\text{opt}}}} y_r^i \cdot u(|(a_i^{\text{opt}}, a_{-i}^{\text{ne}})_r|).$$

Note that, for any resource $r \in \mathcal{R}$, it must hold that $y_r^i \in [(1 - \delta)y_{\text{true},r}, (1 + \delta)y_{\text{true},r}]$. Considering the Nash condition as written above, observe that the tightest constraint arises when $y_r^i = (1 + \delta)y_{\text{true},r}$ for all $r \in \mathcal{R}_{a_i^{\text{ne}}}$, $y_r^i = (1 - \delta)y_{\text{true},r}$ for all $r \in \mathcal{R}_{a_i^{\text{opt}}}$ and $y_r^i = y_{\text{true},r}$ for all other resources. This is the situation where agents overvalue the resources in their equilibrium actions and undervalue the resources in their optimal actions. Thus, we can consider this situation without loss of generality. For each resource $r \in \mathcal{R}$, we define the triplet $(a_r, x_r, b_r) \in \mathbb{N}^3$ as $a_r = |\{i \in \mathcal{I} : r \in \mathcal{R}_{a_i^{\text{ne}}}\}|$, $b_r = |\{i \in \mathcal{I} : r \in \mathcal{R}_{a_i^{\text{opt}}}\}|$, and $x_r = |\{i \in \mathcal{I} : r \in \mathcal{R}_{a_i^{\text{opt}} \cap a_i^{\text{ne}}}\}|$ where $1 \leq a_r + x_r + b_r \leq n$ must hold. Further, we define the map $\theta : \mathbb{N}^3 \rightarrow \mathbb{R}$ such that $\theta(a, x, b)$ is equal to the sum over true values $y_{\text{true},r}$ for all resources with $a_r = a$, $x_r = x$ and $b_r = b$, for all $a, x, b \in \mathbb{N}$ with $1 \leq a + x + b \leq n$. Under this notation, the following expressions hold:

$$W(a^{\text{opt}}) = \sum_{a,x,b} w(b+x)\theta(a,x,b),$$

$$W(a^{\text{ne}}) = \sum_{a,x,b} w(a+x)\theta(a,x,b).$$

We showed above that the tightest Nash condition arises when agents overvalue the resources that they select in their equilibrium actions and undervalue resources in their optimal actions. Under our parameterization, the sum over agents' utilities in this "tightest" scenario are as follows:

$$\sum_{i=1}^n U_i(a^{\text{ne}}; \delta) = \sum_{a,x,b} [(1 + \delta)a + x]u(a+x)\theta(a,x,b),$$

$$\sum_{i=1}^n U_i(a_i^{\text{opt}}, a_{-i}^{\text{ne}}; \delta) = \sum_{a,x,b} (1 - \delta)bu(a+x+1)\theta(a,x,b)$$

$$+ \sum_{a,x,b} xu(a+x)\theta(a,x,b).$$

Observe that if the equilibrium condition in (5.16) holds then the sum over agents' utilities at equilibrium must be greater than or equal to the sum over each of their utilities after they unilaterally deviate. The converse, however, need not hold in general. Thus, the linear program (5.15) in the claim represents a relaxation of the linear program (5.16).

Let V^* and W^* be the optimal values of linear programs (5.15) and (5.16), respectively. According to the proof thus far, we can only say that $V^* \geq W^*$ since V^* is the optimal value of a relaxed linear program, which means that $\text{PoA}(\mathcal{G}_{w,u}^\delta) \geq 1/V^*$. To show that $\text{PoA}(\mathcal{G}_{w,u}^\delta) \leq 1/V^*$ also holds, one can follow the approach outlined in [86]. \square

Following the reasoning detailed in [86], we can also define a linear program that computes the optimal utility design for the class of resource allocation games. Interestingly, we can directly import the techniques in [86] to achieve quite strong answers to questions about utility designs in limited information settings. For a given uncertainty δ and welfare characteristic w , we refer to the optimal utility mechanism as u_δ^{opt} .

Corollary 2. *Consider the class of resource allocation games $\mathcal{G}_{w,u}^\delta$ with n number of agents for a given $w \in \mathbb{R}_{>0}^n$. Additionally, let $\delta \in [0, 1)$ denote the uncertainty. The utility mechanism u_δ^{opt} that maximizes the price of anarchy is given as*

$$(u_\delta^{\text{opt}}, \mu^*) \in \arg \min_{u \in \mathbb{R}^n, \mu \in \mathbb{R}} \mu \quad \text{s.t.}$$

$$w(b+x) - \mu w(a+x) + B_\delta a u(a+x) - b u(a+x+1) \leq 0$$

$$\text{for all } a, x, b \in \mathbb{N} \text{ with } 1 \leq a+x+b \leq n$$

$$u(1) = 1$$

with $\text{PoA}(\mathcal{G}_{w,u_\delta^{\text{opt}}}^\delta) = \frac{1}{\mu^*}$.

In a realistic scenario, the system operator may not know the extent of the informational inconsistencies among the agents (i.e., the precise value of δ). In this case, what can be

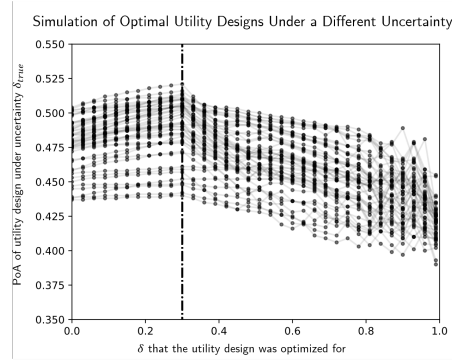


Figure 5.1: The price of anarchy is plotted for the optimal designs for various uncertainties δ under 30 randomly chosen welfare characteristics w and given $\delta_{\text{true}} = 0.3$ (indicated by the dotted line). Spanning all possible δ , it can be verified that indeed $u_{\delta_{\text{true}}}^{\text{opt}}$ designed for $\delta_{\text{true}} = .3$ performs optimally. However, more surprisingly, the degradation of performance (when u_{δ}^{opt} is designed for $\delta \neq .3$) as we move away from δ_{true} is slower on the left than on the right of $\delta = 0.3$. Nonintuitively, this suggests that by underestimating the value of δ_{true} we can achieve higher price of anarchy than overestimating it.

shown about the performance guarantees that a utility u_{δ}^{opt} – designed according to an assumed uncertainty δ – achieves under a realized uncertainty $\delta_{\text{true}} \neq \delta$? In other words, if there is a mismatch between the operator’s assumed δ and the realized δ , is there any loss of performance? In these situations, the following figure shows the quite surprising fact that underestimating the δ actually gives better performance guarantees. In Figure 5.1, we randomly generated 30 different welfare characteristics w , in which $w(j)$ is concave and non-decreasing in j . We assume that $\delta_{\text{true}} = .3$ and the game has $n = 10$ players. For a given w , we computed the optimal utility design $u_{\delta}^{\text{opt}} = \arg \max_u \text{PoA}(\mathcal{G}_{w,u}^{\delta})$ for each δ from $\delta = 0$ to 1. Then we plot the price of anarchy $\text{PoA}(\mathcal{G}_{w,u_{\delta}^{\text{opt}}}^{\delta_{\text{true}}})$ for each design. We see the performance degrades slower to the left of δ_{true} than to the right. In the next section, we formally capture this trend in the well-studied class of set covering games.

5.4 Set Covering Results

In this section, we restrict our analysis to *set covering games*, where the system welfare is the value of the resources covered, i.e. $w_{sc}(k) = 1$ for $k \geq 1$. Set covering games [35] are well studied and known to model a wide variety of practical applications, one of which is detailed in Example 4. We obtain an explicit characterization of the optimal utility design u_δ^{opt} for the class of set covering games and formally demonstrate the phenomenon observed in Figure 5.1 for this class of games.

To characterize the optimal utility design, we first outline the following Proposition to characterize the price of anarchy in set covering games with uncertainty.

Proposition 4. *For given utility u , uncertainty $0 \leq \delta \leq 1$, and number of agents $n \geq 1$, the price of anarchy for the class of set covering games $\mathcal{G}_{w_{sc},u}^\delta$ is*

$$\text{PoA}^{-1}(\mathcal{G}_{w_{sc},u}^\delta) = \max_{j \in [1, n-1]} \{ \max\{ B_\delta(j+1)u(j+1), \\ B_\delta j u(j+1) + 1, B_\delta j u(j) - u(j+1) + 1 \} \}$$

Proof. This proof is inspired by Theorem 3 in [43] with added consideration for the informational inconsistencies between the agents, and is included for completeness. First, we write the Lagrange dual of the linear program (5.15), i.e., $\text{PoA}(\mathcal{G}_{w_{sc},u}^\delta) = 1/\mu^*$ where

$$\begin{aligned} \mu^* &= \min_{\lambda \geq 0, \mu \in \mathbb{R}} \mu \text{ s.t.} \\ \mu w(a+x) &\geq w(b+x) + \lambda [B_\delta a u(a+x) - b u(a+x+1)], \\ \forall a, x, b \in \mathbb{N} \text{ s.t. } &1 \leq a+x+b \leq n. \end{aligned}$$

The rest of the proof involves removing redundant constraints to obtain a closed form expression of the price of anarchy. We first consider the constraints that arise from

$a = x = 0$ and $b \geq 1$. By definition, $w(k) = 1$ if $k \geq 1$ and 0 otherwise, giving the constraint $\lambda \geq \max_{b \in [n]} \frac{1}{b} = 1$. Considering the set of constraints that arise from $b, x = 0$ and $a \geq 1$ gives $\mu \geq \lambda B_\delta a u(a)$ for $a \in [n]$. Now evaluating the set of constraints that arise from $x = 0, a, b \geq 1$ gives

$$\begin{aligned} \mu &\geq \max_{a+b \in [2, n]} 1 + \lambda [B_\delta a u(a) - b u(a+1)] \\ &\geq \max_{a \in [1, n-1]} 1 + \lambda [B_\delta a u(a) - u(a+1)], \end{aligned}$$

which holds since $b = 1$ is the most binding constraint. When we consider the constraints that arise from $a, x, b \geq 1$, the resulting set of constraints can be written as

$$\begin{aligned} \mu &\geq \max_{a+b+x \in [3, n]} 1 + \lambda [B_\delta a u(a+x) - b u(a+x+1)] \\ &\geq \max_{a+x \in [2, n]} 1 + \lambda B_\delta a u(a+x) \\ &\geq \max_{a \in [1, n-1]} 1 + \lambda B_\delta a u(a+1), \end{aligned}$$

where $b = 0$ and $x = 1$ is the most binding constraint. Setting $b = 0$ is binding since it removes the negative term $-b u(a+x+1)$ from the expression. Setting $x = 1$ is binding, since for any pair $\{a, x\} \in [2, n]$, there is another pair $\{a+x-1, 1\}$ that results in stricter constraint. With the nonbinding constraints removed, the program reduces to

$$\begin{aligned} &\min_{\lambda \geq 1, \mu \in \mathbb{R}} \mu \text{ s.t.} \\ &\mu \geq \lambda B_\delta a u(a) && a \in [n] \\ &\mu \geq 1 + \lambda [B_\delta a u(a) - u(a+1)] && a \in [n-1] \\ &\mu \geq 1 + \lambda B_\delta a u(a+1) && a \in [n-1] \end{aligned}$$

The optimal dual variables have $\lambda = 1$, which comes from the tightest constraints. If we assume $a = 1$, this results in the set of constraints

$$\begin{aligned}\mu &\geq B_\delta 1u(1) = B_\delta \\ \mu &\geq 1 + B_\delta 1u(1) - u(2) = B_\delta + (1 - u(2)) \\ \mu &\geq 1 + B_\delta 1u(2) = 1 + B_\delta u(2).\end{aligned}$$

We can see the first constraint is always redundant, no matter if $u(2) \geq 1$ or $u(2) \leq 1$. The expression in the claim follows after removing the last nonbinding constraint and shifting the index from a to $a + 1$ for the first set of constraints. \square

Theorem 10. *For a given δ the optimal utility design u_δ^{opt} for the class of set covering games is*

$$u_\delta^{\text{opt}}(j) = \sum_{k=j}^{\infty} \frac{(j-1)!}{B_\delta^{k-j+1} (e^{\frac{1}{B_\delta}} - 1) k!} \quad (5.17)$$

and has corresponding price of anarchy

$$\text{PoA}(\mathcal{G}_{w_{sc}, u_\delta^{\text{opt}}}^\delta) = 1 - e^{-\frac{1}{B_\delta}}.$$

Proof. First we show that the price of anarchy is lower bounded by the proposed formula with the corresponding proposed f . We assume that the number of agents is n . From Proposition 4, we have that

$$\begin{aligned}\text{PoA}^{-1}(\mathcal{G}_{w_{sc}, u}^\delta) &\leq \mathcal{X} \text{ for any } \mathcal{X} \text{ s.t.} \\ \mathcal{X} &\geq B_\delta(n-1)u(n) + 1, \\ \mathcal{X} &\geq B_\delta j u(j) - u(j+1) + 1 \quad j \in [1, n-1]\end{aligned}$$

where we removed the first set of constraints, and all but the last one of the second constraints. An optimal utility design satisfies the set of inequalities with equality as follows

$$\mathcal{X} = B_\delta(n-1)u_\delta^{\text{opt}}(n) + 1 \quad (5.18)$$

$$\mathcal{X} = B_\delta j u_\delta^{\text{opt}}(j) - u_\delta^{\text{opt}}(j+1) + 1 \quad j \in [1, n-1]. \quad (5.19)$$

We can reformulate this system of equations as a recursive formula to generate the optimal utility design u_δ^{opt} as follows

$$u_\delta^{\text{opt}}(n) = 1 \quad (5.20)$$

$$u_\delta^{\text{opt}}(j) = \frac{u_\delta^{\text{opt}}(j+1)}{B_\delta j} + \frac{1}{j} u_\delta^{\text{opt}}(n)(n-1). \quad (5.21)$$

Iterating through this recursive equation and normalizing so that $u_\delta^{\text{opt}}(1) = 1$ gives

$$u_\delta^{\text{opt}}(j) = \frac{B_\delta^{j-1}(j-1)! \left(\frac{1}{B_\delta^n(n-1)(n-1)!} + \sum_{k=j}^{n-1} \frac{B_\delta^{-k}}{k!} \right)}{\frac{1}{B_\delta^n(n-1)(n-1)!} + \sum_{k=1}^{n-1} \frac{B_\delta^{-k}}{k!}}, \quad (5.22)$$

with a corresponding price of anarchy expression of

$$\text{PoA}(\mathcal{G}_{w_{sc}, u_\delta^{\text{opt}}}^\delta) \geq 1 - \frac{1}{\frac{1}{B_\delta^n(n-1)(n-1)!} + \sum_{k=0}^{n-1} \frac{B_\delta^{-k}}{k!}}.$$

Taking the limit as $n \rightarrow \infty$ and using the identity $\sum_{k=0}^{\infty} \frac{B_\delta^{-k}}{k!} = e^{\frac{1}{B_\delta}}$, we observe that u_δ^{opt} corresponds to the expression in (5.17) and $\text{PoA}(\mathcal{G}_{w_{sc}, u_\delta^{\text{opt}}}^\delta) \geq 1 - e^{-\frac{1}{B_\delta}}$.

For the upper bound, we construct an n agent worst case set covering game G^* inspired by [35]. All agents have two actions with $\mathcal{A}_i = \{a_i^{\text{ne}}, a_i^{\text{opt}}\}$, coinciding with their equilibrium and optimal actions. To state the allocations of a^{ne} and a^{opt} concisely, we

specify each resource with unique label $\ell : \mathcal{R} \rightarrow 2^n$ as follows. First we partition the resources into $n + 1$ groups, $\{\mathcal{R}_0, \dots, \mathcal{R}_n\}$. The true value of each resource $r \in \mathcal{R}_k$ is $y_{\text{true},r} = (B_\delta)^k$. There is one resource $r_0 \in \mathcal{R}_0$ with $\ell(r_0) = \{1\}$. For $k \geq 1$, the set of labels of the resources in \mathcal{R}_k is exactly the set $[2, n] \times {}^{n-1}P_{k-1}$, i.e., the set of permutations without $\{1\}$ as the first element. Therefore, there are $(n-1) \frac{(n-1)!}{(n-k)!}$ in \mathcal{R}_k . For any resource $r \in \mathcal{R}_k$ with $k \geq 1$, the last element of the label $\ell(r)$ denotes which agent selects the resource r in a^{opt} and $\{j \in \mathcal{I} : j \notin \ell(r)\}$ denotes the set of agents that select the resource r in a^{ne} . For the resource $r_0 \in \mathcal{R}_0$, agent 1 selects it in a^{opt} , and every agent selects it in a^{ne} . For example, if the label for the resource r is $\ell(r) = \{2, 3\}$ for a game with 4 agents, then it must be that $r \in \mathcal{R}_k$ with $y_{\text{true},r} = (B_\delta)^2$. Furthermore, agent 3 selects r in a^{opt} , while agents 1 and 4 select r in a^{ne} .

For any resource $r \in \mathcal{R}_k$ in G^* , $n - k$ agents select r in a^{ne} . Furthermore, for any agent $i \geq 2$ and $k \leq n$, the number of resources in \mathcal{R}_k that are selected in a_i^{opt} (denoted as $|\mathcal{R}_{i,k}^{\text{opt}}|$) and the number of resources in \mathcal{R}_{k-1} that are selected in a_i^{ne} (denoted as $|\mathcal{R}_{i,k-1}^{\text{ne}}|$) are equal. For agent 1 and $k \geq 2$, it holds that $|\mathcal{R}_{1,k}^{\text{opt}}| = |\mathcal{R}_{1,k-1}^{\text{ne}}|$. However, it is important to note that for agent 1, $|\mathcal{R}_{1,1}^{\text{opt}}| = 0$ while $|\mathcal{R}_{1,0}^{\text{opt}}| = |\mathcal{R}_{1,0}^{\text{ne}}| = 1$.

The agent valuations $y_{i,r}$ for the resources in G^* are as follows for a fixed δ uncertainty. If $r \in a_i^{\text{ne}}$, then agent i overvalues it to the extreme where $y_{i,r} = (1 + \delta)y_{\text{true},r}$ and if $r \in a_i^{\text{opt}}$, then agent i undervalues it to the extreme, where $y_{i,r} = (1 - \delta)y_{\text{true},r}$. The only exception to this is for agent 1 and the resource $r_0 \in \mathcal{R}_0$ where $y_{1,r_0} = y_{\text{true},r_0}$ since it is selected in both the optimal and equilibrium allocations by agent 1.

Now we can verify a^{ne} is indeed an equilibrium allocation. For any agent $i \in \mathcal{I}$, we

have

$$\begin{aligned}
U_i(a^{\text{ne}}; y_i) &= \sum_{k=0}^n \sum_{r \in \mathcal{R}_{i,k}^{\text{ne}}} y_r^i u(n-k) \\
&= \sum_{k=0}^n \sum_{r \in \mathcal{R}_{i,k}^{\text{ne}}} (1+\delta) y_{\text{true},r} u(n-k) \\
&= \sum_{k=0}^n |\mathcal{R}_{i,k}^{\text{ne}}| (1+\delta) (B_\delta)^k u(n-k) \\
&= \sum_{k=0}^n |\mathcal{R}_{i,k+1}^{\text{opt}}| (1-\delta) (B_\delta)^{k+1} u(n-(k+1)+1) \\
&= \sum_{k=0}^n |\mathcal{R}_{i,k}^{\text{opt}}| (1-\delta) y_{\text{true},r} u(n-k+1) \\
&= U_i(a_i^{\text{opt}}, a_{-i}^{\text{ne}}; y_i),
\end{aligned}$$

where we take advantage of the fact that no resources in \mathcal{R}_n are selected in a^{ne} (i.e., $|\mathcal{R}_{i,n}^{\text{ne}}| = 0$) in the fourth equality and that no agents $i \geq 1$ select the resource $r_0 \in \mathcal{R}_0$ in their optimal allocation (i.e., $|\mathcal{R}_{i,0}^{\text{opt}}| = 0$) for the fifth equality. We can use a similar argument for agent 1 with additional care taken for the resources in \mathcal{R}_0 and \mathcal{R}_1 . Its important to note that under any utility design u , the action a^{ne} is still an equilibrium and the allocations a^{ne} and a^{opt} do not change.

Under allocation a^{ne} in G^* , all resources in \mathcal{R}_k for $k \leq n-1$ are covered while, under the optimal allocation, all resources are covered. We can explicitly write the welfare at both allocations as

$$\begin{aligned}
W(a^{\text{ne}}) &= \sum_{r \in \mathcal{R}} y_{\text{true},r} w(|a_r^{\text{ne}}|) = 1 + \sum_{k=1}^{n-1} (n-1) \frac{B_\delta^k (n-1)!}{(n-k)!} \\
W(a^{\text{opt}}) &= \sum_{r \in \mathcal{R}} y_{\text{true},r} w(|a_r^{\text{opt}}|) = 1 + \sum_{k=1}^n (n-1) \frac{B_\delta^k (n-1)!}{(n-k)!}
\end{aligned}$$

Therefore, a lower bound on the price of anarchy is

$$\text{PoA}(G^*) \geq \frac{W(a^{\text{ne}})}{W(a^{\text{opt}})} = 1 - \frac{1}{\frac{1}{B_\delta^n (n-1)(n-1)!} + \sum_{k=0}^{n-1} \frac{B_\delta^{-k}}{k!}}.$$

Earlier, we showed that $\text{PoA}(\mathcal{G}_{w_{sc}, u}^\delta) \leq \text{PoA}(G^*)$ for any utility design. We just showed that $\text{PoA}(\mathcal{G}_{w_{sc}, u_\delta^{\text{opt}}}^\delta) \geq \text{PoA}(G^*)$. It follows that the utility u_δ^{opt} defined in (5.22) is optimal. Furthermore, taking the limit as $n \rightarrow \infty$ gives $\text{PoA}(\mathcal{G}_{w_{sc}, u_\delta^{\text{opt}}}^\delta) \leq \text{PoA}(G^*) = 1 - e^{-\frac{1}{B_\delta}}$. \square

Now that we have characterized the optimal utility design for set covering games, we can arrive at a closed form expression for the guarantees when there is a mismatch in uncertainty between the system operator and the realized uncertainty.

Proposition 5. *Let u_δ^{opt} be the optimal utility design for $0 \leq \delta \leq 1$ as in (5.17) and $0 \leq \delta_{\text{true}} \leq 1$ be the realized uncertainty. The price of anarchy is $\text{PoA}(\mathcal{G}_{w_{sc}, u_\delta^{\text{opt}}}^{\delta_{\text{true}}}) = V^{-1}$*

$$V = \begin{cases} (B_{\delta_{\text{true}}} B_\delta^{-1} - 1)u_\delta^{\text{opt}}(2) + \\ B_{\delta_{\text{true}}} B_\delta^{-1}(C - 1) + 1 & \text{if } \delta \leq \delta_{\text{true}}, \\ B_{\delta_{\text{true}}} B_{\delta_{\text{true}}}^{-1}(C - 1) + 1 & \text{if } \delta \geq \delta_{\text{true}}, \end{cases} \quad (5.23)$$

where $C = (e^{\frac{1}{B_\delta}} - 1)^{-1}$.

Proof. We first assume there are n agents and note that as $\delta \rightarrow 1$, the recursive formula in (5.20) and (5.21) for the optimal utility design, normalized to $u_\delta^{\text{opt}}(1) = 1$, gives $u_\delta^{\text{opt}}(j) = \frac{1}{j}$ for $j = 1, \dots, n-1$ and $u_\delta^{\text{opt}}(n) = \frac{1}{n-1}$. Additionally, observe that as δ increases, $u_\delta^{\text{opt}}(j)$ increases for any j , since the recursive formula in (5.21) produces a slower increasing sequence for a higher δ , so normalizing to $u_\delta^{\text{opt}}(1) = 1$ gives a larger $u_\delta^{\text{opt}}(j)$. Thus $u_\delta^{\text{opt}}(j) \leq \frac{1}{j}$ for $j = 1, \dots, n-1$ for any δ . Note that based on the recursive formula in (5.21), $u_\delta^{\text{opt}}(j)$ is decreasing in j for any δ . By Proposition 4, the

$\text{PoA}(\mathcal{G}_{w_{sc}, u_{\delta}^{\text{opt}}}^{\delta_{\text{true}}})^{-1} = \mathcal{X}$ where \mathcal{X} is the lowest value satisfying

$$\mathcal{X} \geq B_{\delta_{\text{true}}}(j+1)u_{\delta}^{\text{opt}}(j+1) \quad j \in [1, n-1] \quad (5.24)$$

$$\mathcal{X} \geq B_{\delta_{\text{true}}}ju_{\delta}^{\text{opt}}(j+1) + 1 \quad j \in [1, n-1] \quad (5.25)$$

$$\mathcal{X} \geq B_{\delta_{\text{true}}}ju_{\delta}^{\text{opt}}(j) - u_{\delta}^{\text{opt}}(j+1) + 1 \quad j \in [1, n-1] \quad (5.26)$$

Now the redundant inequalities are eliminated to derive a closed form expression. For the inequalities in (5.24), we have that

$$\begin{aligned} B_{\delta_{\text{true}}}(j+1)u_{\delta}^{\text{opt}}(j+1) &\leq B_{\delta_{\text{true}}} \\ &\leq B_{\delta_{\text{true}}} - u_{\delta}^{\text{opt}}(2) + 1 \quad j \in [1, n-2], \end{aligned}$$

where the first inequality comes from $u_{\delta}^{\text{opt}}(j+1) \leq 1/(j+1)$ and the second inequality comes from $u_{\delta}^{\text{opt}}(2) \leq u_{\delta}^{\text{opt}}(1) = 1$. Note that putting $j = 1$ in the last set of inequalities (5.26) gives the last term. For $j = n-1$,

$$\begin{aligned} B_{\delta_{\text{true}}}nu_{\delta}^{\text{opt}}(n) &= B_{\delta_{\text{true}}}(n-1)u_{\delta}^{\text{opt}}(n) + B_{\delta_{\text{true}}}u_{\delta}^{\text{opt}}(n) \\ &\leq B_{\delta_{\text{true}}}(n-1)u_{\delta}^{\text{opt}}(n) + 1 \\ &\leq B_{\delta_{\text{true}}}B_{\delta}^{-1}(C-1) + 1, \end{aligned} \quad (5.27)$$

where the first inequality comes from the fact that $B_{\delta_{\text{true}}}u_{\delta}^{\text{opt}}(n) \leq n$ for a high enough n and the second inequality comes from the substitution $C-1 = B_{\delta}(n-1)u_{\delta}^{\text{opt}}(n)$ from Equation (5.18). Note that this corresponds to putting $j = n-1$ in the second set of inequalities (5.25). Therefore, we have shown the first set of inequalities is redundant.

For the inequalities in (5.25), we have that for $j \in [1, n - 2]$,

$$\begin{aligned}
B_{\delta_{\text{true}}} j u_{\delta}^{\text{opt}}(j + 1) + 1 &= B_{\delta_{\text{true}}}(j + 1) u_{\delta}^{\text{opt}}(j + 1) \\
&\quad - B_{\delta_{\text{true}}} u_{\delta}^{\text{opt}}(j + 1) + 1 \\
&\leq B_{\delta_{\text{true}}}(j + 1) u_{\delta}^{\text{opt}}(j + 1) \\
&\quad - u_{\delta}^{\text{opt}}(j + 2) + 1,
\end{aligned}$$

where the first inequality comes from the fact that $B_{\delta_{\text{true}}} u_{\delta}^{\text{opt}}(j + 1) \geq 1 \cdot u_{\delta}^{\text{opt}}(j + 2)$. Note that this expression matches the inequalities in (5.26) for $j \in [1, n - 2]$ and therefore are redundant.

We can also reduce the inequalities in (5.26):

$$\begin{aligned}
&B_{\delta_{\text{true}}} j u_{\delta}^{\text{opt}}(j) - u_{\delta}^{\text{opt}}(j + 1) + 1 = \\
&(B_{\delta_{\text{true}}} B_{\delta}^{-1} - 1) u_{\delta}^{\text{opt}}(j + 1) + (B_{\delta_{\text{true}}} B_{\delta}^{-1})(C - 1) + 1
\end{aligned}$$

where $C = \text{PoA}(G_{w_{sc}, u_{\delta}^{\text{opt}}}^{\delta})^{-1}$. The equality comes from the recursive formula in (5.19) with substitution $B_{\delta} j u_{\delta}^{\text{opt}}(j) = u_{\delta}^{\text{opt}}(j + 1) + C - 1$. If $\delta \leq \delta_{\text{true}}$, then $B_{\delta_{\text{true}}} B_{\delta}^{-1} - 1 \geq 0$ and the binding constraint comes from taking $j = 1$,

$$(B_{\delta_{\text{true}}} B_{\delta}^{-1} - 1) u_{\delta}^{\text{opt}}(2) + (B_{\delta_{\text{true}}} B_{\delta}^{-1})(C - 1) + 1.$$

Conversely if $\delta \geq \delta_{\text{true}}$, the binding constraint comes from $j = n - 1$,

$$\begin{aligned}
&(B_{\delta_{\text{true}}} B_{\delta}^{-1} - 1) u_{\delta}^{\text{opt}}(n) + (B_{\delta_{\text{true}}} B_{\delta}^{-1})(C - 1) + 1 \\
&\leq B_{\delta_{\text{true}}} B_{\delta}^{-1}(C - 1) + 1,
\end{aligned}$$




Figure 5.2: We plot the price of anarchy achieved by the utility rules designed for $\delta \in [0, 1]$ within three classes of set covering games corresponding with $\delta_{\text{true}} = 0.2$ (in blue), $\delta_{\text{true}} = 0.3$ (in red), and $\delta_{\text{true}} = 0.4$ (in green). The explicit expression for such curves is provided in (5.23). Observe that underestimating the true level of uncertainty in the class of games gives better price of anarchy guarantees than overestimating it, a trend that we previously noted from the simulation results in Figure 5.1.

where the inequality comes from $(B_{\delta_{\text{true}}} B_{\delta}^{-1} - 1) \leq 0$ for $\delta \geq \delta_{\text{true}}$. Note that this constraint is subsumed by the one in (5.27).

Finally, the resulting set of inequalities is

$$\begin{aligned} \mathcal{X} &\geq B_{\delta_{\text{true}}} B_{\delta_{\text{true}}}^{-1} (C - 1) + 1 \\ \mathcal{X} &\geq (B_{\delta_{\text{true}}} B_{\delta}^{-1} - 1) u_{\delta}^{\text{opt}}(2) + (B_{\delta_{\text{true}}} B_{\delta}^{-1})(C - 1) + 1. \end{aligned}$$

We showed that the first constraint is strictest if $\delta \geq \delta_{\text{true}}$ and that the second constraint is strictest when $\delta < \delta_{\text{true}}$. Taking $n \rightarrow \infty$, we have from Theorem 10 that $C - 1 = \text{PoA}(\mathbf{G}_{w_{sc}, u_{\delta}^{\text{opt}}}^{\delta})^{-1} - 1 = (1 - e^{-\frac{1}{B_{\delta}}})^{-1} - 1 = (e^{\frac{1}{B_{\delta}}} - 1)^{-1}$. \square

In the context of set covering games with uncertainty in the state of the resources, we have formally shown the surprising fact that underestimating these values actually has better performance guarantees than overestimating.

Chapter 6

Nonconvergent Learning Dynamics

Understanding emergent behavior of multi-agent systems is important for various distribution applications, and game theory has been identified as a useful framework for this purpose. The core of game-theoretic analysis is through representing agents as players in a game, which enables prediction of emergent behavior through Nash equilibria. This methodology offers a valuable perspective, where system behavior can be characterized by assuming that self-interested decision-making processes lead to Nash equilibrium without identifying those decision-making processes explicitly. This approach has found applications across diverse domains, such as sensor coverage, traffic networks, auctions, and network coordination. However, guarantees of existence in Nash equilibrium is not universal across problem settings, prompting an exploration into alternative equilibria.

In this chapter, we will instead focus on sink equilibria, which are defined as the attractors of decision-making processes where Nash equilibrium may not be naturally defined. By classifying system outcomes through a global objective function, we can analyze the resulting approximation guarantees that sink equilibria have for a given game. Our main result in this chapter is an approximation guarantee on the sink equilibria through defining an introduced metric of misalignment, which captures how uniform

agents are in their self-interested decision making. Overall, sink equilibria are naturally occurring in many multi-agent contexts, and we display our results on their quality with respect to two practical problem settings. The results and discussion in this chapter is based on the work presented in [87].

6.1 Introduction

In this chapter, we specifically consider a *game theoretic* approach, where the emergent properties of the multi-agent system are studied using tools from game theory. The main idea of this approach is to model agents as self-interested decision makers, where each agent's preference over the collective system outcome is designated through a *utility function*. The agents are then presumed to undergo a *best (or better) response process*, where each agent updates their decision in a self-interested manner to maximize their individual utility. The emergent system outcomes from this process are traditionally expected to be *Nash equilibrium*, which can be expressed as the limit points of the best response process. Thus, many previous works study the emergent system behavior through characterizing properties of the Nash equilibrium; this has been done in many different contexts, such as traffic systems, power networks, etc. [88, 89].

But is it reasonable to expect agents to converge to Nash equilibria from best response processes? In the class of *potential games* [90] and variants [91], best response processes are indeed guaranteed to converge to Nash equilibria [92]. In potential games, agents are fully cooperative, and self-interested decisions made by agents always lead to improvements in a given global objective. However, a variety of natural multi-agent settings fall outside of this class. The system may display competitive interactions; for e.g., business firms may have competing economic interests. In social systems, agents' may be inherently misaligned in their preferences; for e.g., drivers may have different

sensitivities to tolls. Even for multi-agent systems that are fully engineered, there are operational concerns such as prediction errors or informational privacy that must be accounted for. All of these scenarios do not exhibit a potential game structure, and thus guarantees of convergence to Nash equilibria can not be established. We describe two examples of this kind in Example 6 and 7. **In these instances, can the emergent system outcomes still be characterized, where either Nash equilibrium do not exist or best response processes does not converge to Nash equilibria?** This is the main focus of the chapter.

We utilize *sink equilibria* in [22] as an alternate solution concept ¹ to address this, which are specified as the attractors of the best response process. By definition, sink equilibria are well defined and have guarantees of convergence for any given game. Thus, we analyze the behavior of sink equilibria in this chapter. Specifically, we assume that system outcomes are evaluated through a given global objective function. The utility function of each agent may not be aligned with this global objective. In these settings, we can characterize performance guarantees of the induced sink equilibria with respect to the global objective.

To the author’s knowledge, a general approach to studying performance guarantees of sink equilibrium has not been done previously. While sink equilibria have not been studied in as much detail as Nash equilibria, we still highlight an pertinent selection of past literature on sink equilibria. The seminal work in [22] first established the concept of sink equilibria in a game-theoretic context, as well as provide negative results of sink equilibria in valid utility games. Positive results on sink equilibria behavior were shown in [94], where it was seen that sink equilibria perform much better than mixed Nash

¹A popular alternative to Nash equilibria are *coarse correlated equilibria* which have existence and convergence guarantees [93]. However, these dynamics requires full knowledge of the history of decisions of all of the agents. Additionally, only empirical distribution of play is guaranteed to converge to coarse correlated equilibria, which may not correspond to the actualized decisions made by the agents.

equilibria. Analysis of sink equilibria under the name of curb sets was done in [95]. Several complexity results of sink equilibria were introduced in [96]. Sink equilibria were extended in continuous domains in [24]. Recently, design of sink equilibria selection algorithms was done in [97]. While the literature on sink equilibrium is sparse, they naturally emerge when agent utility functions are not aligned perfectly.

6.2 Model and Preliminaries

Consider a general multi-agent scenario with n agents $\mathcal{I} = \{1, \dots, n\}$, where each agent is endowed with a finite decision or action set \mathcal{A}_i . We denote an action as $a_i \in \mathcal{A}_i$, and a joint action profile as $a \in \mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$. The quality of each joint action profile is evaluated with a global objective function $W : \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}$ that characterizes the total system welfare. In other words, the optimal joint decision of the system is described as

$$a^{\text{opt}} \in \arg \max_{a \in \mathcal{A}} W(a). \quad (6.1)$$

We assume that solving the decision problem in Eq. (6.1) cannot be done in a centralized fashion, due to computational, informational, administrative, or communication concerns. Therefore, we assume that each agent selects their decisions in a distributed manner. We assume that each agent i is endowed with a *utility function* $U_i : \mathcal{A} \rightarrow \mathbb{R}$ to classify their preferences over their decision set, resulting in the game tuple $G \triangleq (\mathcal{I}, \mathcal{A}, W, \{U_i\}_{i \in \mathcal{I}})$. Moreover, we assume the utility function U_i depends on the welfare function W , either naturally or by design. In particular, consider a completely cooperative scenario, where the agents exhibit the *common interest* utility design. Here, the utility functions are

completely aligned with the global objective W ², where

$$U_i(a) \equiv W(a) \text{ for all } a \in \mathcal{A} \text{ and } i \in \mathcal{I}. \quad (6.2)$$

Under the common interest utility, the emergent joint decisions coincide with the set of Nash equilibrium NE of the game. A joint action a^{ne} is considered a Nash equilibrium if the following inequality holds

$$U_i(a^{\text{ne}}) \geq U_i(a_i, a_{-i}^{\text{ne}}) \text{ for all } a \in \mathcal{A} \text{ and } i \in \mathcal{I}, \quad (6.3)$$

where $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$ corresponds to the joint action without the decision of agent i . In completely cooperative settings, where they are guaranteed to exist, we can then quantify the emergent behavior through the qualities of the possible resulting Nash equilibrium. This is done through the metric of *price of anarchy* which is defined as

$$\text{PoA}(\text{G}) = \frac{\min_{a^{\text{ne}} \in \text{NE}} W(a^{\text{ne}})}{\max_{a \in \mathcal{A}} W(a)}, \quad (6.4)$$

where we take the worst case ratio of the welfare of Nash equilibria over the optimal welfare. Price of anarchy is a well-studied metric, with many results on its characterization in the literature [43, 44, 88]. Thus a standardized approach can be implemented to characterize system behavior in completely cooperative scenarios.

However, due to operational or design constraints, assuming a common interest utility design may not be feasible (see Examples 6 and 7). Throughout the paper, we then allow $U_i(a) \neq W(a)$ to be misaligned, and focus our attention to sink equilibrium as our standard solution concept. To define sink equilibrium, we first outline the *best response*

²We can relax this constraint to consider utility functions that are *preference equivalent* to the welfare function, where the ordering of preferences over joint actions is maintained. This is indeed the case for potential and weighted-potential games.

process. Under this process, the best decision set for agent i assuming all other agents' decisions are fixed is known as the *best response set*, that is,

$$\text{Br}_i(a) = \arg \max_{\bar{a}_i \in \mathcal{A}_i} U_i(\bar{a}_i, a_{-i}). \quad (6.5)$$

Then for every step, a randomly selected agent picks an action from its best response set uniformly. This induces the following Markovian dynamics on the set of joint actions, describing the best response process, as

$$\Pr(\tilde{a}|a) = \begin{cases} \frac{1}{n \cdot |\text{Br}_i(a)|} & \text{if } \tilde{a} \in (\text{Br}_i(a), a_{-i}) \text{ for some } i \in \mathcal{I} \\ 0 & \text{otherwise,} \end{cases} \quad (6.6)$$

where $\Pr(\tilde{a}|a)$ represents the probability of reaching the joint action \tilde{a} from a in the Markov chain. Note that there is an equal chance for each player i to perform a best response at each time step. We refer to a probability distribution over the joint action set as $\sigma \in \Delta\mathcal{A}$, where $p_\sigma(a)$ denotes the probability of sampling a under σ . We also say that an action $a \in \text{supp}(\sigma) \subseteq \mathcal{A}$ is in the support of σ if the probability $p_\sigma(a) > 0$ is strictly positive. Furthermore, we say that σ is a *stationary distribution* of the Markov chain if the equality $p_\sigma(a) = \mathbb{E}_{\bar{a} \sim \sigma}[\Pr(\tilde{a}|\bar{a}) \cdot p_\sigma(\bar{a})]_a$ for all $a \in \mathcal{A}$ holds. We also specify the *sink strongly connected components* of the Markov chain defined in Eq. (6.6). A set $S \subseteq \mathcal{A}$ is a sink strongly connected component if there exists a path of positive probability from a to \bar{a} under (6.6) for any $a, \bar{a} \in S$ and there are no transitions from S to outside of S . Formally, if $a, \bar{a} \in S$, then there exists a sequence a_0, a_1, \dots, a_m where $\Pr(a_{j+1}|a_j) > 0$ for all j with $a_0 = a$ and $a_m = \bar{a}$. Additionally, if $a \in S$ and $\bar{a} \notin S$, then no such sequence exists.

Definition 4. A probability distribution $\sigma \in \Delta\mathcal{A}$ is a sink equilibrium of the game G if

σ is a stationary distribution of the Markov chain in Eq. (6.6) and if $\text{supp}(\sigma) = S$ is a sink strongly connected component.

In other words, sink equilibria are defined as the attractors of the best response process. Given a game G , we characterize the behavior of the sink equilibria in a similar manner to Eq. (6.4) through the metric of *price of sinking* [22] as

$$\text{PoSE}(G) = \min_{\sigma \in \text{SE}} \mathbb{E}_{a \sim \sigma} [W(a)] / W(a^{\text{opt}}), \quad (6.7)$$

where SE denotes the total set of sink equilibrium of the game G . We note that under a common interest utility design, the set of Nash equilibrium $\text{NE} \simeq \text{SE}$ is equivalent to the set of sink equilibrium.

We now examine the sink equilibria on two illustrative examples in which they naturally emerge. In these examples, we derive guarantees on the sink equilibria. The examples are as follows.

Example 6 (Ecological Monitoring). Ecological monitoring is necessary to understand the well-being of inhabited populations as well as track health of overall ecology. While this can be handled by field ecologists, autonomous agents can supplement or even act as substitutes to gather important ecological data - as was done by the authors in [98]. In this scenario, a high level control objective of the agents is to understand how to orient themselves to monitor the region of interest as best as possible. We can model this as a covering problem [99]. In this way, let $\mathcal{R} = \{r_1, \dots, r_m\}$ define the region monitored by n agents, where we have finitely partitioned the region into possibly arbitrary segments. For each segment r , the importance of monitoring that segment can be associated with a value $v_r \in \mathbb{R}_{\geq 0}$ which defines the intrinsic quality of data that can be collected in that segment. This can be affected by the number or magnitude of populations in the segment, relevant climactic conditions, etc. These parameters are never known before hand, and

thus must be estimated by the agents in the field. Thus each agent has its noisy estimate of the value v_r^i which we assume is drawn from a normal distribution $\mathcal{N}[v_r + c, (d \cdot v_r)^2]$ with bias c and variance $(d \cdot v_r)^2$. Each agent can decide which subset of region to monitor (i.e. $a_i \subset \mathcal{R}$), which depend on its sensor and motor capabilities. Thus, collectively, the goal of the agents is to monitor the most and highest valued portions, as captured by the welfare function below.

$$W(a) = \sum_{r \in \cup_i a_i} v_r \quad (6.8)$$

The objective that each agent witnesses, however is based on their estimate, or that their utility is $U_i(a) = \sum_{r \in \cup_i a_i} v_r^i$, which is potentially different for each agent. Thus, when running a best response algorithm, the agents converge to a sink equilibrium which may not be a Nash equilibrium. We derive a guarantee on price of sinking in Proposition 6.

Proposition 6. *Consider the problem defined above. The expected price of sinking is lower bounded by the following expression*

$$\mathbb{E}[\text{PoSE}(\text{G})] \geq \max\left(\frac{1 - 4n\beta_\Phi}{2}, 0\right), \quad (6.9)$$

where Φ is the normal cumulative distribution function and β_Φ is defined as

$$\beta_\Phi = |\mathcal{R}| \left(d\sqrt{\frac{2}{\pi}} e^{-\frac{c^2}{2d^2}} + c(1 - 2\Phi(-c/d)) \right). \quad (6.10)$$

Example 7 (Radio Signalling). Consider the situation in which n agents have to communicate to each other through k communication channels, as seen in [100]. However, when more than one agent selects a channel to communicate under, the signals experience interference. This can be captured by the parameter $w_{ij} \geq 0$, which dictates the interference between agent i and agent j . The system as a whole, would like to minimize the total interference experienced between all the agents. In turn, the system welfare can be

defined as

$$W(a) = \sum_i \sum_{j: a_j \neq a_i} w_{ij}, \quad (6.11)$$

where $a_i \in \{1, \dots, k\}$ is the channel that agent i decides to transmit their messages on. However, these interference parameters may not be known to the agents and have to be estimated. For simplicity, we can assume that the agents have a margin of error of α or that agent i 's estimate is $w_{ij}^i \in [\alpha \cdot w_{ij}, \alpha^{-1} \cdot w_{ij}]$. Again, when agents run a best response algorithm, they are not guaranteed to converge to a Nash equilibrium due to the informational constraints. Then we can characterize the guarantee on the price of sinking in Proposition 7.

Proposition 7. *Consider the problem defined above with two channels ($k = 2$). The price of sinking is lower bounded by the following expression*

$$\text{PoSE}(\mathcal{G}) \geq \frac{1}{3\alpha^2 + (1 - \alpha^2)n}. \quad (6.12)$$

6.3 Main Results

The main results of this paper are on providing lower bounds for the price of sinking for a given game. To do this, we first recall the notion of *smoothness* [101] as a useful analysis tool for the price of anarchy. In this paper, we use a relaxed version of smoothness to classify a given game. We say that a game is (λ, μ) -*smooth* if, for a fixed $\mu \geq \lambda \geq 0$, we have that

$$\sum_{i \in \mathcal{I}} (U_i(a) - U_i(a_i^{\text{opt}}, a_{-i})) \leq \mu W(a) - \lambda W(a^{\text{opt}}) \quad (6.13)$$

for all actions $a \in \mathcal{A}$. Given these parameters, the efficiency of Nash equilibria for a given game can easily be determined. This is described in Proposition 8, where a proof is included for completeness.

Proposition 8. *Let G be a (λ, μ) -smooth game. Then the price of anarchy is lower bounded by $\text{PoA}(G) \geq \frac{\lambda}{\mu}$.*

Proof. From applying the definition of Nash equilibrium repeatedly for all $i \in \mathcal{I}$ with respect to the deviation a^{opt} , we have the following inequality

$$\sum_i (U_i(a^{\text{ne}}) - U_i(a_i^{\text{opt}}, a_{-i}^{\text{ne}})) \geq 0.$$

Notice that now we can directly substitute the inequality in Eq. (6.13) to get that

$$\mu W(a^{\text{ne}}) - \lambda W(a^{\text{opt}}) \geq 0.$$

Since a^{ne} is any arbitrary Nash equilibrium, we can rearrange the above equation to get that $\text{PoA}(G) \geq \frac{\lambda}{\mu}$ to show the claim. \square

We note that there always exist some λ and μ such that the game is (λ, μ) -smooth, as $\lambda \rightarrow 0$ and $\mu \rightarrow \infty$ will always satisfy the inequality in Eq. (6.13). The main analytical benefit of smoothness analysis is that instead of searching across the set of Nash equilibrium directly, we can instead characterize the price of anarchy through a bi-variable optimization problem (over λ and μ). This can be done in various game-theoretic contexts [101]. The optimization problem is written formally below.

Corollary 3. *The price of anarchy for a given game G is lower bounded by*

$$\text{PoA}(G) \geq \sup_{\mu \geq \lambda \geq 0} \left\{ \frac{\lambda}{\mu} : G \text{ is } (\lambda, \mu) \text{ - smooth} \right\}.$$

However, unlike Nash equilibria, when applying the smoothness inequality directly to the analysis of sink equilibria, it is not possible to get non-trivial guarantees. For

all valid smoothness parameters, it is possible to construct a corresponding game with trivial performance guarantees on sink equilibria, as stated below.

Proposition 9. *For every $\mu > \lambda \geq 0$, there exists a (λ, μ) -smooth game G with a unique sink equilibrium such that the price of sinking is $\text{PoSE}(G) = 0$.*

Proof. Let $\mu > \lambda \geq 0$. Consider the following game G with two agents with the action sets $\mathcal{A}_1 = \{e_1, e_2, e_3\}$ and $\mathcal{A}_2 = \{f_1, f_2, f_3\}$. We define the welfare values $W(a)$ for each joint action through Table 6.1 below. Similarly, we can define the utility values $U_i(a)$ for

	f_1	f_2	f_3
e_1	1	$(\lambda + \varepsilon)/\mu$	0
e_2	$(\lambda + \varepsilon)/\mu$	0	0
e_3	0	0	0

Table 6.1: Welfare $W(a)$ for each joint action $a = (e_i, f_j)$.

each joint action and for each agent in Table 6.2.

	f_1	f_2	f_3
e_1	$(0, 0)$	$(0, \varepsilon)$	$(0, -\varepsilon)$
e_2	$(\varepsilon, 0)$	$(\lambda, -2\lambda)$	$(-2\lambda, \lambda)$
e_3	$(-\varepsilon, 0)$	$(-2\lambda, \lambda)$	$(\lambda, -2\lambda)$

Table 6.2: Welfare $(U_1(a), U_2(a))$ for each joint action a .

We can choose $\varepsilon = (\mu - \lambda)/2 > 0$ such that the optimal joint action is $a^{\text{opt}} = (e_1, f_1)$ with an optimal welfare of $W(a^{\text{opt}}) = 1$. Under the best response dynamics, observe that the set $\{(e_2, f_2), (e_3, f_2), (e_2, f_3), (e_3, f_3)\}$ is the unique strongly connected component. It can be verified that each joint action a satisfies the smoothness condition in Eq. (6.13). Since the welfare of each action in the unique strongly connected component is 0, the price of sinking can be upper bounded by $\text{PoSE}(G) = \mathbb{E}_{a \in \sigma}[W(a)] \leq \sum_{a \in \text{supp}(\sigma)} W(a) = 0$ for the unique sink equilibrium σ . \square

This negative result is similar in spirit to the one presented in [22, Lemma 3.2]. However, we emphasize that the inferior guarantees are more indicative of inefficacy of

a direct approach rather than the intrinsic behavior of sink equilibria. This sentiment is also reflected in [94], where in certain game settings, it is shown that the quality of sink equilibria is arbitrarily better than the quality of any mixed equilibria. In fact, if we consider games with added structure, we can arrive at nontrivial guarantees on sink equilibria.

Therefore, we consider games in which the deviation from the common interest utility $U_i(a) \neq W(a)$ is bounded. We encapsulate the extent of the deviation through the constant $\beta \in [0, 1]$, where $\beta = 0$ signifies no deviation from the common interest utility and $\beta = 1$ signifies the maximum deviation. We define this formally below ³.

Definition 5. A game G is considered to be *β -arithmetically misaligned* if

$$|U_i(a) - W(a)| \leq \beta W(a), \quad (6.14)$$

or *β -geometrically misaligned* if

$$1 - \beta \leq \frac{U_i(a)}{W(a)} \leq \frac{1}{1 - \beta}, \quad (6.15)$$

is satisfied for all actions $a \in \mathcal{A}$ and agents $i \in \mathcal{I}$.

We note that when $\beta = 0$, under the common interest utility, the sink equilibria are equivalent to the Nash equilibria and inherit the price of anarchy guarantees coming from smoothness analysis. Likewise, we will see that the sink equilibria in near-common interest games with β close to 0 inherit similar guarantees dictated by the common interest utility. This observation is also reflected in a different context in [102]. In this vein, let λ_c and μ_c be the parameters that satisfy the smoothness inequality in Eq. (6.13)

³We can generalize the results to instead consider alignment to a potential function. In this way, β characterizes the closeness of the game to a potential game. Near-potential games have been studied in [23].

for the common interest utility

$$\sum_{i \in \mathcal{I}} (W(a) - W(a_i^{\text{opt}}, a_{-i})) \leq \mu_c W(a) - \lambda_c W(a^{\text{opt}}), \quad (6.16)$$

where we have substituted $U_i(a) \equiv W(a)$. With this, we can state the main result of the paper.

Theorem 11. *Let G be a game such that the best response $Br_i(a)$ is always singular valued. Let λ_c and μ_c satisfy Eq. (6.16) for all $a \in \mathcal{A}$. If the game is β -arithmetically misaligned, as in Eq. (6.14), then the price of sinking satisfies*

$$\text{PoSE}(G) \geq \max\left(\frac{\lambda_c - 4\beta n}{\mu_c}, 0\right). \quad (6.17)$$

If the game is β -geometrically misaligned, as in Eq. (6.15), then the price of sinking satisfies

$$\text{PoSE}(G) \geq \frac{\lambda_c}{(1 - \beta)^2 \mu_c + (1 - (1 - \beta)^2) n}. \quad (6.18)$$

Proof. First, we introduce the following lemma to characterize the sink equilibria in an alternative fashion.

Lemma 7. *Let G be a game such that $Br_i(a)$ is always singular valued and let $\sigma \in \text{SE}$ be any sink equilibrium of the game. For any function $g : \mathcal{A} \rightarrow \mathbb{R}$, the following equality must hold*

$$\mathbb{E}_{a \sim \sigma} \left[\sum_{i \in \mathcal{I}} g(a) - g(Br_i(a), a_{-i}) \right] = 0. \quad (6.19)$$

Proof. Let σ be a sink equilibrium of the game G . Since the sink equilibrium is a stationary distribution under the dynamics outlined in Eq. (6.6), we have that $p_\sigma(a) =$

$\sum_{\bar{a} \in \mathcal{A}} \Pr(a|\bar{a})p_\sigma(\bar{a})$. Under this statement, we have the series of equalities below

$$\begin{aligned} n \sum_a p_\sigma(a)g(a) &= n \sum_a \sum_{\bar{a}} \Pr(a|\bar{a})p_\sigma(\bar{a})g(a) \\ \mathbb{E}_{a \sim \sigma} \left[\sum_{i \in \mathcal{I}} g(a) \right] &= \mathbb{E}_{\bar{a} \sim \sigma} \left[n \sum_a \Pr(a|\bar{a})g(a) \right] \\ &= \mathbb{E}_{a \sim \sigma} \left[\sum_{i \in \mathcal{I}} g(\text{Br}_i(a), a_{-i}) \right], \end{aligned}$$

where we change the the naming convention from \bar{a} to a in the last line. Rearranging the terms and using linearity of expectation gives us the claim. \square

Proof of Arithmetic. For ease of notation, let $a_i^{\text{br}} = (\text{Br}_i(a), a_{-i})$. We can apply Lemma 7 with respect to the welfare function W to get

$$\mathbb{E}_{a \sim \sigma} \left[\sum_{i \in \mathcal{I}} W(a) - W(a_i^{\text{br}}) \right] = 0. \quad (6.20)$$

Since we assume the game is β -arithmetically misaligned, we have that $U_i(a) \geq (1 - \beta)W(a) \geq W(a) - \beta W(a^{\text{opt}})$. Likewise, we can also bound $U_i(a_i^{\text{br}}) \leq W(a_i^{\text{br}}) + \beta W(a^{\text{opt}})$. We can substitute these two inequalities in Eq. (6.20) to get

$$W(a) - W(a_i^{\text{br}}) \leq U_i(a) - U_i(a_i^{\text{br}}) + 2\beta W(a^{\text{opt}}) \quad (6.21)$$

We can apply this inequality to Eq. (6.20) for

$$\mathbb{E}_{a \sim \sigma} \left[2\beta n W(a^{\text{opt}}) + \sum_{i \in \mathcal{I}} U_i(a) - U_i(a_i^{\text{br}}) \right] \geq 0.$$

Further, observe that since $U_i(a_i^{\text{br}}) \geq U_i(a_i^{\text{opt}}, a_{-i})$ from the definition of a best response, we can replace $U_i(a_i^{\text{br}})$ $U_i(a_i^{\text{opt}}, a_{-i})$. We can utilize the β -misalignment and substitute

for the utility functions using Eq. (6.21) to arrive at

$$\mathbb{E}_{a \sim \sigma} [4\beta n W(a^{\text{opt}}) + \sum_{i \in \mathcal{I}} W(a) - W(a_i^{\text{opt}}, a_{-i})] \geq 0.$$

Applying the definition of λ_c and μ_c as in Eq. (6.16) results in the final inequality.

$$\mathbb{E}_{a \sim \sigma} [4\beta n W(a^{\text{opt}}) + \mu_c W(a) - \lambda_c W(a^{\text{opt}})] \geq 0.$$

Notice that the above inequality holds for any arbitrary sink equilibrium σ . Thus rearranging terms and using linearity of expectation gives us the price of sinking guarantee in Eq. (6.17). \square

Proof of Geometric. We can apply Lemma 7 with respect to the welfare function W to get Eq. (6.20). For ease of notation, let $\bar{\beta} = 1 - \beta$. We can successively apply the geometric misalignment property in Eq. (6.15), as well as using the fact that $U_i(a_i^{\text{br}}) \geq U_i(a_i^{\text{opt}}, a_{-i})$, to arrive at the following set of inequalities.

$$W(a_i^{\text{br}}) \geq \bar{\beta} U_i(a_i^{\text{br}}) \geq \bar{\beta} U_i(a_i^{\text{opt}}, a_{-i}) \geq \bar{\beta}^2 W(a_i^{\text{opt}}, a_{-i})$$

Substituting these inequalities back into Eq. (6.20) gives

$$\mathbb{E}_{a \sim \sigma} [\sum_i W(a) - \bar{\beta}^2 W(a_i^{\text{opt}}, a_{-i})] \geq 0.$$

Now we can substitute the definitions of λ_c and μ_c in Eq. (6.16) to a portion of the terms and simplify to get

$$\mathbb{E}_{a \sim \sigma} [n(1 - \bar{\beta}^2)W(a) + \bar{\beta}^2 (\mu_c W(a) - \lambda_c W(a^{\text{opt}}))] \geq 0.$$

Notice that the above inequality holds for any arbitrary sink equilibrium σ . Thus rearranging terms and using linearity of expectation gives us the price of sinking guarantee in Eq. (6.18). \square

Thus we have shown the bounds for both geometric and arithmetic misalignment cases. \square

We see that when the utility functions are close to the common interest utility design with $\beta \sim 0$, the price of sinking guarantees match the guarantees for the common interest utility. We note that while our approach allows us to get nontrivial guarantees on the sink equilibria, we still suffer from the degradation of the guarantee as the number of agents $n \rightarrow \infty$ increases arbitrarily. However, we assume worst case deviations (see $W(a_i^{\text{br}}) \geq \bar{\beta}^2 W(a_i^{\text{opt}}, a_{-i})$ in the proof of the geometric misalignment) for all actions in the game, which is not true for most natural games and produces a conservative bound. Thus the focus of future work will be to address this concern to get tighter guarantees. We can also get alternative guarantees if we consider sink induced by better responses. This is discussed in the next section.

6.4 Better Response Sink Equilibrium

In this section we consider sink equilibria that are induced by a better (rather than best) response process. In contrast to the best response set in Eq. (6.5), we consider the *better response set* defined as

$$\text{br}_i(\bar{a}) = \{a_i \in \mathcal{A}_i : U_i(a_i, \bar{a}_{-i}) \geq U_i(\bar{a})\} \quad (6.22)$$

for a given agent i . The better response process is then defined by a random walk, similar to Eq. (6.6) as

$$\Pr(\tilde{a}|\mathbf{a}) = \begin{cases} \frac{1}{n \cdot |\text{br}_i(\mathbf{a})|} & \text{if } \tilde{a} \in (\text{br}_i(\mathbf{a}), a_{-i}) \text{ for some } i \in \mathcal{I} \\ 0 & \text{otherwise.} \end{cases} \quad (6.23)$$

The sink equilibria are similarly defined for these dynamics. If we consider sink equilibria that are induced by better responses, it is possible to get positive guarantees on the behavior of sink equilibrium. More specifically, we show that there always exists a joint action in the support of the sink equilibria that has similar guarantees to the Nash equilibria.

Proposition 10. *Let G be (λ, μ) -smooth. Every sink equilibrium in G induced by better responses contains a joint action $\tilde{a} \in \text{supp}(\sigma)$ in its support such that $W(\tilde{a}) \geq \frac{\lambda}{\mu} W(a^{\text{opt}})$.*

Proof of Proposition 10. We show that $W(a) \geq \frac{\lambda}{\mu} W(a^{\text{opt}})$ for some $a \in \text{supp}(\sigma)$ in the sink induced by better responses. We first claim that for any sink σ , there exists a joint action $a \in \text{supp}(\sigma)$ such that for all $i \in \mathcal{I}$,

$$U_i(a) - U_i(a_i^{\text{opt}}, a_{-i}) \geq 0. \quad (6.24)$$

Consider an arbitrary action $a \in \text{supp}(\sigma)$ in which the condition does not hold true. Let i be the smallest number such that a does not satisfy Eq. (6.24) for agent i . Then the action $\hat{a} = (a_i^{\text{opt}}, a_{-i})$ is a better response to a and therefore $\hat{a} \in \text{supp}(\sigma)$ is in the support of σ as well. Note that \hat{a} also satisfies Eq. (6.24) for agent i . By induction, we can then derive an action $a^* \in \text{supp}(\sigma)$ such that Eq. (6.24) is satisfied for all i . By the

smoothness inequality in Eq. (6.13), we have that

$$\mu W(a^*) - \lambda W(a^{\text{opt}}) \geq \sum_i U_i(a^*) - U_i(a_i^{\text{opt}}, a_{-i}^*) \geq 0.$$

Therefore for some $a^* \in \text{supp}(\sigma)$, the efficiency is lower bounded by $W(a^*) \geq \frac{\lambda}{\mu} W(a^{\text{opt}})$. \square

6.5 Appendix

We outline some extra proofs of Proposition 6 and Proposition 7.

Lemma 8. *Consider the problem defined in Example 6. The welfare function in Eq. (6.8) is β -arithmetically misaligned, with an expected misalignment of*

$$\mathbb{E}[\beta] \leq |\mathcal{R}| \left(d \sqrt{\frac{2}{\pi}} e^{-\frac{c^2}{2d^2}} + c(1 - 2\Phi(-c/d)) \right), \quad (6.25)$$

where Φ is the normal cumulative distribution function.

Proof. We derive a bound for $\frac{|U_i(a) - W(a)|}{W(a)}$ based on the parameters given in Example 6.

From the equations defining the utility and welfare,

$$\begin{aligned} \frac{|U_i(a) - W(a)|}{W(a)} &= \frac{|\sum_{r \in U_i a_i} v_r^i - \sum_{r \in U_i a_i} v_r|}{\sum_{r \in U_i a_i} v_r} \\ &\leq \frac{\sum_r |v_r^i - v_r|}{\sum_r v_r}, \end{aligned}$$

where we use triangle inequality and $U_i a_i \subseteq \mathcal{R}$ to get the inequality on the right hand side. Observe that $\frac{\sum_i x_i}{\sum_i y_i} \leq \sum_i \frac{x_i}{y_i}$. This fact coupled with linearity of expectation and

$v_r \geq 0$ gives the inequality

$$\mathbb{E} \left[\frac{|U_i(a) - W(a)|}{W(a)} \right] \leq \sum_r \mathbb{E} \left[\left| \frac{v_r^i - v_r}{v_r} \right| \right].$$

We assume that $v_r^i \sim \mathcal{N}[v_r + c, (d \cdot v_r)^2]$ is drawn from a normal distribution. Thus the expectation $\mathbb{E} \left[\left| \frac{v_r^i - v_r}{v_r} \right| \right] = \mathbb{E}[\mathcal{N}_f(c, d^2)]$, where \mathcal{N}_f is a folded normal distribution with mean c and variance d^2 . This holds true for any $r \in \mathcal{R}$, so using the equality

$$\mathbb{E} \left[\left| \frac{v_r^i - v_r}{v_r} \right| \right] = \left(d \sqrt{\frac{2}{\pi}} e^{-\frac{c^2}{2d^2}} + c(1 - 2\Phi(-c/d)) \right)$$

results in the bound given in Eq. (6.25). \square

Proof of Proposition 6. We remark that the covering problem in Example 6 is a submodular game [43]. Under the common interest utility, the constants $\lambda_c = 1$ and $\mu_c = 2$ satisfy the inequality in Eq. (6.16) for submodular games [101, Example 2.6]. We can directly apply the guarantee given in Eq. (6.17) for the arithmetic misalignment for

$$\begin{aligned} \mathbb{E}[\text{PoSE}(\text{G})] &\geq \mathbb{E}[\max(\frac{\lambda_c - 4n\beta}{\mu_c}, 0)] \\ &\geq \max(\frac{\lambda_c - 4n\mathbb{E}[\beta]}{\mu_c}, 0), \end{aligned}$$

applying Jensen's inequality. Note that $\mathbb{E}[\beta]$ is given in Eq. (6.25) and substituting $\lambda_c = 1$ and $\mu_c = 2$ gives the lower bound. \square

Lemma 9. *Consider the problem defined in Example 7 with $w_{ij}^i/w_{ij} \in [\alpha, \alpha^{-1}]$ for all interference weight estimates. The welfare function in Eq. (6.11) is $(1 - \alpha)$ -geometrically misaligned.*

Proof of Lemma 9. We derive a bound for $\frac{U_i(a)}{W(a)}$ based on the parameters given in Example

6. First, we verify the identity

$$\min_i \frac{x_i}{y_i} \leq \frac{\sum_i x_i}{\sum_i y_i} \leq \max_i \frac{x_i}{y_i}.$$

Let $m = \min_i \frac{x_i}{y_i}$ and $M = \max_i \frac{x_i}{y_i}$. Then $m \sum_i y_i \leq \sum_i x_i \leq M \sum_i y_i$ and the identity hold true by dividing all sides by $\sum_i y_i$. Now we can use this identity to show that

$$\begin{aligned} \alpha &\leq \min_{i,j} \frac{w_{ij}^i}{w_{ij}} \leq \frac{U_i(a)}{W(a)} = \frac{\sum_i \sum_{j:k_j \neq k_i} w_{ij}^i}{\sum_i \sum_{j:k_j \neq k_i} w_{ij}} \\ &\leq \max_{i,j} \frac{w_{ij}^i}{w_{ij}} \leq \alpha^{-1}, \end{aligned}$$

from the assumption that $w_{ij}^i \in [\alpha \cdot w_{ij}, \alpha^{-1} \cdot w_{ij}]$. Thus, we see that game is geometrically misaligned with $\beta = 1 - \alpha$. \square

Proof of Proposition 6. We claim that $\lambda_c = 1$ and $\mu_c = 3$ are valid constants that satisfy Eq. (6.16) when $k = 1$. Under the claim, subbing λ_c and μ_c in Eq. (6.18) for geometric misalignment with $\beta = (1 - \alpha)$ gives the final expression in Eq. (6.12).

Now we show the claim. The action set for agent i can be defined as $\mathcal{A}_i = \{1, 2\}$, depending on which channel agent i chooses. Let \hat{a} be an arbitrary joint action and a^{opt} be the joint action that maximizes the welfare. We partition the agent set \mathcal{I} with the following subsets

$$D = \{i \in \mathcal{I} : \hat{a}_i = 1 \text{ and } a_i^{\text{opt}} = 2\},$$

$$O = \{i \in \mathcal{I} : \hat{a}_i = 2 \text{ and } a_i^{\text{opt}} = 1\},$$

$$B = \{i \in \mathcal{I} : \hat{a}_i = 1 \text{ and } a_i^{\text{opt}} = 1\},$$

$$N = \{i \in \mathcal{I} : \hat{a}_i = 2 \text{ and } a_i^{\text{opt}} = 2\}.$$

Additionally, for ease of notation, given subsets $S_1, S_2 \subset \mathcal{I}$, we can define $w(S_1, S_2) = \sum_{i \in S_1} \sum_{j \in S_2} w_{ij} + w_{ji}$. Under these definitions, we have that the welfare function can be written as

$$\begin{aligned} W(\hat{a}) &= \sum_i \sum_{j: \hat{a}_i \neq \hat{a}_j} w_{ij} \\ &= w(B, N) + w(D, N) + w(O, B) + w(O, D). \end{aligned}$$

Similarly, it can be verified that $W(a^{\text{opt}}) = w(B, N) + w(O, N) + w(D, B) + w(D, O)$.

Next, we rewrite the sum of deviations as

$$\begin{aligned} \sum_{i \in \mathcal{I}} W(a) - W(a_i^{\text{opt}}, \hat{a}_{-i}) &= \\ w(D, N) + w(O, B) + 2w(O, D) - \\ w(D, D) - w(O, O) - w(D, B) - w(O, N). \end{aligned}$$

From these definitions, we have the following set of inequalities,

$$\begin{aligned} W(a^{\text{opt}}) + \sum_{i \in \mathcal{I}} W(a) - W(a_i^{\text{opt}}, a_{-i}) &\leq \\ 3w(B, O) + w(B, N) + w(D, N) + \\ w(O, B) - w(D, D) - w(O, O) &\leq \\ 3w(B, O) + 3w(B, N) + 3w(D, N) + 3w(O, B) &\leq 3W(\hat{a}). \end{aligned}$$

Since \hat{a} was chosen arbitrarily, we see that Eq. (6.16) is satisfied for $\lambda_c = 1$ and $\mu_c = 3$ for all joint actions. \square

Part II

Industrial Refrigeration

Chapter 7

Load Shifting in Compressor Sequencing

The widespread and significant energy demands of industrial refrigeration have spurred numerous research efforts aimed at exploring various control strategies to reduce energy consumption. In this chapter, we focus on the idea of compressor sequencing, which entails selected the most efficient operational states for the compressors while meeting the required refrigeration load. We will introduce the industrial refrigeration setting and then iterate through two different sequencing problems: static and dynamic. Additionally, we introduce load shifting to address the compressor sequencing problem, which involves pre-cooling to allow for more operationally efficient compressor states.

Our analysis utilizes real-world sensor data from an industrial refrigeration facility operated by Butterball LLC ® in Huntsville, AZ. The findings reveal that, in the absence of load shifting, even optimally sequenced compressors frequently operate inefficiently due to running at intermediate capacity levels. However, by incorporating load shifting, we identify a potential energy savings of up to 20% compared to the optimal compressor sequencing alone. This demonstrates the significant impact of load shifting on improving

energy efficiency in industrial refrigeration systems. This chapter is based on the work in [103].

7.1 Preliminaries on Industrial Refrigeration

Industrial refrigeration systems are present in a multitude of sectors, not limited to food processing, plastics, electronics, and chemical processing [5–7]. Altogether, industrial refrigeration accounts for approximately 8.4% of total energy usage in the U.S [104]. As such, there are tremendous energy saving opportunities available in industrial refrigeration, not only through updating hardware components, but also increasing the sophistication of the implemented control algorithms. Algorithmic improvements are potentially more enticing, as they can realize significant energy savings with minimal capital expenditures to retrofit the system.

The four central components of a prototypical refrigeration system include the evaporators, compressors, condensers, and the expansion valve, with interconnections as illustrated in Figure 7.1. We study a common industrial refrigeration process, where ammonia refrigerant is circulated in a closed loop in a vapor compression cycle to move heat against the thermal gradient of the system. Informally, the main thermodynamic steps in an ideal vapor compression cycle are summarized as follows¹:

(1 \rightarrow 2) The refrigerant vapor flows through a compressor, where it is compressed from a low pressure, referred to as suction pressure, to a high pressure, referred to as discharge pressure. A consequence of this compression is an increase in the temperature of the refrigerant vapor, which now takes the form of a super-heated vapor.

(2 \rightarrow 3) The refrigerant super-heated vapor is then fed to the condenser, where

¹The actualized refrigeration process deviates from this idealization considerably, but we simplify for the purpose of presentation.

constant pressure heat rejection occurs, and heat is released to the ambient environment, resulting in condensation of the ammonia. A consequence of this heat rejection is that the refrigerant transitions from a super-heated vapor to a saturated liquid.

(3 \rightarrow 4) The refrigerant is then expanded adiabatically across an expansion valve, reducing the temperature and pressure and resulting in a vapor-liquid mixture.

(4 \rightarrow 1) Cooled refrigerant liquid flows through the evaporator, where heat absorption from the system via evaporation of the refrigerant occurs, and super-heated vapor is fed back to the compressor, completing the cycle.




Figure 7.1: A simplified diagram of the refrigeration components are depicted showing the flow of ammonia through the vapor compression process.

We direct the interested reader to [7] for a comprehensive review of refrigeration systems. The configuration of the whole refrigeration system can have significant impacts on the cost of operation: this can either be measured through total power, electric cost, carbon emissions, etc. Infrastructural retrofits of the refrigeration system, including changing the choice of refrigerant, hardware specifications of components, or general system layout, can be typically costly to implement. Accordingly, a more viable way to reduce costs is to strategically adjust the control policies of the refrigeration components to meet the required heat extraction while minimizing the operational cost. For example, *thermal load shifting* has received significant attention as a methodology to preemptively cool a facility in order to take financial advantage of dynamic energy cost-rate structures [105–111]. Additionally, another approach is *set point optimization*, where the set points for suction and discharge pressure are dynamically adjusted to drive the refrigeration system towards an energy optimal operating state while maintaining the desired achievable cooling demands [112–114]. For these domains, standard control and optimization techniques, such as model predictive control and set-point tracking, can be implemented for attaining favorable control strategies.




Figure 7.2: This figure highlights the cumulative distribution functions for the slide valve position for four compressors operating at the Butterball facility during the month of June, 2023. Here, the slide valve position is associated with compressor capacity, where 100% means that the compressor is running at full capacity. We also highlight the percentages in which each compressor is operating at full capacity (where the slide valve sensor is measured above 99%) or trim as well as the percentage of time the compressor is turned on and off. Note that the compressors are often operating in trim, suggesting that there are potential opportunities to save energy by operating the compressors at full capacity more often.

While most of the existing control approaches for energy optimization focus on the evaporators, e.g., thermal load shifting, it is important to highlight that the compressors represent the dominant energy expenditure (around 85%) in most refrigeration systems. For example, at an industrial refrigeration site of Butterball (a large poultry processing facility) during the month of June 2023, 40% of total energy usage is attributed to the compressors, 5% is attributed to the rest of the refrigeration process, and 55% is attributed to the non-refrigeration components of the facility. It is widely known that compressors are operated most efficiently when running at full capacity [115,116]; however, the typical control objective for the compressors is suction pressure stabilization. In this way, the operational state of the compressors is directly dependent on the state of the evaporators and this can ultimately lead to the compressors operating in an inefficient manner, i.e., at partial capacity. Figure 7.2 confirms this phenomena, directly highlighting the cumulative distribution functions for the slide valve position of the four compressors operating at the Butterball facility during this time period.² Here, the slide valve position can be viewed in the same light as capacity, where 100% means that the compressor is running at full capacity. Note that the slide valve positions are often significantly below 100%, suggesting that there are potential energy saving opportunities in algorithmic improvements for compressor scheduling and control.

We will shift the control focus from the evaporators directly to the compressors, where the goal is to optimize the operational state of the compressors to serve the required refrigeration load. We formalize this optimization problem as the *compressor sequencing problem* [115]. We begin by characterizing the optimal solution to the *static* compressor sequencing problem, which focuses on satisfying a given refrigeration load at a single

²This data was acquired through direct partnership with CrossnoKaye (see crossnokaye.com), which focuses on the derivation and implementation of intelligent control systems for industrial refrigeration systems in the cold food and beverage domain. CrossnoKaye has been monitoring and controlling the refrigeration system at the Butterball facility since April 2023




Figure 7.3: Approximate power breakdown for a Butterball facility at Huntsville.

time instance. We show that the optimal solution can actually be derived via a simple water-filling algorithm. However, the water filling algorithm must be executed with the correct compressor sequence which can be highly dependent on the refrigeration load, posing implementation and computational hurdles.

Given these limitations, we shift focus to the *dynamic* compressor sequencing problem. Here, we are provided with a given time-dependent profile of the refrigeration load that we need to serve over a given horizon. See forthcoming Figure 7.5 as an illustration of a typical refrigeration load over a month long horizon at the Butterball facility. Unlike the static problem, this dynamic formulation gives us the flexibility to exploit load shifting, where one preemptively cools the facility, so that the compressors can operate in a more efficient fashion, i.e., more often at full capacity. Not only does load shifting provide substantial potential for energy savings, interestingly, the resulting optimal solution is again a simple water filling algorithm with an order that is fixed and can be easily computed. Hence with load shifting, the optimal compressor sequencing becomes more amenable to real world implementations.

In order to practically assess the energy saving opportunities associated with compressor sequencing, we implement a numerical case study on the Butterball facility in Section

7.1. Using collected time-series data on compressor configurations and refrigeration load estimates, our initial results suggest that the potential energy savings could be significant, with upwards of 20% reduction in total energy expenditure when comparing optimal compressor sequencing with load shifting to optimal compressor sequencing without load shifting. Furthermore, this chapter provides a number of supporting results characterizing properties of the optimal and near-optimal online load shifting algorithms.

While this chapter introduces load shifting as a novel approach to the compressor sequencing problem, practical implementation requires careful consideration. Adjusting compressor capacities haphazardly can lead to system instability, but a potential implementation strategy is to have the evaporators directly respond to changes in compressor optimization. Nevertheless, this chapter primarily focuses on assessing the potential benefits of load shifting, leaving the development of practical control strategies for future investigation.

7.2 Mathematical Model

In this chapter, we formalize the control problem for optimal compressor sequencing. Here, the operational state of the compressors (e.g. the on/off status as well as the slide valve position) is chosen such that the thermal demands are met with the least cost, which we measure in terms of energy usage. We discuss potential opportunities for algorithmic improvements in this chapter.

Many large scale refrigeration systems employ algorithms for intelligently choosing the operational state of the compressors. In refrigeration systems with multiple compressors, one must decide the operational state of these compressors that is necessary to service the underlying refrigeration load. More formally, let \mathcal{C} denote a finite set of compressors (for Butterball, $\mathcal{C} = \{C_1, C_2, C_3, C_4\}$), where each compressor $c \in \mathcal{C}$ is associated with a minimum and maximum heat capacity, q_c^- and q_c^+ respectively, as well as a power-heat

curve $P_c : Q_c \rightarrow \mathbb{R}_{\geq 0}$ where $Q_c = 0 \cup [q_c^-, q_c^+]$ designates the viable refrigeration loads on compressor c , with 0 indicating the compressor is turned off. Here, $P_c(q_c) \geq 0$ is the power required to serve heating load $q_c \in Q_c$ through compressor c . We assume that $P_c(0) = 0$ and P_c is concave and increasing over the interval $[q_c^-, q_c^+]$, which implies that compressors operate more efficiently at higher capacities. Specifically for the compressors in operation at Butterball, we assume an affine structure for the power-heat curves, where for any compressor $c \in \mathcal{C}$ and thermal load $q_c \in [q_c^-, q_c^+]$ we have

$$P_c(q_c) = P_c(q_c^-) + \left(\frac{q_c - q_c^-}{q_c^+ - q_c^-} \right) (P_c(q_c^+) - P_c(q_c^-)).$$

We validate the affine models for the power-heat curves against collected data on estimated refrigeration load and compressor power, which is shown in Figure 7.4. The extreme points of these power-heat curves is summarized in Table 7.1. Note that for

Compressor	C1	C2	C3	C4
Model	Screw	Screw	Screw	Screw
q_c^- (kW)	220	239	165	284
q_c^+ (kW)	3000	2126	1760	2351
$P(q_c^-)$ (kW)	124	173	142	181
$P(q_c^+)$ (kW)	262	427	356	494

Table 7.1: Compressor Characteristics

simplicity, we removed the dependence on slide valve position to provide a direct relationship between thermal load and power.

The problem of compressor sequencing centers on the goal of meeting the incoming refrigeration load, which we denote by $q^{\text{in}} \in \mathbb{R}_{\geq 0}$, with the least possible energy expenditure. More formally, the goal is to identify compressor loads $\{q_c\}_{c \in \mathcal{C}}$ that satisfy the incoming refrigeration load, i.e., $\sum_{c \in \mathcal{C}} q_c \geq q^{\text{in}}$, and minimize the total work expenditure as measured by the total power usage by the compressors, i.e., $\sum_{c \in \mathcal{C}} P_c(q_c)$. We denote this compressor assignment by the policy $\pi : \mathbb{R}_{\geq 0} \rightarrow \prod_{c \in \mathcal{C}} Q_c$, where $\pi(q^{\text{in}}) = \{q_c\}_{c \in \mathcal{C}}$




Figure 7.4: For each of the compressors, the estimated power and heat capacity for each minute in the month of June was recorded for a Butterball facility. We depict the resulting spread in the given figure and notice a fairly affine relationship, which we denote in red. This is also supported from manufacturing simulation software for the compressors.

designates the refrigeration loads for each compressor $c \in \mathcal{C}$. The cost of a policy π for a given q^{in} is defined by $J_\pi(q^{\text{in}}) = \sum_{c \in \mathcal{C}} P_c(\pi_c(q^{\text{in}}))$. We will henceforth remove the dependence on the compressor set, i.e., denote $\{\cdot\}_{c \in \mathcal{C}}$ as merely $\{\cdot\}$, for notational simplicity.

7.3 Results on Compressor Sequencing

7.3.1 Fixed Order Compressor Sequencing

The industrial standard for compressor operation is to meet a given refrigeration load q^{in} through a water filling algorithm with a pre-determined order of compressors \mathcal{C} . For ease of presentation, let the set of compressors $\mathcal{C} = \{c_1, c_2, \dots, c_m\}$ naturally denote the order of the compressors, i.e., c_1 first, c_2 second, etc. Then the operation of the compressors according to this policy, represented by π^{FO} , is given by Algorithm 3. We will denote the fixed order policy as $\pi^{\text{FO}}(q^{\text{in}}; \mathcal{O})$, where \mathcal{O} describes a specific ordering of the compressor set \mathcal{C} . Note that Algorithm 3 returns a thermal load profile $\{q_c\}$ that is guaranteed to satisfy the inequality $\sum_{c \in \mathcal{C}} q_c \geq q^{\text{in}}$ provided that we assume that $\min_c q_c^- \leq q^{\text{in}} \leq \sum_{c \in \mathcal{C}} q_c^+$. When $q_c = q_c^+$, we say that the compressor is operating at full capacity. Alternatively, when $q_c^+ > q_c \geq q_c^-$, we say that the compressor is operating

Algorithm 3 Water Filling Algorithm

Require: \mathcal{O} , q^{in} , $q^{\text{tot}} \leftarrow 0$, $q_c \leftarrow 0$ for all $c \in \mathcal{C}$

for c in \mathcal{O} **do**

if $q^{\text{in}} > q^{\text{tot}}$ **then**

$q_c \leftarrow q_c^+$

$q^{\text{tot}} \leftarrow q^{\text{tot}} + q_c^+$

end if

end for

for c in reverse(\mathcal{O}) **do**

if $q^{\text{in}} \leq q^{\text{tot}}$ and $q_c \neq 0$ **then**

$d \leftarrow \min\{q_c^+ - q_c^-, q^{\text{tot}} - q^{\text{in}}\}$

$q_c \leftarrow q_c - d$

$q^{\text{tot}} \leftarrow q^{\text{tot}} - d$

end if

end for

return $\{q_c\}_{c \in \mathcal{C}}$

in trim. The central tuning parameter of the policy $\pi^{\text{FO}}(q^{\text{in}}; \mathcal{O})$ is the order \mathcal{O} that the water filling algorithm is run under. There may be significant differences in energy usage for different orders, especially if compressors vary in the efficiency.

7.3.2 Optimal Compressor Sequencing

While the standard water-filling algorithms provide a straightforward approach to compressor sequencing, moving away from a fixed order scheme may lead to more efficient compressor operation. Hence, we consider the problem of optimal compressor sequencing in this chapter, where the goal is to determine the compressor state that meets the refrigeration load with the least possible energy expenditure. More formally, the operation of the compressors would be determined by the solution of the following non-convex optimization problem.

$$\begin{aligned}
 J^*(q^{\text{in}}) &= \min_{q_c \in Q_c} \sum_{c \in \mathcal{C}} P_c(q_c) \\
 \text{s.t.} \quad &\sum_{c \in \mathcal{C}} q_c \geq q^{\text{in}}
 \end{aligned} \tag{7.1}$$

The following proposition characterizes the structure of optimal solution to the compressor sequencing problem in Eq. (7.1). In fact, regardless of the refrigeration load q^{in} , the optimal compressor state can be realized by a water-filling algorithm with a specific order that depends on q^{in} .

Proposition 11. *Let q^{in} be the incoming refrigeration load. The optimal compressor state, as given by the solution of Eq. (7.1), can be realized by the water filling algorithm given in Algorithm 3 with a specific order \mathcal{O} that depends on q^{in} .*

Proof. To show this statement, we show equivalently that for all the compressors, only one compressor $c \in \mathcal{C}$ has $q_c^+ > q_c^* > q_c^-$ in the optimal solution to Eq. (7.1). Note that if this is true, the optimal order \mathcal{O} coincides to when the compressors are ordered decreasing in their capacities $\{q_c^*\}$. Given this order, Algorithm 3 will produce the equivalent capacity $\{q_c^{\text{alg}}\} = \{q_c^*\}$ that match the optimal solution.

We show the claim that at most one $q_c^* \notin \{q_c^-, q_c^+, 0\}$ is not at the endpoints through contradiction. Let i and j be the compressors at partial capacity in the optimal solution. Notice that the function $P_i(q_i + d) + P_j(q_j - d)$ is a concave function of d , by assumption of concavity of P_i and P_j and preservation of concavity under affine transformations. For any feasible d , note that the constraint in Eq. (7.1) is always satisfied if we change q_i to $q_i + d$ and q_j to $q_j - d$. Additionally, since $P_i(q_i + d) + P_j(q_j - d)$ is concave, the optimal value for d must occur at either endpoints of the feasible interval. Thus in the optimal solution, q_i or q_j must be either at q_c^- or q_c^+ or 0, ensuing in contradiction. \square

This structural result demonstrates that the optimal control algorithm for compressor sequencing could be realized from the perspective of a partitioning process where the admissible refrigeration loads are partitioned into various regions, and each region is associated with a distinct ordering. However, this approach to compressor sequencing has significant problems from both a computation and implementation perspective. First,

solving the optimization problem in Eq. (7.1) represents a mixed integer optimization that grows exponentially in complexity in the size of the compressor set \mathcal{C} . Furthermore, the optimal order can change drastically as a function of the refrigeration load. This means that the state of the compressors could shift wildly during operation, which may be infeasible due to delays in changing compressor capacities and causing unnecessary variability in compressor operation. Thus, we look to load shifting as a medium to smooth out the compressor sequencing problem - we discuss this next.

7.3.3 Optimal Compressor Sequencing with Load Shifting

Load shifting is common practice in refrigeration systems for reducing operational costs. Load shifting involves the process of preemptively cooling a facility, thereby using the product within the facility as a thermal battery to save on future cooling demands. Accordingly, for this setting we will think about cooling needs over a given discrete horizon $[0, 1, \dots, T]$ where the refrigeration load at each stage k is given by $q^{\text{in}}(k)$ and $\mathbf{q}^{\text{in}} = \{\mathbf{q}^{\text{in}}(\mathbf{k})\}_{0 \leq \mathbf{k} \leq \mathbf{T}}$. Here, we will assume that there is complete knowledge of the refrigeration load over the horizon at the initial stage $k = 0$. This assumption will allow us to hypothetically assess the potential opportunities associated with load shifting for compressor sequencing on realistic refrigeration load profiles as provided in forthcoming Figure 7.5, which highlights the refrigeration load over the month of June, 2023.

The goal of optimal compressor sequencing is to establish a new shifted thermal load demand trajectory $\mathbf{q}^{\text{sh}} = \{\mathbf{q}^{\text{sh}}(\mathbf{0}), \dots, \mathbf{q}^{\text{sh}}(\mathbf{T})\}$ and dynamic compressor states $\mathbf{q}_{\mathbf{c}} = \{\mathbf{q}_{\mathbf{c}}(\mathbf{0}), \dots, \mathbf{q}_{\mathbf{c}}(\mathbf{T})\}$ that minimize the cumulative energy expenditure. Here, we require the shifted load demand trajectory satisfies

$$\sum_{k=0}^{\tau} q^{\text{sh}}(k) \geq \sum_{k=0}^{\tau} q^{\text{in}}(k), \quad \forall \tau \in [0, T],$$

where the provided cooling exceeds the refrigeration load required for any horizon $[0, \tau]$ with $\tau \in \{0, \dots, T\}$. Accordingly, our new optimization takes on the following form:

$$\begin{aligned}
J^*(\mathbf{q}^{\text{in}}) &= \min_{\mathbf{q}^{\text{sh}}, \{\mathbf{q}_c\}} \frac{1}{T} \sum_{k=0}^T \sum_{c \in \mathcal{C}} P_c(q_c(k)) \\
\text{s.t. } & q_c(k) \in Q_c \text{ for all } c \in \mathcal{C}, k \in [0, T], \\
& \sum_{k=0}^{\tau} q^{\text{sh}}(k) \geq \sum_{k=0}^{\tau} q^{\text{in}}(k) \text{ for all } \tau \in [0, T], \\
& \sum_{c \in \mathcal{C}} q_c(k) \geq q^{\text{sh}}(k) \text{ for all } k \in [0, T].
\end{aligned} \tag{7.2}$$

This optimization has two sets of decision variables: the shifted thermal load trajectory \mathbf{q}^{sh} and the dynamic compressor loads $\{\mathbf{q}_c\}$. The first constraint ensures that the compressors loads are viable for every stage k . The second constraint dictates that the shifted load trajectory \mathbf{q}^{sh} needs to deliver at least as much cooling as any nominal thermal profile \mathbf{q}^{in} for any horizon $\tau \in [0, T]$. The last constraint ensures that the total compressor load needs to match the shifted thermal load \mathbf{q}^{sh} at each stage k . The optimization problem in Eq. (7.2) is determined through $\{\mathbf{q}_c\}$ and \mathbf{q}^{sh} , whereas the static compressor sequencing problem in Eq. (7.1) fixes $q^{\text{sh}}(k) = q^{\text{in}}(k)$ for all k . While this results in a seemingly more complex optimization problem, the following proposition demonstrates that the optimal solution is actually attained by a fixed order water-filling algorithm.

Proposition 12. *Let \mathbf{q}^{in} be the dynamic refrigeration load. Furthermore, consider a set of compressors $\mathcal{C} = \{c_1, \dots, c_m\}$ that are ordered in terms of increasing marginal costs of cooling at full capacity, i.e., if compressor i comes before compressor j in the order \mathcal{O}^{sh} , we have that*

$$\frac{q_i^+}{P_i(q_i^+)} \leq \frac{q_j^+}{P_j(q_j^+)}. \tag{7.3}$$

Then the following optimization problem yields the same optimal cost in Eq. (7.2) when $T \rightarrow \infty$:

$$\begin{aligned}
J^*(\mathbf{q}^{\text{in}}) &= \min_{\mathbf{q}^{\text{sh}}} \frac{1}{T} \sum_{k=0}^T \sum_{c \in \mathcal{C}} P_c(q_c(k)) \\
\text{s.t. } q_c(k) &= \pi_c^{\text{FO}}(q^{\text{sh}}(k)) \text{ for all } c \in \mathcal{C}, k \in [0, T], \\
\sum_{k=0}^{\tau} q^{\text{sh}}(k) &\geq \sum_{k=0}^{\tau} q^{\text{in}}(k), \text{ for all } \tau \in [0, T]
\end{aligned} \tag{7.4}$$

where $\pi_c^{\text{FO}}(\cdot)$ comes from the water filling algorithm in Algorithm 3 with the above ordering \mathcal{O}^{sh} .

Proof. Let $\hat{\mathbf{q}}^{\text{sh}}, \{\hat{\mathbf{q}}_{\mathbf{c}}\}$ be the optimal solution for Eq. (7.2). As P_c is monotonic, we note that $\sum_{c \in \mathcal{C}} \hat{q}_c(k) = \hat{q}^{\text{sh}}(k)$ must hold with equality for all k in the optimal solution. Then, for a given $\hat{q}^{\text{sh}}(k)$, the optimal compressor loads $\{\hat{q}_c(k)\}$ can be given through the water filling algorithm in Algorithm 3 for each k for some order $\mathcal{O}(k)$. This can be shown with arguments similar to the proof of Proposition 11.

Now we show that $\mathcal{O}(k) = \mathcal{O}^{\text{sh}}$ for all k . Let i and j be compressors such that i comes before j in \mathcal{O}^{sh} but j comes before i in $\mathcal{O}(k)$. We claim that this is only possible a finite number of times. If not, there are an infinite number of times where $\hat{q}_i(k) = 0$, but $\hat{q}_j(k) > 0$. However, there exist a load shift and time points $\{k_\ell\}_{1 \leq \ell \leq N}$ time steps in which $q_i(k_\ell) \rightarrow q \leq q_i^+$ for $1 \leq \ell \leq M < N$ and $q_j(k_\ell) \rightarrow 0$ for $M < \ell \leq N$ which produces a more efficient solution, since the marginal cost of cooling for compressor i is less than compressor j . Thus the water filling algorithm with order \mathcal{O}^{sh} recovers an optimal solution to Eq. (7.2) when $T \rightarrow \infty$. \square

From Proposition 12, we see that a simple fixed order water-filling algorithm achieves the optimal solution to Eq. (7.2) given the correct shifted load \mathbf{q}^{sh} . We can also evaluate the potential cost benefits between compressor sequencing with and without load shifting. We characterize the greatest possible difference in cost in the next proposition. For

notational ease, we define the ratios $R_{\max} = \max_c P_c(q_c^-)/q_c^-$ and $R_{\min} = \min_c P_c(q_c^+)/q_c^+$.

Proposition 13. *For a given dynamic refrigeration load \mathbf{q}^{in} , let $J^{\text{cs}}(\mathbf{q}^{\text{in}})$ be the trajectory cost in Eq. (7.2) associated with compressor sequencing without load shifting and let $J^*(\mathbf{q}^{\text{in}})$ be the optimal trajectory cost with load shifting with $T \rightarrow \infty$. The fractional difference between the costs is upper and lower bounded by*

$$0 \leq \frac{J^{\text{cs}}(\mathbf{q}^{\text{in}}) - \mathbf{J}^*(\mathbf{q}^{\text{in}})}{J^*(\mathbf{q}^{\text{in}})} \leq \frac{R_{\max} - R_{\min}}{R_{\min}} \quad (7.5)$$

Proof. We first note that in Eq. (7.2), we recover the optimization without load shifting (or Eq. (7.1)) if we impose the constraint $\mathbf{q}^{\text{sh}} = \mathbf{q}^{\text{in}}$ directly. Therefore, we have that $J^*(\mathbf{q}^{\text{in}}) \leq \mathbf{J}^{\text{cs}}(\mathbf{q}^{\text{in}})$ necessarily, and the ratio is always non-negative. We first show the upper bound when considering one compressor; the ratios simplify to $R_{\min} = P_c(q_c^-)/q_c^-$ and $R_{\max} = P_c(q_c^+)/q_c^+$. Consider two thermal load profiles \mathbf{q}^1 and \mathbf{q}^2 of length $T \gg 0$, where $q^1(k) = q_c^-$ for $T - \text{round}(D \cdot q_c^+) \leq k < T$ and 0 elsewhere and $q^2(k) = q_c^+$ for $0 \leq k < \text{round}(D \cdot q_c^-)$ and 0 elsewhere. We choose D and T large enough that $\sum_k q^1(k)$ is approximately equal to $\sum_k q^2(k)$ and $\text{round}(D \cdot q_c^-) < T - \text{round}(D \cdot q_c^+)$. By our definitions, \mathbf{q}^2 is a load shifted version of \mathbf{q}^1 , in that it satisfies $\sum_{k=0}^{\tau} q^2(k) \geq \sum_{k=0}^{\tau} q^1(k)$ for all $\tau \in [0, T]$. Thus without load shifting for \mathbf{q}^1 , the power usage is approximately $P_c(q_c^-) \times D \cdot q_c^+$ and with load shifting to \mathbf{q}^2 , the power usage is approximately $P_c(q_c^+) \times D \cdot q_c^-$. Thus the ratio of power use matches the upper bound $(R_{\max} - R_{\min})/R_{\min}$. It can be easily verified that this example attains the worst case ratio via convexity of P_c . Extension to multiple compressors follows an analogous argument. A similar construction can be assembled, where in \mathbf{q}^2 , the most efficient compressor satisfies the required load at full capacity and in \mathbf{q}^1 , the least efficient compressor satisfies the required load at the minimum capacity. \square

Proposition 13 provides a possible range of energy savings when utilizing load shifting. By the worst-case constructions of \mathbf{q}^{in} , we see that if the profile \mathbf{q}^{in} fluctuates significantly, implementing load shifting can generate the most energy savings; however, if q^{in} is relatively constant, the energy savings may be less. For the compressors operating in Butterball, the fractional difference is between $0 \leq \frac{J^{\text{cs}}(\mathbf{q}^{\text{in}}) - J^*(\mathbf{q}^{\text{in}})}{J^*(\mathbf{q}^{\text{in}})} \leq 8.85$, suggesting significant energy savings through load shifting. We validate this in the next discussion.

Algorithm 4 Online Load-Shifting Algorithm

Require: $\mathbf{q}^{\text{in}}, (k, q^k, q^{\text{tot}}) \leftarrow (0, 0, 0), \mathbf{q}_c \leftarrow \mathbf{0}$ for all $c \in \mathcal{C}$

```

for  $k \leq T$  do
   $q^k \leftarrow q^k + q^{\text{in}}(k)$ 
   $q^{\text{tot}} \leftarrow \max\{q^k, \text{mean}(\mathbf{q}^{\text{in}})\}$ 
  for  $c$  in  $\mathcal{O}^{\text{sh}}$  do
    if  $q^{\text{tot}} > 0$  then
       $q_c \leftarrow q_c^+$ 
       $q^{\text{tot}} \leftarrow q^{\text{tot}} - q_c^+$ 
       $q^k \leftarrow q^k - q_c^+$ 
    end if
  end for
   $k \leftarrow k + 1$ 
end for

```

7.3.4 Online Compressor Sequencing with Load Shifting

While our results clearly highlights the potential opportunities of compressor sequencing with load shifting, there are still several practical concerns before a stable implementation can be employed onto a real refrigeration facility. One important issue is that a future trajectory of refrigeration load \mathbf{q}^{in} is usually not fully known. While the general daily trends for refrigeration load are relatively predictable, (for e.g., see that the curve in Figure 7.5 is rather periodic) estimating the exact future refrigeration load may be hindered by uncertainties in weather, product load, and other disturbances. Thus, in this chapter, we introduce an online version of a compressor sequencing algorithm with load




Figure 7.5: Over a month, we display the estimated refrigeration load serviced by the compressors, where each data point is associated with every minute in June 2023. In this figure, we also apply an average filter of 20 min. to smooth out the data. Data points for refrigeration load were calculated through estimates of compressor power and COP through sensor readings for each time point.

shifting, where at each time step, we only assume knowledge of the current refrigeration load and the total time-average of the refrigeration load.

This simple, online compressor sequencing algorithm is given in Algorithm 4. In this online algorithm, the compressors service either the time-average refrigeration load or the current required refrigeration load at full capacity. In this fashion, the constraints in Eq. (7.2) are satisfied with a certain degree of pre-cooling. We will see that this simple implementation achieves similar energy savings to the optimal compressor sequencing with load shifting in the provided simulations.

7.4 Simulations

In this chapter, we evaluate the potential energy savings possible through compressor sequencing and load shifting on a case study of the Butterball facility. Here, we use the predicted refrigeration load profile from the Butterball facility during the month of June 2023, depicted in Figure 7.5, to serve as a prototypical example for our case study. This profile was estimated using direct measurements of power usage of the compressors, as well derived coefficient of performance or COP (sometimes CP or CoP) of the refrigeration system, which is defined as the ratio of useful cooling to work done by the compressors.

We observe that the thermal profile is relatively cyclic, where the peaks correspond to working hours during the work week (when new product requiring cooling typically enters during standard operating hours), and the lower plateaus represent weekends (representing off hours of the facility).

In the Butterball facility, all compressors in operation are screw compressors, which use a screw thread to trap and compress a volume of gas. For screw compressors, the main control parameter to modulate the capacity is a continuous slide valve control.³ Increasing the slide valve will expose more of the screw thread, increasing the volume of gas to be compressed and increases the capacity of the compressor. As highlighted previously, compressors operate at their highest efficiency when running at full capacity. The slide valve being at the minimum position corresponds to the minimum heat capacity q_c^- (and similarly for q_c^+) for compressor c .

Methodology	Average Power
Worst Fixed Order	856.7 kW
Best Fixed Order	562.3 kW
Compressor Sequencing (C.S)	551.0 kW
Online C.S with Load Shifting	444.3 kW
C.S with Load Shifting	443.5 kW

Table 7.2: Cost of Algorithms

The main energy savings studied in this chapter is summarized in Table 7.2. First, we examine the average power usage when using a standard water-filling approach with a fixed order. We see significant energy savings if a good fixed compressor order is utilized over a bad compressor order; thus, compressor order has substantial effects on energy usage. When using optimal compressor sequencing without load shifting, surprisingly, we see the possible energy savings are only up to 2%, as compared to the

³It is also sometimes possible to modulate the capacity of a screw compressor through speed control with a variable speed drive. While we do not explicitly consider this method, our analysis extends to this scenario as well.

best fixed order algorithm. However, when utilizing compressor sequencing with load shifting, we see the energy savings jump to 20% when compared to the best fixed order algorithm. Furthermore, we find that the online compressor sequencing algorithm achieves a similar energy cost to optimal compressor sequencing with load shifting, suggesting the viability of online extensions. Thus, we validate load shifting as a viable mechanism for garnering energy savings with regards to compressor sequencing. In future work, we will extend these results to construct realizable control algorithms to offer actualized energy savings.

Chapter 8

Inventory Control and Peak Pricing

The extensive deployment of industrial refrigeration systems across various sectors substantially contributes to global energy consumption, underscoring significant opportunities for energy conservation through advanced control designs. This chapter concentrates on developing control algorithms for industrial refrigeration systems that aim to minimize operational costs while ensuring efficient heat extraction. Leveraging concepts from inventory control, we analyze the structure of optimal control policies and examine the influence of different energy cost-rate structures, including time-of-use (TOU) pricing and peak pricing.

Our findings reveal that classical threshold policies are optimal under TOU pricing schemes. However, the introduction of peak pricing disrupts their optimality, highlighting the necessity for meticulously designed control strategies when facing substantial peak costs. We present both theoretical results and simulation studies to illustrate this phenomenon, providing valuable insights for enhancing the efficiency of industrial refrigeration management. This chapter is based on the work in [117].

8.1 Introduction to Peak Pricing

Industrial refrigeration systems are widely utilized across diverse sectors, not limited to plastics manufacturing, chemical processing, food storage, and electronics production [6, 7, 118, 119]. Collectively, industrial refrigeration contributes to approximately 8.4% of total energy consumption in the United States [104]. Consequently, there exist significant opportunities for energy conservation within industrial refrigeration. These opportunities extend beyond hardware upgrades to include enhancements in the control algorithms implemented in these systems. Improvement in the algorithm design can potentially be more appealing, as they can yield substantial energy savings with minimal capital investment required for system retrofitting.

The intention behind improvements in algorithm design is to strategically adjust the control strategies of the components within the industrial refrigeration process to achieve the necessary heat extraction while minimizing operation costs. These costs can be measured through total power, electric costs, carbon emissions, or other relevant metrics. For example, the set points or steady state configurations of the components can be optimized to raise the energy efficiency [112–114, 120]. Furthermore, model predictive control or trajectory optimization can be utilized to dynamically optimize the energy efficiency of the refrigeration cycle over a time horizon [121–123]. Another approach that has garnered significant attention is *thermal load shifting*, where refrigeration loads are dynamically managed to leverage variable energy cost-rate structures [105, 108–111]. While these results are quite encouraging in conventional cost structures, there remains a need to characterize the qualitative behavior of optimal control policies when considering more varied cost structures.

There exist many different energy cost-rate structures that may depend on various factors. *Fixed rate* pricing, the most traditional rate structure, involves charging a flat

rate per unit of energy consumed, irrespective of external conditions. In contrast, *time-of-use* (TOU) pricing varies the per-unit rate based on the time or season. Typically, this results in higher energy costs during peak demand periods and lower energy costs during off-peak periods. In this setup, there may be significant economic benefit for the operator to shift their energy usage to off-peak hours. Deciding on how to shift is exactly the focus of the thermal load shifting literature. To further regulate energy usage, *peak pricing* can also be introduced, where the cost is dependent on the maximum power usage of the system over a time period and spikes in energy consumption are highly disincentivated. In many industrial refrigeration systems, peak pricing may comprise of a large portion of the energy costs. In this chapter, we focus on when both peak pricing and TOU costs are present and quantify the impact on optimizing the scheduling of refrigeration loads.

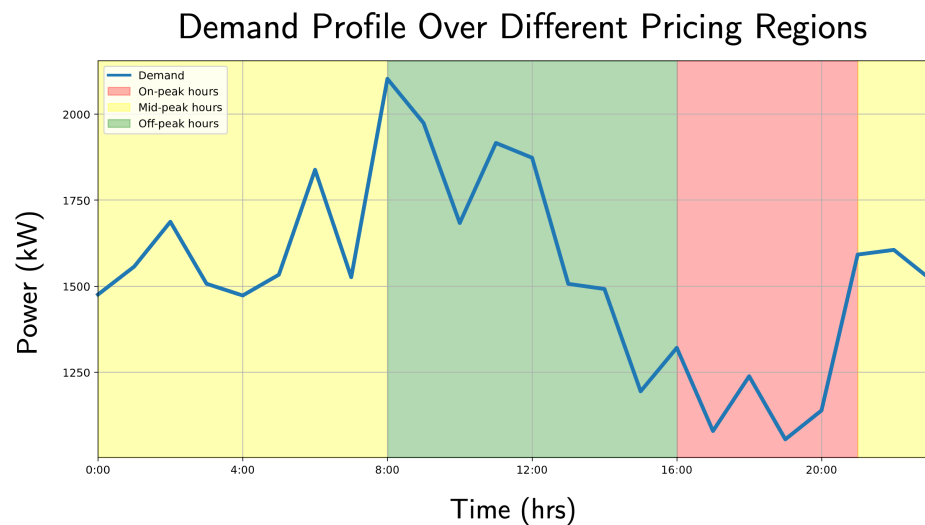


Figure 8.1: We depict a prototypical power consumption profile over different pricing regions.

Case Study. We introduce an example of a rate structure that a refrigeration facility may be charged with for its energy consumption.¹ The time-of-use costs are dependent

¹This data was obtained through a direct collaboration with CrossnoKaye (see crossnokaye.com), a

on time of day, where the cost rates vary between three time windows: on-peak hours (4:00PM - 9:00PM), mid-peak hours (9:00PM - 8:00AM), and off-peak hours (8:00AM - 4:00PM). Typically, energy-costs are higher in on-peak hours and lower in off-peak hours. Additionally, the facility can incur additional peak costs: a peak charge over the maximum energy consumption over a month and a peak charge over the month for each of the specific time windows (on, mid, off). We display a possible demand profile under this cost structure in Figure 8.1.

Designing control policies using optimization methods that account for peak pricing, though not as common as TOU pricing, has been studied previously [124–126]. However, the primary objective of this chapter is to theoretically evaluate how peak pricing influences the qualitative structure of the optimal control policies. We approach this through uniquely adopting the perspective of *inventory control*. Inventory control [127, 128] is a classical branch of multi-stage decision problems that explores purchasing policies of inventory to ensure optimal warehouse stock levels. When energy rates solely consist of time-of-use (TOU) costs, classical findings from inventory control suggest that the optimal control algorithms should adopt a threshold approach: when the facility’s temperature exceeds a specified threshold, the refrigeration load is increased to return the facility to a desired buffer temperature. Threshold policies are commonly implemented in practice for industrial refrigeration; we verify its optimality under TOU costs in Proposition 14. However, the introduction of peak pricing disrupts the optimality of threshold policies. In fact, in Proposition 16, we characterize the structure of the optimal control policy under peak pricing; we see that the optimality of simple threshold policies is lost in these settings. We also verify our findings through simulation in Chapter 8.4. Our results suggests that designing control policies should be done carefully if significant peak costs

company that focuses in integrating control systems for industrial refrigeration systems within the cold food and beverage sector.

are present.

8.2 Mathematical Model

In this chapter, we aim to devise scheduling strategies for refrigeration systems that strike a balance between minimizing overall costs and adequately meeting the necessary cooling demands of the facility. Additionally, we investigate the impact of *peak pricing*, a common energy pricing mechanism involving charging based on the maximum energy consumption over a specific period. While such pricing mechanisms are commonly employed, strategies that directly account for these costs are not as well explored. Therefore we leverage concepts from *inventory control* to analyze how peak pricing influences the optimal scheduling policies.

To focus on this, we simplify the operational process of the industrial refrigeration, and solely focus on the relationship between the total cooling and energy expenditures of the refrigeration process. The primary system state is the facility temperature, denoted as $x_t \in \mathbb{R}$ for a given time $t \in \{1, \dots, T\}$. Here, T represents the length of the horizon under consideration. Using a first-order model of specific heat², we can succinctly describe the dynamics of the facility's temperature as follows:

$$x_{t+1} = x_t - u_t + q_t, \quad (8.1)$$

where $u_t \geq 0$ represents the heat removed from the facility via the refrigeration system and q_t represents the heat influx from the surrounding environment. We assume that over the horizon, the incoming heat $q_t \geq 0$ is a non-negative random variable that is drawn from a known distribution \mathbf{Q}_t that may be time varying. As an illustration, we

²According to the specific heat equation, $x_{t+1} - x_t = C_f Q_{\text{net}}$, where Q_{net} is the net heat transfer and C_f is the heat capacity. For simplicity of notation, we assume that $C_f = 1$ for the manuscript.

present a scatter plot of possible heat demands over a day³ in Figure 8.2.




Figure 8.2: We present distributions and respective averages of \mathbf{Q}_t over the horizon of a day for a particular refrigeration facility.

The control task of the refrigeration system is to generate a sequence of refrigeration loads $\mathbf{u} \equiv \mathbf{u}_1, \dots, \mathbf{u}_T$ that minimizes the overall system cost, dependent on power consumption, while providing the necessary heat extraction. For a food storage facility, the required refrigeration comes in the form of a temperature constraint $x_t \leq 0$, where we would like to maintain temperatures to be below freezing. We capture both the energy costs and temperature constraint violations in the following stage cost,

$$c_t(x_t, u_t) = o_t(u_t) + \mathbb{E}_{q_t} [h_t(x_{t+1})]. \quad (8.2)$$

Here, $o_t : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ describes the cost associated with running the refrigeration at a load of u_t and $h_t : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ represents the penalty costs for set-point deviations from

³Over June 2023, power and coefficient of power (COP) estimates were collected from a refrigeration site of Butterball LLC $\text{\textcircled{R}}$ to estimate the incoming refrigeration loads.




Figure 8.3: We display the power-heat curves of different compressor types as taken from [1]. We see that for compressors without variable frequency drives (VFDs), the power draw and the respective thermal capacity share an affine relationship.

the desired temperature $x_t = 0$. For each t , we assume that h_t is continuous, convex, and has a minimum at 0, i.e. our desired temperature set point. If $x > 0$ is positive, $h_t(x)$ represents the penalty cost for temperature constraint violations. If $x < 0$ is negative, then $h_t(x)$ represents the cost of excessive cooling which may lead to system inefficiencies. An example of a reasonable penalty function is

$$h_t(x) = \begin{cases} b_t \cdot x^2 & \text{if } x \geq 0 \\ d_t \cdot x^2 & \text{if } x < 0. \end{cases} \quad (8.3)$$

with $b_t \gg d_t$ for each t . Since the energy costs are proportional to the total power consumption of the refrigeration system, we assume an affine structure for the energy

costs with respect to a given refrigeration load, written as

$$o_t(u_t) = \begin{cases} K + a_t \cdot u_t & \text{if } u_t > 0 \\ 0 & \text{if } u_t = 0, \end{cases} \quad (8.4)$$

where K represents the setup cost of having the refrigeration system in operation and a_t determines the per-unit cost for refrigeration capacity at time t . The per-unit cost a_t may depend heavily on t , representing a potential time-of-use cost structure. This cost model is reflected in the power-heat curves of the compressors, which represent the majority ($\sim 90\%$) of the energy consumption of the refrigeration process, as shown in Figure 8.3.

The cost for peak pricing is reflected in the maximum refrigeration load over the horizon. More formally, the peak price can be written as $P \cdot \max\{u_t \text{ for } t \leq T\}$, where $P \geq 0$ is the scaling factor associated with the peak costs. To represent peak cost as a terminal cost, we can introduce an auxiliary state variable $y \in \mathbb{R}_{\geq 0}$ with dynamics

$$y_{t+1} = \max\{y_t, u_t\}. \quad (8.5)$$

We note that while the structure of problem without peak pricing is classical in the inventory control literature (for this, see [129–131] for similar results with regards to convex, piecewise affine o_t), the addition of peak pricing is a novel consideration. While not as studied, peak pricing is extremely significant to the cost structures for industrial processes.

Consolidating the stage and terminal costs, the total cost of refrigeration process

under the sequence of loads \mathbf{u} with a given peak y is

$$J(x, y, \mathbf{u}) = \mathbb{E} \left[\sum_{t=1}^T c_t(x_t, u_t) \right] + P \max \left\{ y, \{u\}_{t \leq T} \right\}, \quad (8.6)$$

where x_t follows from the respective state transition probabilities from the initial state $x_1 = x$ and the expectation is taken over the possible incoming heat $q_t \sim \mathbf{Q}_t$. The optimal total cost can be written recursively via the Bellman equation,

$$V_t(x, y) = \min_{u \geq 0} \left\{ c_t(x, u) + \mathbb{E}_{q_t} \left[V_{t+1}(x^+ + q_t, y^+) \right] \right\}, \quad (8.7)$$

$$V_{T+1}(x, y) = P \cdot y. \quad (8.8)$$

V_{T+1} represents the terminal cost of the dynamic program. We use $x^+ = x - u$ and $y^+ = \max\{y, u\}$ to denote the successor states for simplicity of notation. The optimal loads can be written in feedback form with a policy function $\pi_t^* : \mathbb{R} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ as the argument to the previous optimization formulation:

$$\pi_t^*(x, y) \in \arg \min_{u \geq 0} \left\{ c_t(x, u) + \mathbb{E}_{q_t} \left[V_{t+1}(x^+, y^+) \right] \right\}. \quad (8.9)$$

The main concern of this chapter is on characterizing the structure of these optimal policies with respect to peak pricing. We do this through analytical characterizations in Section 8.3 and through simulations in Section 8.4.

8.3 Results on Compressor Scheduling

We first characterize the structure of optimal refrigeration policies with no peak pricing costs, i.e. when $P = 0$ in Eq. (8.6), where there are only TOU costs present. In this case, our models align with the standard ones present in inventory control, and we can

directly invoke classical results to get the structure of the optimal policies. Interestingly, the optimal policies simplify to a threshold strategy.

Proposition 14 ([132]). *Consider the industrial refrigeration problem with a total cost in Eq. (8.6) with no peak cost ($P = 0$). The optimal policy π_t^* is a threshold policy of the form*

$$\pi_t^*(x, y) = \begin{cases} x - S_t & \text{if } x > s_t, \\ 0 & \text{if } x \leq s_t, \end{cases} \quad (8.10)$$

for some $s_t \geq S_t \in \mathbb{R}$ for every $t \leq T$.

Proof. We can directly use classical results in inventory control with sign changes (for e.g., see [132], [128, Chapter 2.6], or [133, Chapter 4]) to show the claim. \square

From the above proposition, we see that the optimal structure comes in the form of a simple threshold policy, where the refrigeration system is turned off if the facility is below a certain buffer temperature s_t but if turned on, sets the facility to a lower buffer temperature S_t . We remark that if there is no setup costs, or that $K = 0$ in Eq. (8.6), the two buffer temperatures align, with $s_t = S_t$.

Threshold policies are commonly implemented in industrial refrigeration to schedule refrigeration loads. While we have shown that these threshold policies are optimal without peak pricing, computing closed form solutions for s_t and S_t is not possible in general. However, these buffer temperatures can be derived through data-driven strategies to produce well-performing threshold parameters. We do this in Section 8.4.

When peak pricing is introduced to the total cost, however, we will see that the optimal policies stop corresponding to threshold policies. In fact, for simple models of refrigeration, we see that the optimal policies can become quite complicated in the next example.

Example 8 (Peak Pricing). Consider a toy refrigeration scenario, where there is external heat input $q_t \equiv 0$ for all t . Let the temperature penalty function be $h_t(x) = b|x|$ for all $t \leq T$. We set the parameters $P \geq b \geq a = 1$ for the peak, penalty and refrigeration costs, and fix the setup cost to $K = 0$ to frame the problem naturally. Even within this basic model, the optimal policy can be surprisingly intricate.

To see this, we solve for the value function explicitly in Eq. (8.7) via the backwards recursion. Under the model assumptions, the value function V_T at time T can be simply expressed as

$$V_T(x, y) = \min_{u \geq 0} \left\{ u + b|x - u| + P \max\{y, u\} \right\}.$$

From the value function, there are three regimes to describe the optimal policy. These three regimes are depicted in Figure 8.4, where the optimal policy is explicitly characterized below.

$$\pi_T^*(x, y) = \begin{cases} 0 & \text{if } x \leq 0, \\ x & \text{if } 0 \leq x \leq y, \\ y & \text{if } x \geq y. \end{cases}$$

The optimal policy across one time step is a threshold policy with a cap at $u = y$ due to the peak cost. However, adding another step in the backwards recursion complicates the optimal policy π_{T-1}^* . For the next step in the value recursion, we have that

$$V_{T-1}(x, y) = \min_{u \geq 0} \left\{ u + b|x - u| + V_T(x^+, y^+) \right\},$$

where the value function V_T can be described in closed form via the characterization of

Optimal Policy For Different Regimes




Figure 8.4: We depict the possible optimal inputs for each of the three regimes, dependent on x . Note that since $P \geq b$, the optimal u^* must live between $[0, y]$, as delineated by the orange markers.

the optimal policy π_T^* as

$$V_T(x, y) = \begin{cases} Py - bx & \text{if } x \leq 0, \\ Py + x & \text{if } 0 \leq x \leq y, \\ (P + 1 - b)y + bx & \text{if } x \geq y. \end{cases}$$

With the expression of the value function V_T , we can describe the optimal policy for a horizon of 2. Since V_T is piecewise-affine, the optimal inputs must occur at the boundary conditions, where $u = 0$ or $u = x/2$ or $u = y$ or $u = x$. According to this, we can solve

for the optimal policy algebraically to be

$$\pi_{T-1}^*(x, y) = \begin{cases} 0 & \text{if } x \leq 0, \\ x & \text{if } 0 \leq x \leq y, \\ y & \text{if } y \leq x \leq 2y \text{ or} \\ & \text{if } 2y \leq x \text{ and} \\ & (P + 2 - 3b)y \leq \left(\frac{P}{2} - \frac{3b}{2} + 1\right)x, \\ \frac{x}{2} & \text{otherwise.} \end{cases}$$

As we can see, even with a horizon of two steps, the optimal policy differs greatly from the original threshold policy in Eq. (8.10).

While generating closed form expressions for optimal policies are hard to do in general (as seen in Example 8), if we limit to a horizon of 1, we have that a modified threshold policy depicted in Figure 8.5 is optimal. We characterize the structure in the following Proposition.

Optimal Policy For 1-step Horizon




Figure 8.5: We depict the modified policy.

Proposition 15. *Consider the industrial refrigeration problem with a total cost in Eq. (8.6) with a horizon $T = 1$ with no set up cost ($K = 0$ in Eq. (8.4)). The optimal policy π_1^* is characterized by*

$$\pi_1^*(x, y) = \begin{cases} x - \hat{S} & \text{if } x \geq \hat{S} + y, \\ y & \text{if } x \in [S + y, \hat{S} + y], \\ x - S & \text{if } x \in [S, S + y], \\ 0 & \text{if } x \leq S. \end{cases} \quad (8.11)$$

for some $s \geq S \in \mathbb{R}_{\geq 0}$ and $\hat{S} \geq S \in \mathbb{R}_{\geq 0}$.

Proof. First, we express the value function V_1 as below

$$V_1(x, y) = \min_{u > 0} \left\{ f(x^+) + Py^+ \right\} + ax,$$

with $f(z) \triangleq \mathbb{E}[h_t(z + q_t)] - az$ being a convex scalar function. Let S be the minimum of f . Note that if $x < S$, then by convexity of $f(z)$, the optimal input is $u = 0$. Additionally, if $S \leq x \leq S + y$, the optimal input is $u = x - S$. Now let \hat{S} be such that the subderivative $\partial f / \partial z|_{z=\hat{S}} \ni P$. Since f is convex, the minimum is defined by the first order condition on the subderivative $-\partial f / \partial z|_{x-u} + P \ni 0$. This gives the first two conditions of the optimal policy. \square

While we can compute the optimal one-step policy in closed form, this is not true in general, as shown in Example 8. However, as shown in the next Proposition, we can still derive a qualitative threshold-like characterization of the optimal policy in general. We depict the optimal policy pictorially in Figure 8.6.

Proposition 16. *Consider the industrial refrigeration problem with a total cost in Eq.*

(8.6) with no set up cost ($K = 0$ in Eq. (8.4)). The optimal policy π_t^* is characterized by

$$\pi_t^*(x, y) \in \begin{cases} [y, z_t^*] & \text{if } g_t(y) + y \leq x, \\ \{x - g_t(y)\} & \text{if } x - y \leq g_t(y) \leq x, \\ \{0\} & \text{if } x \leq g_t(y), \end{cases} \quad (8.12)$$

where $g_t(y) : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ and z_t^* is such that $x = g_t(z_t^*) + z_t^*$.

Proof. By assumption, we have that $K = 0$ in Eq. (8.6). First we show that for every t , V_t is non-decreasing in y via backwards induction. For the base case, $V_{T+1}(x, y) = Py$ is non-decreasing in y , as $P > 0$. From Eq. (8.7), we have that $V_{t+1}(x^+, y^+)$ is a composition of non-decreasing functions in y by the induction assumption. Since monotonicity is preserved under expectations and infimum projections, the iterate V_t is thus non-decreasing in y as well, and we have the claim.

Now we show that V_t is convex for every t . Similarly as before, we show this via backwards induction. For the base case, $V_{T+1}(x, y) = Py$ is affine in y and thus convex. For the inductive case, we have that c_t is convex in x and u . Additionally, V_{t+1} is convex by the induction assumption. Moreover, since x^+ is affine in x and u ; y^+ is convex in y and u ; and V_{t+1} is non-decreasing in y , $V_{t+1}(x^+, y^+)$ is a convex function of x , y , and u . Since convexity is preserved under expectation, $V_t(x, y)$ can be concisely expressed as

$$V_t(x, y) = \min_{u \geq 0} f_t(x^+, y^+) + ax \quad (8.13)$$

where $f_t(x, y) = \mathbb{E}_{q_t}[\mathbf{h}_t(x + q_t) + V_{t+1}(x + q_t, y)] - ax$ is a convex function. Since $\{u \in \mathbb{R} : u \geq 0\}$ is convex, convexity is preserved under the minimization, and thus $V_t(x, y)$ is a convex function. Thus the claim is shown.

Now we describe the optimal policy $\pi_t(x, y)$. Note that we can describe the optimal

policy as

$$\pi_t(x, y) = \arg \min_{u \geq 0} \{f_t(x^+, y^+)\} + ax.$$

where $f_t(x, y)$ is defined as in Eq. (8.13). For each y , define the function $g_t(y)$ to be

$$g_t(y) = \max \left\{ \arg \min_x f_t(x, y) \right\}. \quad (8.14)$$

We describe the three cases for the optimal policy as shown in Figure 8.6 and Eq. (8.12). If $x \leq g_t(y)$, then observe that $f_t(x^+, y^+) \geq f_t(x^+, y)$, since V_{t+1} is non-decreasing in y . Additionally, as f_t is convex, and $x \leq g_t(y)$, we have that $f_t(x^+, y) \geq f_t(x, y)$ for all $u \geq 0$, and thus $\pi_t^*(x, y) = 0$ when $x \leq g_t(y)$. Likewise, if $x - y \leq g_t(y) \leq x$ we have that $f_t(x^+, y^+) \geq f_t(x^+, y) \geq f_t(g_t(y), y)$ for all $u \geq 0$, and thus $\pi_t(x, y) = x - g_t(y)$.

For the third condition $g_t(y) + y \leq x$, we first note that for $u = y$ generates a cost of $f_t(x - y, y) \leq f_t(x - u, y)$ for any $u \leq y$. Thus the optimal policy must satisfy $\pi^*(x, y) \geq y$. Moreover, let $u = z^*$ be the input in which (x^+, y^+) intersects the curve $x = g_t(y)$. Note that $f(x^+, y^+) \geq f(x^+, z^*) \geq f(g_t(z^*), z^*)$ for any $u \geq z^*$, since f is non-decreasing in y and f is convex. Thus the optimal policy must satisfy $\pi^*(x, y) \leq z^*$. \square

Remark 9. From Proposition 16, we note that if $g_t(y) + y \leq x$, the optimal control action u may lie anywhere in the interval $[y, z_t^*]$. Where on this interval depends non-trivially on the relative peak cost P and the specific holding cost h_t and the current time t . However, we maintain a threshold-like policy structure, where below a buffer temperature $g_t(y)$ (dependent on the current peak), the optimal decision is to turn off the refrigeration system.

Example 9. The optimal policy $\pi^*(x, y)$ in Eq. (8.12) may incur non-intuitive dependencies on the current peak value y ; we outline a simplified scenario that outlines this phenomena. Consider a horizon of $T = 2$, where the starting temperature is $x_1 = -2$ and the incoming




Figure 8.6: We depict the possible optimal inputs for each of the three regimes corresponding to Eq. (8.12), given the curve $g_t(y)$.

heat is $q_t = 2$ for all t . Let there be no per-unit cost ($a = 0$) and setup cost ($K = 0$). Furthermore, we set the peak costs to $P \gg 1$ and the temperature penalty function to be

$$h_t(x) = \begin{cases} bx & \text{if } x \geq 0 \\ -x & \text{if } x < 0, \end{cases}$$

where $b \gg 1$. We compare two scenarios: where the current peak is $y = 1$ and the current peak is $y = 2$. In the first case ($y = 1$), observe that the optimal control sequence is $u_1^* = 1$ and $u_2^* = 1$ to not incur any temperature violation costs. However, when $y = 2$, the optimal control sequence is $u_1^* = 0$ and $u_2^* = 2$. As such, the optimal policy $\pi_1(x, y)$ actually increases with the current peak y , counter to intuition that increasing that the peak should introduce more conservatism to the optimal policy.

8.4 Simulations

In this section, we evaluate the performance of different policy designs through a case scenario based on the thermal data (also represented in Figure 8.2) acquired from a refrigeration facility owned by Butterball LLC®. We depict the possible incoming heat distributions during the on and off peak hours in Figure 8.7.




Figure 8.7: Incoming heat distributions of on and off peak time periods.

The specifics used in the case study are outlined as follows. We evaluate the total cost as according to Eq. (8.6), where each time step t approximately corresponds to a 12-hour time window and the trajectory cost is evaluated over horizon of a month. If the time t corresponds to the on-peak, then the incoming heat distribution \mathbf{Q}_t is exactly the one depicted in Figure 8.7 (and respectively for the off-peak time periods). We additionally normalize the incoming heat values to Megawatts to prevent numerical issues. We assume that the starting temperature and starting peak value are both 0. We use the temperature violation cost of the form in Eq. (8.3) with $b_t = 20$ and $d_t = 1$. We additionally set the respective coefficients for peak costs to $P = 30$, setup costs to $K = 0$, and per-unit costs to $a = 1$ to reasonably model a possible refrigeration scenario.

Methodology	Total Cost
Static Threshold Policy	36.1
Dynamic Threshold Policy	26.7
Modified Threshold Policy	30.7
Dynamic Modified Threshold Policy	25.5

Table 8.1: Cost of Algorithms

Under this example setup, we evaluate the performance of different types of threshold policies according to the total cost in Eq. (8.6). We first examine the performance of the threshold policy in Eq. (8.10), where $s_t = S_t$ is not time varying. Additionally, we examine the dynamic version, where s_t is allowed to depend on if t corresponds to an on-peak or off-peak period. We also examine the performance of the policies stated in Proposition 14 for time-dependent and time-independent values as well. For these policies, we optimize for the coefficients through a Monte Carlo search. The performance of these algorithms are depicted in Table 8.1.

Bibliography

- [1] J. J. Downs and E. F. Vogel, *A plant-wide industrial process control problem*, *Computers & chemical engineering* **17** (1993), no. 3 245–255.
- [2] M. G. Bell, Y. Iida, *et. al.*, *Transportation network analysis*, .
- [3] T. Le, H. L. Vu, N. Walton, S. P. Hoogendoorn, P. Kovács, and R. N. Queija, *Utility optimization framework for a distributed traffic control of urban road networks*, *Transportation Research Part B: Methodological* **105** (2017) 539–558.
- [4] K. Wang, Z. Ouyang, R. Krishnan, L. Shu, and L. He, *A game theory-based energy management system using price elasticity for smart grids*, *IEEE Transactions on Industrial Informatics* **11** (2015), no. 6 1607–1616.
- [5] Z. Sun *et. al.*, *Comprehensive performance analysis of cascade refrigeration system with two-stage compression for industrial refrigeration*, *Case Studies in Thermal Engineering* **39** (2022) 102400.
- [6] F. Fabrega, J. Rossi, and J. d’Angelo, *Exergetic analysis of the refrigeration system in ethylene and propylene production process*, *Energy* **35** (2010), no. 3 1224–1231.
- [7] W. F. Stoecker, *Industrial refrigeration handbook*. McGraw-Hill Education, 1998.
- [8] M. Leng and M. Parlar, *Game theoretic applications in supply chain management: a review*, *INFOR: Information Systems and Operational Research* **43** (2005), no. 3 187–220.
- [9] C. Papadimitriou and G. Piliouras, *From nash equilibria to chain recurrent sets: An algorithmic solution concept for game theory*, *Entropy* **20** (2018), no. 10 782.
- [10] J. He, Y. Li, H. Li, H. Tong, Z. Yuan, X. Yang, and W. Huang, *Application of game theory in integrated energy system systems: a review*, *Ieee Access* **8** (2020) 93380–93397.
- [11] J. F. Nash *et. al.*, *Equilibrium points in n-person games*, *Proceedings of the national academy of sciences* **36** (1950), no. 1 48–49.

- [12] A. Krause, A. Singh, and C. Guestrin, *Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies*, *Journal of Machine Learning Research* **9** (2008) 235–284.
- [13] M. Stoica and M. Schindehutte, *Understanding adaptation in small firms: Links to culture and performance*, *Journal of Developmental Entrepreneurship* **4** (1999), no. 1 1.
- [14] S. Jeong, R. McGrew, E. Nudelman, Y. Shoham, and Q. Sun, *Fast and compact: A simple class of congestion games*, in *Proceedings of the 20th national conference on Artificial intelligence-Volume 2*, pp. 489–494, 2005.
- [15] P. Avila and C. Mullon, *Evolutionary game theory and the adaptive dynamics approach: adaptation where individuals interact*, *Philosophical Transactions of the Royal Society B* **378** (2023), no. 1876 20210502.
- [16] 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.
- [17] C. Daskalakis, M. Fishelson, and N. Golowich, *Near-optimal no-regret learning in general games*, *Advances in Neural Information Processing Systems* **34** (2021) 27604–27616.
- [18] L. S. Shapley *et. al.*, *A value for n-person games*, *Princeton University Press Princeton* (1953).
- [19] R. Konda, R. Chandan, D. Grimsman, and J. R. Marden, *Balancing asymptotic and transient efficiency guarantees in set covering games*, in *2022 American Control Conference*, IEEE, 2022.
- [20] R. Konda, R. Chandan, D. Grimsman, and J. R. Marden, *Optimal utility design of greedy algorithms in resource allocation games*, *IEEE Transactions on Automatic Control* (2024).
- [21] 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, 2004.
- [22] M. Goemans, V. Mirrokni, and A. Vetta, *Sink equilibria and convergence*, in *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pp. 142–151, IEEE, 2005.
- [23] O. Candogan, I. Menache, A. Ozdaglar, and P. A. Parrilo, *Flows and decompositions of games: Harmonic and potential games*, *Mathematics of Operations Research* **36** (2011), no. 3 474–503.

- [24] C. Papadimitriou and G. Piliouras, *Game dynamics as the meaning of a game*, *ACM SIGecom Exchanges* **16** (2019), no. 2 53–63.
- [25] G. Christodoulou, V. S. Mirrokni, and A. Sidiropoulos, *Convergence and approximation in potential games*, in *Annual Symposium on Theoretical Aspects of Computer Science*, pp. 349–360, Springer, 2006.
- [26] V. Bilò, A. Fanelli, M. Flammini, and L. Moscardelli, *Performances of one-round walks in linear congestion games*, in *International Symposium on Algorithmic Game Theory*, pp. 311–322, Springer, 2009.
- [27] A. Fanelli, M. Flammini, and L. Moscardelli, *The speed of convergence in congestion games under best-response dynamics*, in *International Colloquium on Automata, Languages, and Programming*, pp. 796–807, Springer, 2008.
- [28] 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.
- [29] V. S. Mirrokni and A. Vetta, *Convergence issues in competitive games*, in *Approximation, randomization, and combinatorial optimization. algorithms and techniques*, pp. 183–194. Springer, 2004.
- [30] S. Suri, C. D. Tóth, and Y. Zhou, *Selfish load balancing and atomic congestion games*, *Algorithmica* **47** (2007), no. 1 79–96.
- [31] I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli, *Tight bounds for selfish and greedy load balancing*, in *International Colloquium on Automata, Languages, and Programming*, pp. 311–322, Springer, 2006.
- [32] J. R. Marden and A. Wierman, *Distributed welfare games*, *Operations Research* **61** (2013), no. 1 155–168.
- [33] A. Krause and D. Golovin, *Submodular function maximization.*, *Tractability* **3** (2014) 71–104.
- [34] A. S. Manne, *A target-assignment problem*, *Operations research* **6** (1958), no. 3 346–351.
- [35] M. Gairing, *Covering games: Approximation through non-cooperation*, in *International Workshop on Internet and Network Economics*, pp. 184–195, Springer, 2009.
- [36] J. R. Marden and M. Effros, *The price of selfishness in network coding*, *IEEE Transactions on Information Theory* **58** (2012), no. 4 2349–2361.

- [37] M. Narasimhan, N. Jojic, and J. A. Bilmes, *Q-clustering*, *Advances in Neural Information Processing Systems* **18** (2005).
- [38] 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.
- [39] 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.
- [40] 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.
- [41] 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.
- [42] R. Chandan, D. Paccagnan, and J. R. Marden, *Tractable mechanisms for computing near-optimal utility functions*, *arXiv preprint arXiv:2102.04542* (2021).
- [43] 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* (2021).
- [44] T. Roughgarden, *Intrinsic robustness of the price of anarchy*, in *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pp. 513–522, 2009.
- [45] R. Konda, D. Grimsman, and J. Marden, *Execution order matters in greedy algorithms with limited information*, *arXiv preprint arXiv:2111.09154* (2021).
- [46] G. T. Nguyen and K. Kim, *A survey about consensus algorithms used in Blockchain*, *Journal of Information Processing Systems* **14** (2018), no. 1 101–128.
- [47] S. Kar and J. M. Moura, *Distributed consensus algorithms in sensor networks: Quantized data and random link failures*, *IEEE Transactions on Signal Processing* **58** (2010), no. 3 PART 1 1383–1400, [arXiv:0712.1609].
- [48] D. Ongaro and J. Ousterhout, *In search of an understandable consensus algorithm*, *Proceedings of the 2014 USENIX Annual Technical Conference, USENIX ATC 2014* (2019) 305–319.
- [49] J. G. Wardrop, *Road paper. some theoretical aspects of road traffic research.*, *Proceedings of the institution of civil engineers* **1** (1952), no. 3 325–362.

- [50] H. Von Stackelberg, *Market structure and equilibrium*. Springer Science & Business Media, 2010.
- [51] B. Marzouki, O. Belkahla Driss, and K. Ghédira, *Multi Agent model based on Chemical Reaction Optimization with Greedy algorithm for Flexible Job shop Scheduling Problem*, *Procedia Computer Science* **112** (2017) 81–90.
- [52] B. Gharesifard and S. L. Smith, *On distributed submodular maximization with limited information*, *American Control Conference* **2016-July** (2016) 1048–1053.
- [53] A. Badanidiyuru, B. Mirzasoleiman, A. Karbasi, and A. Krause, *Streaming submodular maximization: Massive data summarization on the fly*, in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 671–680, 2014.
- [54] H. Lin and J. Bilmes, *A class of submodular functions for document summarization*, in *Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, vol. 1, pp. 510–520, 2011.
- [55] A. Singh, W. Kaiser, M. Batalin, A. Krause, and C. Guestrin, *Efficient planning of informative paths for multiple robots*, in *International Joint Conference on Artificial Intelligence*, pp. 2204–2211, 2007.
- [56] M. Corah and N. Michael, *Distributed matroid-constrained submodular maximization for multi-robot exploration: theory and practice*, *Autonomous Robots* **43** (2019), no. 2 485–501.
- [57] G. Arslan, J. R. Marden, and J. S. Shamma, *Autonomous vehicle-target assignment: a game-theoretical formulation*, *Journal of Dynamic Systems, Measurement and Control, Transactions of the ASME* **129** (2007), no. 5 584–596.
- [58] M. Gomez-Rodriguez, J. Leskovec, and A. Krause, *Inferring networks of diffusion and influence*, *ACM Transactions on Knowledge Discovery from Data* **5** (2012), no. 4.
- [59] G. Kim, E. P. Xing, L. Fei-Fei, and T. Kanade, *Distributed cosegmentation via submodular optimization on anisotropic diffusion*, *Proceedings of the IEEE International Conference on Computer Vision* (2011) 169–176.
- [60] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. Vanbriesen, and N. Glance, *Cost-effective outbreak detection in networks*, in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 420–429, 2007.
- [61] A. Clark and R. Poovendran, *A submodular optimization framework for leader selection in linear multi-agent systems*, in *IEEE Conference on Decision and Control*, pp. 3614–3621, IEEE, 2011.

- [62] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák, *Maximizing a monotone submodular function subject to a matroid constraint*, *SIAM Journal on Computing* **40** (2011), no. 6 1740–1766, [9780201398298].
- [63] Y. Filmus and J. Ward, *The power of local search: maximum coverage over a matroid*, in *Symposium on Theoretical Aspects of Computer Science*, pp. 601–612, LIPIcs, 2012.
- [64] U. Feige, *A threshold of $\ln n$ for approximating set cover*, *Journal of the ACM* **45** (1998), no. 4 634–652.
- [65] D. Grimsman, M. S. Ali, J. P. Hespanha, and J. R. Marden, *Impact of information in greedy submodular maximization*, in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pp. 2900–2905, IEEE, 2017.
- [66] L. Bulteau, S. Giraud, and S. Vialette, *Disorders and Permutations*, in *32nd Annual Symposium on Combinatorial Pattern Matching (CPM 2021)* (P. Gawrychowski and T. Starikovskaya, eds.), vol. 191 of *Leibniz International Proceedings in Informatics (LIPIcs)*, (Dagstuhl, Germany), pp. 11:1–11:15, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.
- [67] P. Erdos, A. Rényi, *et. al.*, *On the evolution of random graphs*, *Publ. Math. Inst. Hung. Acad. Sci* **5** (1960), no. 1 17–60.
- [68] N. Reza zadeh and S. S. Kia, *A sub-modular receding horizon solution for mobile multi-agent persistent monitoring*, *Automatica* **127** (2021) 109460.
- [69] B. Du, K. Qian, C. Claudel, and D. Sun, *Jacobi-style iteration for distributed submodular maximization*, *arXiv preprint arXiv:2010.14082* (2020).
- [70] A. Robey, A. Adibi, B. Schlotfeldt, G. J. Pappas, and H. Hassani, *Optimal algorithms for submodular maximization with distributed constraints*, *arXiv xxx* (2019) 1–20, [arXiv:1909.1367].
- [71] R. Konda, R. Chandan, and J. R. Marden, *Mission level uncertainty in multi-agent resource allocation*, in *2021 60th IEEE Conference on Decision and Control (CDC)*, pp. 4521–4526, IEEE, 2021.
- [72] A. Nedić and A. Olshevsky, *Distributed optimization over time-varying directed graphs*, *IEEE Transactions on Automatic Control* **60** (2014), no. 3 601–615.
- [73] N. Li and J. R. Marden, *Designing games for distributed optimization with a time varying communication graph*, in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pp. 7764–7769, IEEE, 2012.

- [74] E. Koutsoupias and C. Papadimitriou, *Worst-case equilibria*, in *Annual Symposium on Theoretical Aspects of Computer Science*, pp. 404–413, Springer, 1999.
- [75] J. C. Harsanyi, *Games with incomplete information played by “bayesian” players, i–iii part i. the basic model*, *Management science* **14** (1967), no. 3 159–182.
- [76] Y. E. Sagduyu, R. A. Berry, and A. Ephremides, *Jamming games in wireless networks with incomplete information*, *IEEE Communications Magazine* **49** (2011), no. 8 112–118.
- [77] E. Billard and S. Lakshmivarahan, *Learning in multilevel games with incomplete information. i*, *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* **29** (1999), no. 3 329–339.
- [78] A. Antoniadis, H. J. Kim, and S. Sastry, *Pursuit-evasion strategies for teams of multiple agents with incomplete information*, in *42nd IEEE International Conference on Decision and Control (IEEE Cat. No. 03CH37475)*, vol. 1, pp. 756–761, IEEE, 2003.
- [79] M. Weber, *A method of multiattribute decision making with incomplete information*, *Management Science* **31** (1985), no. 11 1365–1371.
- [80] M. Aghassi and D. Bertsimas, *Robust game theory*, *Mathematical Programming* **107** (2006), no. 1-2 231–273.
- [81] R. Meir and D. Parkes, *Playing the wrong game: Smoothness bounds for congestion games with behavioral biases*, *ACM SIGMETRICS Performance Evaluation Review* **43** (2015), no. 3 67–70.
- [82] J. R. Marden and A. Wierman, *Distributed welfare games*, *Operations Research* (2008) 1–25.
- [83] R. Chandan, D. Paccagnan, and J. R. Marden, *When smoothness is not enough: toward exact quantification and optimization of the price-of-anarchy*, in *2019 IEEE 58th Conference on Decision and Control (CDC)*, pp. 4041–4046, IEEE, 2019.
- [84] M. Hefeeda and M. Bagheri, *Forest fire modeling and early detection using wireless sensor networks.*, *Ad Hoc Sens. Wirel. Networks* **7** (2009), no. 3-4 169–224.
- [85] R. A. Murphey, *Target-based weapon target assignment problems*, in *Nonlinear assignment problems*, pp. 39–53. Springer, 2000.
- [86] D. Paccagnan, R. Chandan, and J. R. Marden, *Distributed resource allocation through utility design-part i: optimizing the performance certificates via the price of anarchy*, *arXiv preprint arXiv:1807.01333* (2018).

- [87] R. Konda, R. Chandan, and J. R. Marden, *Quality of non-convergent best response processes in multi-agent systems through sink equilibria*, in *2023 62nd IEEE Conference on Decision and Control (CDC)*, pp. 6996–7001, IEEE, 2023.
- [88] T. Roughgarden, *Selfish routing and the price of anarchy*. MIT press, 2005.
- [89] I. Atzeni, L. G. Ordóñez, G. Scutari, D. P. Palomar, and J. R. Fonollosa, *Noncooperative and cooperative optimization of distributed energy generation and storage in the demand-side of the smart grid*, *IEEE transactions on signal processing* **61** (2013), no. 10 2454–2472.
- [90] D. Monderer and L. S. Shapley, *Potential games*, *Games and economic behavior* **14** (1996), no. 1 124–143.
- [91] J. R. Marden, G. Arslan, and J. S. Shamma, *Regret based dynamics: convergence in weakly acyclic games*, in *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, (New York, NY, USA), pp. 1–8, Association for Computing Machinery, 2007.
- [92] B. Swenson, R. Murray, and S. Kar, *On best-response dynamics in potential games*, *SIAM Journal on Control and Optimization* **56** (2018), no. 4 2734–2767.
- [93] S. Hart and A. Mas-Colell, *A simple adaptive procedure leading to correlated equilibrium*, *Econometrica* **68** (2000), no. 5 1127–1150.
- [94] R. D. Kleinberg, K. Ligett, G. Piliouras, and É. Tardos, *Beyond the nash equilibrium barrier.*, in *ICS*, pp. 125–140, 2011.
- [95] K. Basu and J. W. Weibull, *Strategy subsets closed under rational behavior*, *Economics Letters* **36** (1991), no. 2 141–146.
- [96] A. Fabrikant and C. H. Papadimitriou, *The complexity of game dynamics: Bgp oscillations, sink equilibria, and beyond.*, in *SODA*, vol. 8, (USA), pp. 844–853, Citeseer, Society for Industrial and Applied Mathematics, 2008.
- [97] R. Yan, X. Duan, Z. Shi, Y. Zhong, J. R. Marden, and F. Bullo, *Policy evaluation and seeking for multi-agent reinforcement learning via best response*, *IEEE Transactions on Automatic Control* **67** (2021), no. 4 1898–1913.
- [98] G. Notomista, Y. Emam, and M. Egerstedt, *The slothbot: A novel design for a wire-traversing robot*, *IEEE Robotics and Automation Letters* **4** (2019), no. 2 1993–1998.
- [99] V. Chvatal, *A greedy heuristic for the set-covering problem*, *Mathematics of operations research* **4** (1979), no. 3 233–235.

- [100] L. Gourves and J. Monnot, *On strong equilibria in the max cut game*, in *International Workshop on Internet and Network Economics*, pp. 608–615, Springer, 2009.
- [101] T. Roughgarden, *The price of anarchy in games of incomplete information*, *ACM Transactions on Economics and Computation (TEAC)* **3** (2015), no. 1 1–20.
- [102] O. Candogan, A. Ozdaglar, and P. A. Parrilo, *Dynamics in near-potential games*, *Games and Economic Behavior* **82** (2013) 66–90.
- [103] R. Konda, V. Chandan, J. Crossno, B. Pollard, D. Walsh, R. Bohonek, and J. R. Marden, *Utilizing load shifting for optimal compressor sequencing in industrial refrigeration*, *arXiv preprint arXiv:2403.07831* (2024).
- [104] *Manufacturing energy consumption survey*, tech. rep., US Energy Information Administration, 2018.
- [105] Y. Sun, S. Wang, F. Xiao, and D. Gao, *Peak load shifting control using different cold thermal energy storage facilities in commercial buildings: A review*, *Energy conversion and management* **71** (2013) 101–114.
- [106] J. E. Braun, *Load control using building thermal mass*, *J. Sol. Energy Eng.* **125** (2003), no. 3 292–301.
- [107] A. Afram and F. Janabi-Sharifi, *Theory and applications of hvac control systems—a review of model predictive control (mpc)*, *Building and Environment* **72** (2014) 343–355.
- [108] Y. Yao and D. K. Shekhar, *State of the art review on model predictive control (mpc) in heating ventilation and air-conditioning (hvac) field*, *Building and Environment* **200** (2021) 107952.
- [109] R. Pattison, C. Touretzky, T. Johansson, I. Harjunoski, and M. Baldea, *Optimal process operations in fast-changing electricity markets: framework for scheduling with low-order dynamic models and an air separation application*, *Industrial & Engineering Chemistry Research* **55** (2016) 4562–4584.
- [110] R. Pattison, C. Touretzky, I. Harjunoski, and M. Baldea, *Moving horizon closed-loop production scheduling using dynamic process models*, *AIChE Journal* **63** (2017) 639–651.
- [111] A. Vishwanath, V. Chandan, and K. Saurav, *An iot-based data driven precooling solution for electricity cost savings in commercial buildings*, *IEEE Internet of Things Journal* **6** (2019), no. 5 7337–7347.

- [112] L. S. Larsen, C. Thybo, J. Stoustrup, and H. Rasmussen, *Control methods utilizing energy optimizing schemes in refrigeration systems*, in *2003 European Control Conference (ECC)*, pp. 1973–1977, IEEE, 2003.
- [113] K. A. Manske, *Performance optimization of industrial refrigeration systems*, .
- [114] L. F. S. Larsen, *Model based control of refrigeration systems*. Department of Control Engineering, Aalborg University, 2006.
- [115] D. T. Reindl *et. al.*, *Sequencing & control of compressors*, *ASHRAE Journal* **55** (2013), no. 11 14.
- [116] K. Manske, D. Reindl, and S. Klein, *Evaporative condenser control in industrial refrigeration systems*, *International journal of refrigeration* **24** (2001), no. 7 676–691.
- [117] R. Konda, J. Prescott, V. Chandan, J. Crossno, B. Pollard, D. Walsh, R. Bohonek, and J. R. Marden, *Efficient industrial refrigeration scheduling with peak pricing*, *arXiv preprint arXiv:2405.20433* (2024).
- [118] I. Dincer, *Refrigeration systems and applications*. John Wiley & Sons, 2017.
- [119] S. Tassou, J. S. Lewis, Y. Ge, A. Hadawey, and I. Chaer, *A review of emerging technologies for food refrigeration applications*, *Applied Thermal Engineering* **30** (2010), no. 4 263–276.
- [120] L. Zhao, W. Cai, X. Ding, and W. Chang, *Model-based optimization for vapor compression refrigeration cycle*, *Energy* **55** (2013) 392–402.
- [121] T. G. Hovgaard, S. Boyd, L. F. Larsen, and J. B. Jorgensen, *Nonconvex model predictive control for commercial refrigeration*, *International Journal of Control* **86** (2013), no. 8 1349–1366.
- [122] X.-H. Yin and S.-Y. Li, *Model predictive control for vapor compression cycle of refrigeration process*, *International Journal of Automation and Computing* **15** (2018), no. 6 707–715.
- [123] S. E. Shafiei, J. Stoustrup, and H. Rasmussen, *Model predictive control for flexible power consumption of large-scale refrigeration systems*, in *2014 American Control Conference*, pp. 412–417, IEEE, 2014.
- [124] M. J. Risbeck and J. B. Rawlings, *Economic model predictive control for time-varying cost and peak demand charge optimization*, *IEEE Transactions on Automatic Control* **65** (2019), no. 7 2957–2968.

- [125] Y. Mo, Q. Lin, M. Chen, and S.-Z. J. Qin, *Optimal online algorithms for peak-demand reduction maximization with energy storage*, in *Proceedings of the twelfth ACM international conference on future energy systems*, pp. 73–83, 2021.
- [126] F. Oldewurtel, A. Ulbig, A. Parisio, G. Andersson, and M. Morari, *Reducing peak electricity demand in building climate control using real-time pricing and model predictive control*, in *49th IEEE conference on decision and control (CDC)*, pp. 1927–1932, IEEE, 2010.
- [127] S. Axsäter, *Inventory control*, vol. 225. Springer, 2015.
- [128] B. Liu and A. O. Esogbue, *Decision criteria and optimal inventory processes*, vol. 20. Springer Science & Business Media, 2012.
- [129] E. L. Porteus, *Stochastic inventory theory, Handbooks in operations research and management science* **2** (1990) 605–652.
- [130] Y. Lu and M. Song, *Inventory control with a fixed cost and a piecewise linear convex cost*, *Production and Operations Management* **23** (2014), no. 11 1966–1984.
- [131] M. J. Sobel, *Making short-run changes in production when the employment level is fixed*, *Operations Research* **18** (1970), no. 1 35–51.
- [132] H. Scarf, K. Arrow, S. Karlin, and P. Suppes, *The optimality of (s, s) policies in the dynamic inventory problem*, *Optimal pricing, inflation, and the cost of price adjustment* (1960) 49–56.
- [133] J.-S. J. Song, *Research handbook on inventory management*. Edward Elgar Publishing, 2023.