Title
On the Role of Information in the Control of Multi-Agent Systems: A Game Theoretic Approach

Author
Ferguson, Bryce Lowell

Publication Date
2024

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

On the Role of Information in the Control of Multi-Agent Systems: A Game Theoretic Approach

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

Doctor of Philosophy
in
Electrical and Computer Engineering

by

Bryce L. Ferguson

Committee in charge:

Professor Jason R. Marden, Chair
Professor João P. Hespanha
Professor Mahnoosh Alizadeh
Professor Philip N. Brown, University of Colorado, Colorado Springs
Professor Bary S. R. Pradelski, National Centre for Scientific Research
Professor Ketan Savla, University of Southern California

September 2024

The Dissertation of Bryce L. Ferguson is approved.

Professor João P. Hespanha

Professor Mahnoosh Alizadeh

Professor Philip N. Brown, University of Colorado, Colorado Springs

Professor Bary S. R. Pradelski, National Centre for Scientific Research

Professor Ketan Savla, University of Southern California

Professor Jason R. Marden, Committee Chair

May 2024

On the Role of Information in the Control of Multi-Agent Systems:

A Game Theoretic Approach

A smooth sea never made a skilled sailor.

- Franklin D. Roosevelt

To my many supporters,

Acknowledgements

The contributions of this thesis were not accomplished alone, both in practice and in spirit. I have been honored to work with and alongside many amazing researchers, and I have also been encouraged and supported by many friends and family members. I would first like to acknowledge the researchers I have had the pleasure of interacting with, discussing research topics, and often coauthoring papers; there are too many to name individually, but each of you has contributed to the success of and satisfaction with my time in graduate school. I would like to acknowledge my committee members for their support and the guidance they have provided me in writing the works contained in this thesis. Mahnoosh, thank you for introducing me to research by inviting me into your group during my undergraduate; without you, I would have never had the knowledge or confidence to apply to a PhD program. Philip, thank you for working with me so closely at the beginning of my PhD and being a mentor; your advice and guidance have helped me to understand how to be an effective researcher. Bary, thank you for hosting me in Grenoble; it was a fantastic research and cultural opportunity. Ketan, thank you for the insightful research discussions and for providing feedback on my work. Joao, thank you for teaching me various topics in control and demonstrating the excellence as a professor I one day hope to resemble. Most importantly, I would like to acknowledge my PhD adviser, Jason. I have had a fantastic PhD experience both professionally and personally, and that is largely thanks to you. I have always appreciated that you always put the interests of your students first. I could not imagine a better relationship with a PhD adviser, and I will be forever grateful to have spent these years of my life working with you.

Outside of my committee, I would like to thank my labmates and the students of the CCDC at UCSB. Carpooling down to SoCal workshops together and sharing tea after seminars are just some of the memories I will cherish. I have made many friendships that I hope will last a lifetime. I also owe significant appreciation to my family. I am very fortunate to be supported by so many loving people. My cousin, Sergei, has always been a role model to me, and I am so grateful that I got to spend so much time with him growing up. My sister, Shannon, is one of the most intelligent people I have ever met, and growing up with her gave me the perspective and passion to continue my education to this point. My father, Mark, showed me what it meant to be responsible, confident, and diligent. The principles you have instilled in me have guided me through my life. My mother, Laurie, has always shown me such incredible love and support. You being there to regularly talk through my thoughts and help center my focus has been critical to dealing with the ups and downs of this PhD. Finally, Jenny, without going into too much detail about how much you mean to me, I will simply state that this PhD is only the second most important thing I have gained from my time in Santa Barbara.

Curriculum Vitæ
Bryce L. Ferguson

Education

2024	Ph.D. in Electrical and Computer Engineering, University of California, Santa Barbara.
2019	M.S. in Electrical and Computer Engineering, University of California, Santa Barbara.
2018	B.S. in Electrical and Computer Engineering, University of California, Santa Barbara.
2016	A.A. in Mathematics, Santa Rosa Junior College.

Publications

Journal Articles

1. B. L. Ferguson, D. Paccagnan, B. S. R. Pradelski, and J. R. Marde, "Collaborative Decision-Making and the k-Strong Price of Anarchy in Common Interest Games," *preprint available on ArXiv* [1]
2. B. L. Ferguson, P. N. Brown, and J. R. Marden, "Information Signalling with Concurrent Monetary Incentives in Bayesian Congestion Games," *IEEE Transactions on Intelligent Transportation Systems (T-ITS)*, 2024 [2]
3. B. L. Ferguson, D. Paccagnan, and J. R. Marden, "The Cost of Informed Decision Making in Multi-Agent Maximum Coverage Problems," *IEEE Control Systems Letters (L-CSS)*, 2023 [3]
4. B. L. Ferguson and J. R. Marden, "Robust Utility Design in Distributed Resource Allocation Problems with Defective Agents," *Dynamic Games and Applications (DGAA)*, 2022 [4]
5. B. L. Ferguson, P. N. Brown, and J. R. Marden, "The Effectiveness of Subsidies and Tolls in Congestion Games," *IEEE Transactions on Automatic Control (TAC)*, 2021 [5]
6. B. L. Ferguson, P. N. Brown, and J. R. Marden, "The Effectiveness of Subsidies and Taxes in Atomic Congestion Games," *IEEE Control Systems Letters (L-CSS)*, 2021 [6]
7. D. Paccagnan, R. Chandan, B. L. Ferguson, and J. R. Marden, "Optimal taxes in atomic congestion games," *ACM Transactions on Economics and Computation (TEAC)*, 2021 [7]
8. B. L. Ferguson, P. N. Brown, and J. R. Marden, "Value of Information in Incentive Design: A Case-Study in Simple Congestion Networks," *IEEE Transaction on Computational Social Systems (TCSS)*, 2023 [8]

Proceedings and Refereed Conferences

1. B. L. Ferguson, D. Paccagnan, B. S. R. Pradelski, and J. R. Marden, "Bridging the Gap Between Central and Local Decision-Making: The Efficacy of Collaborative Equilibria in Altruistic Congestion Games," *IEEE Conference on Decision and Control (CDC)*, 2024. (under review)
2. V. Shah, B. L. Ferguson, and J. R. Marden, "Learning Optimal Stable Matches in Decentralized Markets with Unknown Preferences," *IEEE Conference on Decision and Control (CDC)*, 2024. (under review)
3. B. L. Ferguson, D. Paccagnan, Bary S. R. Pradelski, and J. R. Marden, "Collaborative Coalitions in Multi-Agent Systems: Quantifying the Strong Price of Anarchy for Resource Allocation Games," IEEE Conference on Decision and Control (CDC), 2023 [9]
4. R. Zhang, Y. Zhang, R. Konda, B. L. Ferguson, J. R. Marden, N. Li, "Markov Games with Decoupled Dynamics: Price of Anarchy and Sample Complexity," *IEEE Conference on Decision and Control (CDC)*, 2023 [10]
5. A. K. Chen*, B. L. Ferguson*, D. Shishika, M. Dorothy, J. R. Marden, G. J. Pappas, and V. Kumar, "Path Defense in Dynamic Defender-Attacker Blotto Games (dDAB) with Limited Information," *American Control Conference (ACC)*, 2023 [11]
6. B. L. Ferguson, D. Shishika and J. R. Marden, "Ensuring the Defense of Paths and Perimeters in Dynamic Defender-Attacker Blotto Games (dDAB) on Graphs," *58th Annual Allerton Conference on Communication, Control, and Computing*, 2022 [12]
7. B. L. Ferguson, P. N. Brown and J. R. Marden, "Avoiding Unintended Consequences: How Incentives Aid Information Provisioning in Bayesian Congestion Games," *IEEE Conference on Decision and Control (CDC)*, 2022 [13]
8. B. L. Ferguson and J. R. Marden, "Robust Utility Design in Distributed Resource Allocation Problems with Defective Agents," *IEEE Conference on Decision and Control (CDC)*, 2021 [14]
9. Yixiao Yue, B. L. Ferguson and J. R. Marden, "Incentive Design for Congestion Games with Unincentivizable Users," *IEEE Conference on Decision and Control (CDC)*, 2021 [15]
10. B. L. Ferguson and J. R. Marden, "The Impact of Fairness on Performance in Congestion Networks," *American Control Conference (ACC)*, 2021 [16]
11. B. L. Ferguson, P. N. Brown, and J. R. Marden, "Carrots or Sticks? The Effectiveness of Subsidies and Tolls in Congestion Games," *American Control Conference (ACC)*, 2020 [17]
12. R. Chandan, D. Paccagnan, B. L. Ferguson, and J. R. Marden, "Computing Optimal Taxes in Atomic Congestion Games," *NetEcon*, 2019 [18]

13. B. L. Ferguson, P. N. Brown, and J. R. Marden, "Utilizing Information Optimally to Influence Distributed Network Routing," *IEEE Conference on Decision and Control (CDC)*, 2019 [19]
14. N. Tucker, B. Ferguson, and M. Alizadeh, "An Online Pricing Mechanism for Electric Vehicle Parking Assignment and Charge Scheduling," *American Control Conference (ACC)*, 2019 [20]
15. B. Ferguson, V. Nagaraj, E. C. Kara, and M. Alizadeh, "Optimal Planning of Workplace Electric Vehicle Charging Infrastructure with Smart Charging Opportunities," *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, 2018 [21]

Awards & Honors

- 2023 CSS Student Travel Award for the IEEE Conference on Decision and Control
- 2023 UCSB Graduate Division Dissertation Fellowship Award
- 2023, 2022, 2021, 2020, 2019 UCSB ECE Outstanding TA Award
- 2022 Invited Research Fellow from the French Centre National de la Recherche Scientifique
- 2022 Cyber-Physical Systems Rising Star
- 2022 NSF ASI Scholar on Industrial Risk Management at INSA CVL and Inria Lille, France
- 2022-2023 UCSB ECE Department Dissertation Fellowship Award
- 2020 ACC Best Student Paper Award Finalist
- 2019 CSS Student Travel Award for the IEEE Conference on Decision and Control
- 2018 Graduated with Highest Honors, University of California, Santa Barbara
- 2016-2018 ESTEEM Scholar with full tuition scholarship, University of California, Santa Barbara
- 2016-2018 Engineering Honor Society, University of California, Santa Barbara
- 2016-2018 Dean's Honors, College of Engineering, University of California, Santa Barbara
- 2016 Graduated with Highest Honors, Santa Rosa Junior College
- 2014-2016 Dean's List, Santa Rosa Junior College

Abstract

On the Role of Information in the Control of Multi-Agent Systems:
A Game Theoretic Approach

by

Bryce L. Ferguson

Due to the emergence of new communication and computation technologies, many existing systems and infrastructures are experiencing revolutions in their behavior and capabilities. For example, with the development of self-driving vehicles (which possess the power to automate driving decisions and coordinate with other automated vehicles on the road), new traffic patterns take place, and the opportunity to introduce safer and more efficient driving behavior presents itself. Similarly, with the proliferation of large-language AI models and access to social media, information can be garnered and exchanged in new (sometimes unreliable) ways. These systems, among many others, consist of many human and engineered entities (or agents) interacting and making decisions with the limited information they possess about one another and the environment they are in; we will refer to these systems as *multi-agent systems*. The main goal of this thesis is to introduce new understandings of the role information plays in several aspects of multi-agent systems, i.e., the information a system designer possesses about the agents, the agents' knowledge of the environment, and the agents' ability to share information and coordinate behavior among themselves. The contributions are split into two parts: the first studies the interactions of designed decision-makers in distributed autonomous systems, and the second studies behavior in social systems in which decision-makers are human users. In each setting, we can model the interactions between agents (human or engineered) through the mathematical framework of game theory. The findings of this

work reveal several insights on the role of information in multi-agent systems, including 1) showing to what extent greater information on human-agent preferences and engineered-agent capabilities can aid in the design phase, 2) proving that revealing information to agents (human or designed) can worsen system performance unless done carefully, and 3) quantifying the benefits and costs incurred by increasing the level of communication and collaboration among agents. These insights will be shown through rigorous analysis of several game-theoretic models.

Contents

Chapter 1

Introduction

1.1 Motivation of This Thesis

The proliferation of high-speed communication and embedded computational intelligence has ushered in a new era for engineers, presenting both opportunities and challenges. In various engineered domains, once complicated and independent processes are being re-imagined by the development of technologies that automate decision-making and the now accessible ability to communicate vast amounts of information very quickly. For example, in the field of transportation, the development of smart and autonomous vehicles has led researchers and city planners to consider new driving behaviors that may be brought on from fleets or platoons of vehicles enabled to communicate and coordinate driving behavior on the road. Similarly, the use cases of robotics increase immensely with the ability to have many small, inexpensive robots cooperatively perform tasks as a fleet. Additionally, individuals' investments in home and commercial power storage and generation threaten to disrupt traditional understandings of demand in the power grid by allowing internet-connected devices to intelligently coordinate charging schedules. These (and many other) emerging technologies form the basis of a new type of system that consists of many interconnected components (termed agents); throughout this work, we will term systems henceforth referred to as *multi-agent systems*.

Though forming such systems is an achievement in and of itself, the ability to effectively design and operate a multi-agent system requires a fundamental understanding of the emergent *collective* behavior of the many sub-systems and agent decision-making processes, which may not align with preexisting intuition. This challenge underscores the purpose of this thesis: developing the fundamental theory of multi-agent systems and garnering a greater understanding of control opportunities therein. The broad goal of studying the behavior of multi-agent systems has been the focus of much recent research; the chapters of this thesis primarily focus on the many roles information plays in determining the emergent behavior of a multi-agent system. Particularly, the formal analysis outlined in this thesis will seek to address the following questions:

- How does access to greater information about a system/environment aid in designing a multi-agent system?
- How does revealing information to individual agents affect collective behavior?
- What improvements in multi-agent coordination are attainable by allowing agents to communicate and collaborate in their decision-making?

These questions on the role of information in multi-agent systems have wide-reaching implications and are the broad focus of this thesis. In each of the previously mentioned multi-agent systems (fleet robotics, traffic routing, and power grids), several aspects of behavior and decision-making can be associated with or attributed to different forms of information. For example, the local sensor measurements for individual drones in a fleet may be a form of collected information about the environment that can serve as an input to designed control algorithms. Deciding how these drones interpret, implement, and share this information will impact the decisions of each individual drone (or agent) and, thus, the emergent behavior of the fleet. Alternatively, in transportation, each human driver may have some personalized response to a new toll lane being added to their commute. The designer of this toll lane may benefit from modeling the resulting traffic patterns, which relies on the amount of information they possess about the population of drivers responding to different incentives that may be deployed. In these two settings, information plays disparate but vital roles in our understanding of collective behavior in multi-agent systems.

To study the behavior of a multi-agent system at a more fundamental level, we require a rigorous set of tools that allow us to model the interactions of many decision-making entities. For this, the analysis presented in this thesis adopts the framework of game theory.

1.2 Preliminaries on Game Theory

At a high level, game theory can be described as a mathematical model for the interactions of rational decision-makers. Though drawing its name from a friendly competition, in the field of game theory, a game is any setting where a set of players interact and make decisions for their own benefit. This can certainly capture two children playing checkers, taking turns moving pieces to improve their chances of victory; however, it can also capture a variety of other interactions, such as those of the agents in our aforementioned multi-agent systems. Formally, a game is defined by a small list of mathematical objects: a set of players $N = \{1, \dots, n\}$, a set of actions for each player $\mathcal{A}_i$, and a utility or cost function that maps the collection of each players' action to some reward $U_i : \mathcal{A}_1 \times \cdots \times \mathcal{A}_n \to \mathbb{R}$. In itself, this model is very simple, but the particular nuances of any multi-agent system can be embedded in these primitives (namely, the utility functions and action sets), making this framework fit for studying a wide variety of decision-making problems.

As a field, game theory has developed rapidly over the last century. In 1928, in the seminal work “On the Theory of Games of Strategy” [22], John Von Neumann studied the interactions of two directly competing players in zero-sum games and proved (under certain conditions) the existence of states where neither player could individually alter their behavior and improve their individual reward. In 1950, John Nash took this idea further and showed that in games with finitely many players (under certain conditions), there exist states (now termed Nash equilibria) where no player can unilaterally deviate their action to improve their reward [23]. Later still, in 1996, Dov Monderer and Lloyd Shapley introduced the class of potential games, in which there exists an alignment of the players’ objectives that gives rise to the existence of pure Nash equilibria and desirable convergence properties [24]; many of the studies in this thesis are within classes of potential games in part thanks to the natural alignment of decision-making in engineered and social multi-agent systems. Great amounts of research have gone on to study the repeated [25], sequential [26], and dynamic [27] nature of decision making, different markets and environments in which agents interact [28], and, more recently, the interactions of multiple learners [29].

Game theory has long had a presence in the field of economics. The immediate application of game theory to problems of competing firms or corporations initiated new disciplines of microeconomics, including mechanism design [30], matching markets [31], the formation of conventions [27], and more. Game theory also helped unify several existing principles such as Cournot competitions of competitively priced goods [32] and welfare economics [33]. However, despite its common association, game theory is not a tool restricted only to economists; the general mathematical framework has more recently been used in several new domains.

At the turn of the millennium, algorithmic game theory emerged in computer science departments. As an extension of research on computer algorithms, game theory provided a blueprint to describe the solution concepts of algorithms run in a distributed fashion. One key concept that propelled the growth of algorithmic game theory was the introduction of the price of anarchy [34], which–defined informally–measures the degradation in performance caused by a system's agents making decisions locally as opposed to centrally. Such a performance metric (or approximation ratio of selfish/distributed behavior) allowed researchers to quantify the effects of previously acknowledged phenomena such as the tragedy of the commons [143], the cost of selfish routing [35], and the inefficiency of distributed algorithms [36]. At the same time, researchers in the field of automation and control were embracing the new solution concept of Nash equilibrium within the context of distributed dynamical systems [37, 38, 39], where largely the stability of and convergence to Nash equilibria was evaluated in various dynamical systems. With the evolution of these two new focuses of game theory research, when used in tandem, one can formally describe the behavior and performance of a variety of multi-agent systems. The expansion of game theory into engineering fields not only allows for cross-disciplinary research but also enables researchers to more effectively address a variety of emerging problems.

1.3 Game Theoretic Problems in Controls

The field of controls seeks to develop rules for algorithms that guide the behavior of various cyber and physical systems. In expanding the breadth of systems to which control theory can be applied and developed, game theory emerges as a framework to help formalize several new challenges:

Distributed Control - In large-scale autonomous systems, it may be difficult or infeasible to design and implement a control policy that simultaneously regulates the numerous system components (e.g., a fleet of autonomous vehicles). To reduce complexity and communication costs, a control designer may opt to implement more simple, local algorithms that will allow each device to operate (e.g., each autonomous vehicle possessing onboard self-driving software). If each local control algorithm is designed to optimize a defined objective, then the interactions of the many automated components can be studied as a mathematical game. Thus, the Nash equilibrium becomes a solution concept that can be used to understand the emergent behavior of these new distributed autonomous systems, such as autonomous vehicles, fleet robotics, and the Internet of Things.

Social Modeling - In many societal-level systems, a system operator may not be empowered to design and automate every facet of the system, rather some decisions may be made by the system's human users. To develop control algorithms that operate in these environments, a system operator needs some understanding of not only how humans have previously behaved but also how they will react to various design choices. For this, game theory again offers relevant opportunities, now in its ability to model the interactions of human decision-makers, around which a control engineer can design. This role of game theory has already proven useful in a variety of domains, particularly in studying shared utilization and congestion in traffic, supply chain resource management, and power demand, where congestion game models the negative consequences of users' overlapping decisions and Wardrop equilibria serve as a simple model for the system's emergent behavior [151].

Robust Optimization - Some problems are posed and solved without even knowing their game theoretic connections. A common problem in engineering is to select system parameters whilst being cognizant of various risks. To formulate this, one may transcribe a minimax problem (of the form $\min_{a \in A} \max_{b \in B} C(a, b)$), in which one seeks to find an optimal parameter (a) that minimizes a cost function (C) given the worst possible realization of some hazard or unknown feature (b). The solution and value of this optimization problem provide a decision and guarantee that is robust to the risk of any hazard occurring. This form of worst-case analysis is valuable in many settings like risk minimization, Chebyshev centers and other spatial reasoning, pursuit evasion, and adversarial security. The minimax problem has grown from the study of zero-sum games [22], in which two players make decisions with oppositely aligned objectives. The equilibria of these games are precisely the solutions to the minimax problem.

These emerging and existing problems have made game theory increasingly relevant in many areas of engineering and computer science. As such, the contributions of this thesis aim to further develop the existing understanding of game theory in control, particularly revolving around newly introduced aspects of information and communication. Specifically, the forthcoming chapters will study how altering the amount of available information and introducing new communication channels can affect the quality of solutions offered via game theoretic approaches. This thesis will be divided into two parts: the first will focus on the interactions of designed, autonomous agents, and the second will focus on human agents with which we have a limited ability to influence.

1.4 Outline of Chapters

The thesis is divided into two parts, each consisting of three chapters. Part 1 is focused on the use of game theory to understand the performance of distributed algorithms in multi-agent systems. Part 2 studies control and influencing mechanisms in social systems modeled by games. In each, the aforementioned aspects of information are explored, namely, what value information possesses at the design phase and what opportunities a system operator has in exploiting communication channels as a method of control. For posterity and the ability of each chapter to stand alone, the chapters are largely unaltered from the published work on which they are based. The contributions of each are outlined as follows:

Part 1: Multi-Agent Optimization and Distributed Resource Allocation

The first part of this thesis contains three chapters, each studying the distinct role of information in systems of many agents designed to optimize a prescribed objective function. Each work offers some general insights but delves further into the model of distributed resource allocation. In a distributed resource allocation problem or resource allocation game, each agent must select which resources to 'utilize' from a shared set of resources. Agents' alignment in their decisions can be penalized or rewarded to model settings where agents' overlap is desirable or to be avoided. In this class of problems, the three chapters are as follows:

Chapter 2: Collaborative Decision-Making and the k-Strong Price of Anarchy in Common Interest Games - This chapter explores the benefits and costs of increasing communication and coordination among agents. By allowing agents to coordinate in their decision-making, new, stronger notions of equilibria become stable. This chapter studies the added benefit to equilibrium welfare and added computational complexity incurred by designing agents to coordinate. The contributions of this chapter originally appeared in [1].

Chapter 3: The Cost of Informed Decision Making in Multi-Agent Maximum Coverage Problems - In many settings, agents need not have full information about their environment. This chapter studies whether increasing the amount of information the agents of a multi-agent system have is beneficial or detrimental. It is shown that increasing the knowledge of the system state can actually induce less desirable equilibria and thus harm system performance. However, the best-case equilibrium always improves. The trade-off between improving the performance of best-case equilibria and worsening worst-case equilibria is explored via utility design. The contributions of this chapter originally appeared in [3].

Chapter 4: Robust Utility Design in Distributed Resource Allocation Problems with Defective Agents - If agents become defective, they may not just cause issues locally but can cascade issues across the agents they interact with. This chapter studies how a subset of the agents becoming defective (i.e., not contributing to the system welfare but still interacting with other agents) can alter the global dynamics. First, the drop in equilibrium welfare caused by the presence of defective agents is bounded. Then, agents' utilities are designed to encourage overlap in the case of defective agents among the group. Finally, the trade-off between designing for the presence or absence of defective agents is explored via utility design. The contributions of this chapter originally appeared in [4].

Part 2: Congestion Control and Social Influencing

In the second part of this thesis, rather than focusing on the interactions of designed autonomous agents, we focus on the challenge of influencing human agents. Human agents possess their own individual preferences and cannot have their behavior directly dictated to them. Instead, to alter the decision-making and collective behavior of social systems, we need to consider the design of mechanisms that influence behavior. Successfully performing this influence relies on the amount of information available to the system operator. We study the facets of information in social influence in several ways, largely in the context of congestion games.

Chapter 5: The Effectiveness of Subsidies and Tolls in Congestion Games - Monetary incentives are a direct approach to influencing human behavior; however, two such incentives exist: subsidies and tolls. In this chapter, we study the relative performance of subsidies and tolls in influencing social behavior under budgetary constraints and when users' responses to incentives are uncertain. Subsidies outperform tolls under similar budgetary constraints with little heterogeneity; however, as the amount of heterogeneity increases, tolls prove to be more robust. The contributions of this chapter originally appeared in [5].

Chapter 6: Information Signalling with Concurrent Monetary Incentives in Bayesian Congestion Games - A less direct mechanism of influencing human behavior is to strategically reveal information as a way to shape the population's beliefs. In this work, we consider a model of Bayesian persuasion where an information source can strategically signal information to agents as a mechanism to shape posterior beliefs and behavior. We show that revealing truthful information can sometimes be harmful. To remedy this, we consider the concurrent use of monetary incentives and show that with appropriately designed incentives, system performance is guaranteed to improve. The contributions of this chapter originally appeared in [2].

Chapter 7: Value of Information in Incentive Design: A Case-Study in Simple Congestion Networks - In designing any influencing mechanism, a system designer's effectiveness is conditioned on their understanding of the system in which they work. In this chapter, we study the capabilities of a toll designer in various informational environments. We quantify the value of different pieces of information (such as information about the population of users or information about the network in which they interact) and provide a comparative analysis of which is more valuable. The contributions of this chapter originally appeared in [8].

This work was supported by Office of Naval Research under Grant #N00014-20-1-2359, Air Force Office of Scientific Research under Grants #FA95550-20-1-0054 and #FA9550-21-1-0203, Army Research Lab through ARL DCIST CRA #W911NF-17-2-0181, National Science Foundation under Grant #ECCS-1638214, French National Research Agency (ANR) under Grant # ANR-19-CE48-0018-01.

Part I

Multi-Agent Optimization and Distributed Resource Allocation

Chapter 2

Collaborative Decision-Making and the k-Strong Price of Anarchy in Common Interest Games

The control of large-scale, multi-agent systems often entails distributing decision-making across the system components. However, with advances in communication and computation technologies, we can consider new collaborative decision-making paradigms that bridge centralized and distributed control architectures. In this section, we seek to understand the benefits and costs of increased collaborative communication in multi-agent systems. We specifically study this in the context of common interest games in which groups of up to k agents can coordinate their actions in maximizing a common objective function. The equilibria that emerge in these systems are the k-strong Nash equilibria of the common interest game; studying the properties of these states provide relevant insights into the efficacy of inter-agent collaboration. Our contributions come threefold: 1) provide bounds on how well k-strong Nash equilibria approximate the optimal system welfare, formalized by the k-strong price of anarchy, 2) prove the run-time and transient performance of collaborative agent-based dynamics, and 3) introduce techniques of redesigning objectives for groups of agents which improve system performance. We study these three facets generally as well as in the context of resource allocation problems, in which we provide tractable linear programs that give tight bounds on the k-strong price of anarchy.

2.1 Introduction

Large-scale systems such as transportation services [40], robotic fleets [41], supply chains [42], or cloud computing services [43] can be challenging to design effective control schemes for due to their many components and vast scale. The two prevailing paradigms to design control schemes are centralized control [44, 45, 46], which guides behavior across the entire system and distributed control [47, 48, 49], which allows local components to guide their own behavior. Each of these approaches possesses respective pros and cons: centralization allows for more direct manipulation of system behavior at the cost of greater communication and computation requirements, while decentralization reduces the communication and computation requirements but cannot always attain the desired system behavior. With advancements of small computers and communication technologies [50, 3, 51, 52], we are enabled to design new paradigms that *exist between centralized and distributed control.*

Specifically, we study the efficacy of learning in multi-agent systems when individual system components (or agents) can partially communicate and thus coordinate their behavior. Many engineering domains are on the precipice of enabling these collaborative paradigms; for example, autonomous vehicle platoons with connected cruise control [53], unmanned aerial surveillance vehicles with range-limited communication [54], and cloud computing networks with emerging distributed learning techniques [55]. In

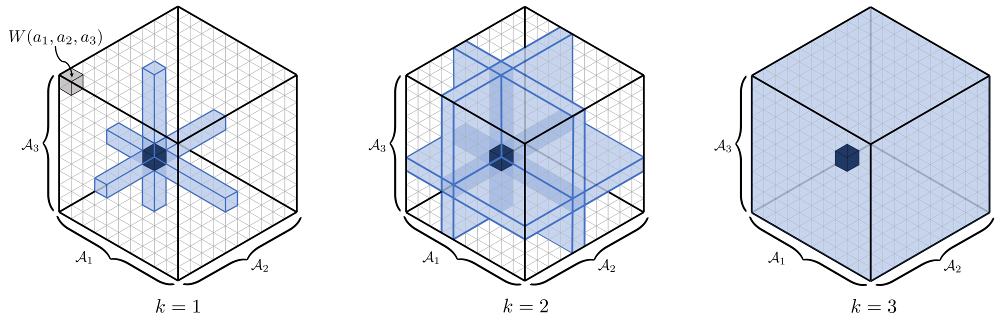

Figure 2.1: Illustration of the k-strong Nash equilibrium local optimality guarantee for a 3 agent matrix game where $k \in \{1, 2, 3\}$. In each case, if the dark cube is a k-strong Nash equilibrium, then it is an optimal joint action over the highlighted region. As k (the size of collaborative groups) increases, the local optimality is strengthened by holding over all k-lateral deviations.

each of these settings, inter-agent communication and collaboration offer the opportunity to improve the performance attainable by the system as a whole; however, implementing these frameworks incurs costs that are both monetary–in the form of the additional technology required–and computational–in the form of more complex decision-making algorithms. In this section, we provide tools to help better understand the *benefits and costs associated with collaborative communication* in multi-agent systems.

We model a multi-agent system as a common interest game where some (but not all) groups of agents can collaborate in selecting their actions to maximize the system welfare. We particularly focus on the case where a collaborative action takes the form of a group best response, i.e., a group of agents updating their actions in response to the remaining players' actions. As the size and number of these collaborative groups increase, a coordinated group decision has a larger impact on system behavior. To range the level of collaboration between the fully distributed setting (where no agents can collaborate) and the fully centralized setting (where all agents can collaborate collectively), we consider the cases where groups of up to k agents can collaborate. In these collaborative environments, a stable state of the system is that of the k-strong Nash equilibrium [56]. Researchers have studied the existence [57] and computation [58] of strong Nash equilibria in settings including congestion games [59], lexicographical games [60], and Markov games [61]. This section applies these concepts to multi-agent systems. To understand the possible benefits of collaboration to system performance, we quantify how well k-strong Nash equilibria approximate the optimal welfare, termed the k-strong price of anarchy [62, 63]. To understand the possible cost of collaboration, we analyze the running time and transient performance of agent-based dynamics, which converge to k-strong Nash equilibria.

Distributed learning in games has been a widely studied area in controls [64], but the ability to reach equilibrium with coalitional best responses has not yet been studied; similarly, the efficient computation of a k-strong Nash equilibrium has largely been studied from a centralized perspective [58]. Quantifying the k-strong price of anarchy has been studied in network formation games [62, 63] and load balancing games [65, 66, 67], as well as more general utility maximizing games [68, 69]. In many of these, the bounds are either not tight or only hold in the limit of large numbers of players.

***Organization* -** In this section, we provide tools to understand the benefits and costs of collaborative communication by studying the qualities of k-strong Nash equilibria. In Section 2.3.1, we consider the case where groups of agents are designed to maximize the system welfare and introduce the notion of (λ, μ)-k-coalitionally smooth games (a generalization of smooth games [70] and coalitionally smooth games [68]), and provide bounds on the k-strong price of anarchy. Then, in Section 2.3.2, we focus on the well studied setting of distributed resource allocation problems [36, 71, 72, 73, 74, 4], and provide tight bounds on the k-strong price of anarchy via the solution of a tractable linear program. Fig. 2.3 plots these bounds and demonstrates how increased collaboration improves efficiency guarantees in several classes of resource allocation problems. In Section 2.4, we consider the effects of group decision-making on agent-based dynamics; specifically, we show the added run-time complexity of coalitional round-robin dynamics and provide transient performance guarantees of asynchronous best response dynamics.

We support our findings with numerical examples, which highlight that collaborative agent-based dynamics provide better performance but require more evaluations of the system's welfare. In Section 2.5, we consider that the system operator may be able to design the agents' objective separately from the system welfare; we provide a generalized technique for bounding the k-strong price of anarchy in this setting. In Section 2.5.2, we again focus on the setting of resource allocation and provide two linear programs to lower and upper bound the attainable k-strong price of anarchy guarantee via utility design.

2.2 Preliminaries

Throughout, we will denote $[n] = \{1, \ldots, n\}$. We will regularly use the binomial coefficient $\binom{n}{k} = \frac{n!}{(n-k)!k!}$ in constructing optimization problems; we define this value as 0 when $n < k$ for ease of notation.

2.2.1 Collaborative Decision Making

Consider a finite set of agents $N = \{1, \ldots, n\}$. Each agent $i \in N$ selects an action a_i from a finite action set $\mathcal{A}_i$. When each agent selects an action, we will denote their joint action by the tuple $a = (a_1, \ldots, a_n) \in \mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$. Let $G = (N, \mathcal{A})$ be a tuple encoding the components of the agent environment. The system's performance is dictated by the agents' actions; as such, for each joint-action a we assign a system welfare $W(a)$ where $W : \mathcal{A} \to \mathbb{R}_{\geq 0}$ is the system designer's objective function. With this, we let the tuple (G, W) denote a *multi-agent system* (often referred to as a system) which defines the primitives of the system designer's problem of designing an effective control algorithm.

The system designer would like to configure the agents to reach a joint action that maximizes the system welfare, i.e.,

$$a^{\text{opt}} \in \arg\max_{a \in \mathcal{A}} W(a). \tag{2.1}$$

Though this system state is ideal, it may be difficult to attain as 1) solving for the optimal allocation can be combinatorial and in some cases (including those from Section 2.3.2) NP-hard [36], and 2) it requires a centralized authority to control all agents, which may be practically or logistically difficult. To resolve this, we will consider that agents make decisions in a decentralized manner.

Fully distributing the decision-making involves designing each agent to update their action locally and has been widely studied and developed to guarantee reasonable system behavior [47]; however, fully distributing decision-making may often become unnecessary as emerging communication technologies enable *collaborative inter-agent decision-making*[50]. To implement one such collaborative system architecture, a system operator must make two decisions: 1) which group of agents can collaborate on their decisions (possibly subject to some operational constraints), and 2) how the agents should collaborate on their decisions. A natural choice for the latter is a group best response. Let $\Gamma \subseteq N$ be a group of agents endowed with the ability to collaboratively select a *group action* $a_\Gamma \in \mathcal{A}_\Gamma = \prod_{i \in \Gamma} \mathcal{A}_i$, which they select by maximizing the system welfare over their group action-set,

$$a_\Gamma \in \arg\max_{a'_\Gamma \in \mathcal{A}_\Gamma} W(a'_\Gamma, a_{-\Gamma}), \tag{2.2}$$

where $a_{-\Gamma}$ denotes the actions of the players $i \in N \setminus \Gamma$. If there are multiple elements in the argmax, the group breaks them at random unless they can remain with their current action.

Intuitively, a group best responding and collaboratively maximizing the system welfare should lead to direct improvements to system performance; however, one can consider other group decision-making rules as well. In particular, in Section 2.5, we will consider that the system designer can design the agents' objective separately from the system objective as a means to further shape system behavior. In either case, one would imagine that the greater the collaborative structure, the greater the impact on emergent behavior.

For the system operator's decision over which groups should collaborate, let $\mathcal{C} \subseteq 2^N$ denote the *collaboration set*, or the set of groups of agents ($\Gamma \in \mathcal{C}$) able to collaborate their decisions. These collaborations can overlap–where agents can partake in multiple, disparate collaborations–and vary in size. For example, if agents send signals through a communication network [75], we will have $\mathcal{C} = \{(i,j) \in N^2 \mid (i,j) \in E\}$ where E are the edges in a communication graph. If agents are allowed to communicate with each other one at a time and make pairwise decisions [76], then $\mathcal{C} = \{(i,j) \in N^2\}$. If agents can only communicate with others within a local proximity [77], then $\mathcal{C} = \{\Gamma \subseteq N \mid \rho(i,j) \leq d \ \forall i,j \in \Gamma\}$ where ρ measures the distance between two agents and d is a maximum communication range. Once the system operator decides on the collaborative structure and the group decision-making protocol, the agents' decision-making process forms a collaborative multi-agent system, denoted by the tuple $(G, W, \mathcal{C})$.

As we vary the number and size of collaborative sets, we can consider control paradigms somewhere between centralized (i.e., $\{N\} \in \mathcal{C}$) and fully distributed (i.e., $\mathcal{C} = \bigcup_{i \in N} \{i\}$). This section seeks to understand the efficacy of different levels of communication/ collaboration. To more effectively quantify this, we consider a specific type of collaboration set in which we can range between the centralized and distributed extremes.

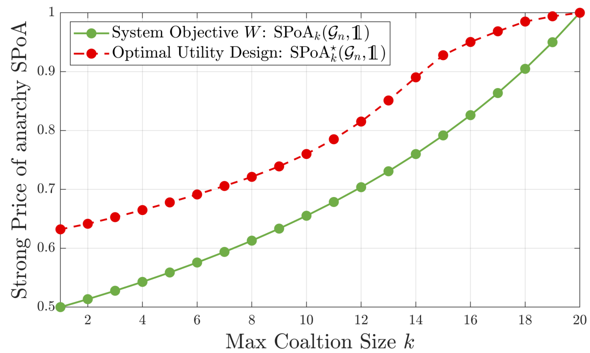

Figure 2.2: Strong Price of Anarchy in resource covering games with $n = 20$ players and coalitions up to size k (horizontal axis). As the size of groups that are allowed to collaborate grows, so too does the approximation ratio (i.e., strong price of anarchy) of a k-strong Nash equilibrium. The efficiency of an equilibrium can be further improved by designing the utility functions agents are set to maximize. The solid green line is the k-strong price of anarchy when agents maximize the system objective (generated by Theorem 2.3.1). The dashed red line is an upper bound on the k-strong price of anarchy while using an optimal utility design (generated by Proposition 2.5.2).

2.2.2 k-Strong Nash Equilibria

We consider the collaboration sets that contain groups of agents up to size k. Let $\mathcal{C}_k = \{\Gamma \subseteq N \mid |\Gamma| = k\}$ denote the subsets of exactly k agents and $\mathcal{C}_{[k]} = \bigcup_{\zeta \in [k]} \mathcal{C}_\zeta$ be the subsets that contain at most k agents. When $k = 1$, we recover the fully distributed setting, and when $k = n$, we recover the fully centralized setting. As we vary k between 1 and n, we sweep through different levels of communication and collaboration.

In the game-theoretic approach to multi-agent systems, a Nash equilibrium is a joint action where no agent can unilaterally deviate their action to improve the system welfare [56]. We generalize this concept to the setting of collaborative decision-making by considering a *k-strong Nash equilibrium* as a joint action where no group of k agents can deviate their group's actions to improve the welfare.

Definition 1. *A joint-action $a^{k\mathrm{SNE}} \in \mathcal{A}$ is a k-strong Nash equilibrium for the common-*

interest game $(G, W, \mathcal{C}_{[k]})$ *if*

$$W(a^{k\text{SNE}}) \geq W(a'_\Gamma, a^{k\text{SNE}}_{-\Gamma}),\ \forall a'_\Gamma \in \mathcal{A}_\Gamma,\ \Gamma \in \mathcal{C}_{[k]}. \tag{2.3}$$

Let $k\text{SNE}(G, W) \subseteq \mathcal{A}$ denote the set of all k-strong Nash equilibria. This definition differs slightly from the literature, where k-strong Nash equilibria are defined by no group of agents deviating to a new group action that is Pareto-optimal for the group (i.e., no agent receives a lower payoff with respect to their individual utility function) [56]; when the agents respond to a common interest objective, the definitions are equivalent. Additionally, in general games, k-strong Nash equilibria need not exist; however, that is not the case in our setting due to the common-interest structure we impose on agent decision-making.

Proposition 2.2.1. *In a system* (G, W) *with collaboration set* $\mathcal{C}_{[k]}$ *for any* $k \in [n]$*, a* k*-strong Nash equilibrium exists.*

The proof appears in the Appendix A.1.

The main focus of this section is understanding how equilibrium performance changes with the level of collaborative communication. Notice that (2.3) serves as a local optimality guarantee in the neighborhood of k-lateral deviations. Fig. 2.1 depicts this for a three-player matrix game; when $k = 1$, a 1-strong Nash equilibrium is optimal over the unilateral deviations, when $k = 2$ a 2-strong Nash equilibrium is optimal over the bilateral deviations, and when $k = 3 = n$, the 3-strong Nash equilibrium is optimal over the whole joint-action space. From this, we observe that the local optimality guarantee is strengthened as we increase the level of collaboration k (i.e., $k'\text{SNE} \subseteq k\text{SNE}$ for $k' > k$).

To quantify the effect of varying k on equilibrium performance, we consider the ratio of worst-case equilibrium welfare and the optimal attainable welfare, termed the

k-strong price of anarchy.

$$\mathrm{SPoA}_k(G,W) = \frac{\min_{a^{k\mathrm{SNE}} \in k\mathrm{SNE}(G,W)} W(a^{k\mathrm{SNE}})}{\max_{a^{\mathrm{opt}} \in \mathcal{A}} W(a^{\mathrm{opt}})} \in [0,1], \tag{2.4}$$

where we let 0/0 be defined as 1 to ignore the degenerate case when no welfare is attainable. In the multi-agent system (G, W) with communication structure $\mathcal{C}_{[k]}$, every k-strong Nash equilibrium approximates the optimal solution at least as well as $\mathrm{SPoA}_k(G, W)$. Accordingly, we will use the k-strong price of anarchy to understand the efficiency associated with collaborative decision-making. For example, in Fig. 2.2, we depict the k-strong price of anarchy in resource covering games [36] for $1 \leq k \leq n$, illustrating the performance guarantees attainable between centralized and distributed control paradigms.

2.2.3 Summary of Contributions

This section studies the benefits and costs of increased collaborative communication within multi-agent systems. Our contributions come threefold:

1) In Section 2.3, we provide tools to quantify the k-strong price of anarchy when agents optimize the system objective. We introduce (λ, μ)-k-coalitionally smooth games and provide a k-strong price of anarchy guarantee using the parameters λ and μ. We then focus on the class of resource allocation games, where in Proposition 2.3.2, we show that these parameters can be found via the solution to a tractable linear program. In Theorem 2.3.1, we show that combining the constraints of each of the k linear programs gives a tight bound. Figure 2.3 depicts the k-strong price of anarchy for several classes of resource allocation games.

2) In Section 2.4, we study collaborative dynamics that reach these equilibria. In Section 2.4.1, we introduce the coalitional round-robin dynamics and show that an equi-

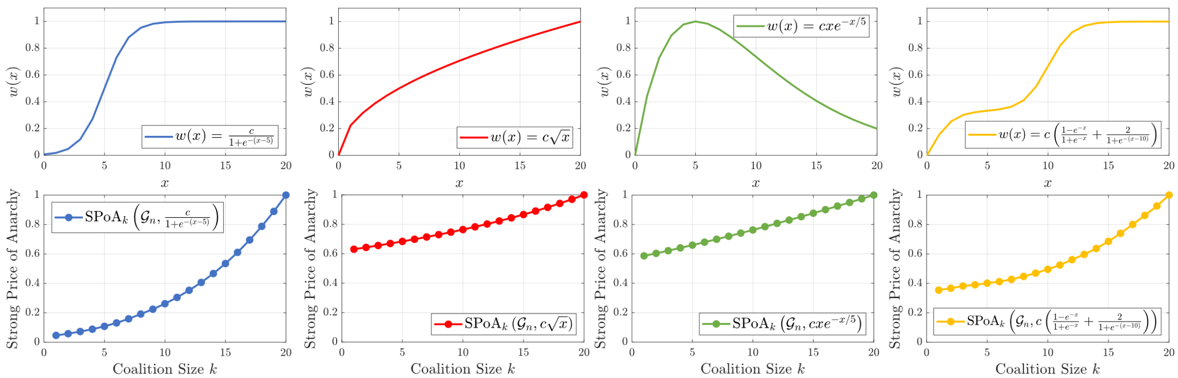

Figure 2.3: Tight k-strong price of anarchy bounds for resource allocation games with various welfare functions. We illustrate four settings of local welfare function (top, left to right), and for each, we use Theorem 2.3.1 to generate tight bounds on the k-strong price of anarchy for all $1 \leq k \leq n$. The bottom figures show these bounds and illustrate how increased inter-agent collaboration increases our efficiency guarantees on equilibrium system welfare.

librium is reached in a finite number of best responses and that the number of welfare comparisons grows with a small-base exponential of k. In Section 2.4.2, we introduce the asynchronous coalitional best response dynamics, which we show converge almost surely. Further, if the game is (λ, μ)-k-coalitionally smooth , then we provide a bound on the transient performance (or the cumulative welfare along the dynamics). We support these findings with a numerical study in Section 2.4.3.

3) In Section 2.5, we consider how to improve the design of a group's decision-making process. By providing the agents with a new, designed objective function, the system designer may alter the set of equilibria and ideally increase the k-strong price of anarchy. In Section 2.5.1, we generalize the notion of coalitional smoothness to the setting where the agents' objective differs from the system welfare, and in Theorem 2.5.1, we show how we can construct an optimal utility rule. Fig. 2.5 shows the k-strong price of anarchy under the optimal utility design for resource allocation games, demonstrating the added benefit of designing how groups of agents make decisions.

2.3 Quantifying k-Strong Price of Anarchy

2.3.1 Coalitionally Smooth Games

We first consider the efficiency of k-strong Nash equilibria for general multi-agent systems. This efficiency–quantified by the k-strong price of anarchy–is conditioned on the system welfare W and the agent decision-making environment G. In Definition 2, we provide a condition on a system (G, W) that will be useful in bounding the k-strong price of anarchy.

Definition 2. *A system (G, W) is (λ, μ)-k-coalitionally smooth , where $\lambda, \mu \in \mathbb{R}^k_{\geq 0}$, if for all $a, a' \in \mathcal{A}$*

$$\frac{1}{\binom{n}{\zeta}} \sum_{\Gamma \in \mathcal{C}_\zeta} W(a'_\Gamma, a_{-\Gamma}) \geq \lambda_\zeta W(a') - \mu_\zeta W(a), \ \forall \zeta \in [k]. \tag{2.5}$$

In (2.5), we provide a constraint on the welfare function stating that the average effect of a group of size ζ deviating their action from a to a' is lower bounded by a linear combination of the welfare of a and a'. The term smooth is in reference to the welfare function's change over the joint-action space being bounded by (2.5). Additionally, Definition 2 extends the classic notion of smooth games [70] and coalitional smoothness for strong equilibria [68] to the setting of k-coalitions in common interest games.

In effect, every system (G, W) is smooth with $\lambda_\zeta = \mu_\zeta = 0$ for all $\zeta \in [k]$, but some parameters (λ, μ) are more useful than others. In Proposition 2.3.1, we show that the parameters λ and μ from Definition 2 can be used to lower bound the k-strong price of anarchy.

Proposition 2.3.1. *A system (G, W) that is (λ, μ)-k-coalitionally smooth has k-strong* *price of anarchy satisfying*

$$\mathrm{SPoA}_k(G,W) \geq \frac{\lambda_\zeta}{1+\mu_\zeta}, \ \forall \zeta \in [k]. \tag{2.6}$$

Proof. Let $a^{k\mathrm{SNE}} \in \mathcal{A}$ denote a k-strong Nash equilibrium in the system (G, W) (i.e., satisfying Definition 1), and let $a^{\mathrm{opt}} \in \arg\max_{a\in\mathcal{A}} W(a)$ denote an optimal joint action. For any $\zeta \in [k]$, we have

$$W(a^{k\mathrm{SNE}}) = \frac{1}{\binom{n}{\zeta}} \sum_{\Gamma \in \mathcal{C}_\zeta} W(a^{k\mathrm{SNE}}) \tag{2.7a}$$

$$\geq \frac{1}{\binom{n}{\zeta}} \sum_{\Gamma \in \mathcal{C}_\zeta} W(a^{\mathrm{opt}}_\Gamma, a^{k\mathrm{SNE}}_{-\Gamma}) \tag{2.7b}$$

$$\geq \lambda_\zeta W(a^{\mathrm{opt}}) - \mu_\zeta W(a^{k\mathrm{SNE}}). \tag{2.7c}$$

Where (2.7a) holds from $|\mathcal{C}_\zeta| = \binom{n}{\zeta}$, (2.7b) holds from Definition 1, and (2.7c) holds from Definition 2. Rearranging, we get $W(a^{k\mathrm{SNE}})/W(a^{\mathrm{opt}}) \geq \lambda_\zeta/(1+\mu_\zeta)$. □

(2.6) provides k lower bounds on the k-strong price of anarchy, i.e., a k-strong Nash equilibrium approximates the system optimal at least as well as $\max_{\zeta\in[k]}\{\lambda_\zeta/(1+\mu_\zeta)\}$. Often, the best lower bound is provided by $\zeta = k$; however, this is not true in general. As such, we must consider each of the constraints in (2.5) to derive the best bounds.

The efficiency bounds of this form are valuable for several reasons, including: 1) they can be used to provide insights on the transient guarantees of various multi-agent dynamics (see Section 2.4, 2) they easily generalize to broader equilibrium concepts (subject of future work), and 3) if parameters (λ, μ) can be shown to satisfy (2.5) for a set of systems $\mathcal{S}$, then each system $(G, W) \in \mathcal{S}$ inherits the efficiency guarantee of (2.6). This last point is particularly pertinent, as system models may be subject to noise, mischaracterizations, or changes over time. If the efficiency guarantee holds across many similar systems, then the guarantees are essentially robust to these issues.

In the Section 2.3.2, we will provide methods to find coalitional smoothness parameters for classes of resource allocation games via tractable linear programs.

2.3.2 Resource Allocation Games

In this section, we consider the well studied class of resource allocation games [72, 36, 73, 74, 78]. Consider a set of resources or tasks $\mathcal{R} = \{1, \ldots, R\}$, to which agents are assigned, i.e., agent $i \in N$ selects a subset of these resources as its action $a_i \subseteq \mathcal{R}$ from a constrained set of subsets $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$. Each resource $r \in \mathcal{R}$ has a value $v_r \geq 0$; the welfare contributed by a resource is $v_r w(|a|_r)$, where $w : \{0, \ldots, n\} \to \mathbb{R}_{\geq 0}$ captures the added benefit of having multiple agents assigned to the same resources and $|a|_r$ is the number of agents assigned to r in allocation a. Assume that $w(0) = 0$ as no welfare is contributed by resources assigned to zero agents and further that $w(y) > 0$ for all $y > 0$. The system welfare is thus

$$W(a) = \sum_{r \in \mathcal{R}} v_r w(|a|_r). \tag{2.8}$$

For ease of notation, we will refer to the system welfare by the local welfare rule w, noting in the agent-environment G, it generates a welfare function W via (2.8).

As discussed in 2.3.1, we wish to find efficiency bounds that hold over a class of resource allocation problems. Let $G = (\mathcal{R}, N, \mathcal{A}, \{v_r\}_{r \in \mathcal{R}})$ denote a resource allocation problem, and let $\mathcal{G}_n$ denote the set of all such resource allocation problems with at most n agents. In Proposition 2.3.2, we propose a tractable linear program whose solution provides parameters (λ, μ) which satisfy Definition 2 for every system $(G, w) \in \mathcal{G}_n \times \{w\}$. From Proposition 2.3.1, this also provides a lower bound on the k-strong price of anarchy for the class of resource allocation problems with local welfare w.

Proposition 2.3.2. *Each resource allocation problem* $(G, w) \in \mathcal{G}_n \times \{w\}$ *is* (λ, μ)-k-

coalitionally smooth with $\lambda_\zeta = 1/\nu_\zeta^\star$ *and* $\mu_\zeta = \rho_\zeta^\star/\nu_\zeta^\star - 1$*, where* $(\rho_\zeta^\star, \nu_\zeta^\star)$ *is a solution to the linear program* (Pζ)*:*

$$
\begin{aligned}
&(\rho_\zeta^\star, \nu_\zeta^\star) \in \underset{\rho \geq \nu \geq 0}{\arg\min} \quad \rho \\
&\text{s.t.} \;\; 0 \geq w(o{+}x) - \rho w(e{+}x) + \\
&\nu \left(\binom{n}{\zeta} w(e{+}x) - \sum_{\substack{0 \leq \alpha \leq e \\ 0 \leq \beta \leq o \\ \alpha+\beta \leq \zeta}} \binom{e}{\alpha}\binom{o}{\beta}\binom{n-e-o}{\zeta-\alpha-\beta} w(e{+}x{+}\beta{-}\alpha) \right) \\
&\qquad\qquad\qquad\qquad \forall (e,x,o) \in \mathcal{I}
\end{aligned} \tag{Pζ}
$$

The constraints are parameterized by the triples $\mathcal{I} := \{(e,x,o) \in \mathbb{N}^3_{\geq 0} \mid 1 \leq e + x + o \leq n\}$. With the possibility of collaboration, an equilibrium becomes more difficult to characterize than in a fully distributed setting. We circumvent this by introducing a parameterization which allows us to generalize the $\mathcal{O}\left(\sum_{\zeta=1}^{k} \binom{n}{\zeta} m^\zeta\right)$ comparisons of (2.3) (where $m := \max_{i \in N} |\mathcal{A}_i|$) into $\mathcal{O}(n^3)$ linear inequalities. Further, satisfying these inequalities provides parameters $(\lambda_\zeta, \mu_\zeta)$ that satisfy Definition 2, leading to (Pζ) as a search for such parameters with the best k-strong price of anarchy guarantee.

Proof of Proposition 2.3.2: The proof largely relies on introducing a parameterization that lets us treat (2.5) as a set of linear constraints. Consider a resource allocation game $(G, w) \in \mathcal{G}_n \times \{w\}$ and any two actions $a, a' \in \mathcal{A}$. To each resource $r \in \mathcal{R}$, we assign a label (e_r, x_r, o_r), where

$$
\begin{aligned}
e_r &= |\{i \in N \mid r \in a_i \setminus a_i'\}| \\
x_r &= |\{i \in N \mid r \in a_i \cap a_i'\}| \\
o_r &= |\{i \in N \mid r \in a_i' \setminus a_i\}|.
\end{aligned}
$$

This is to say, e_r denotes the number of agents utilizing resource r in joint action a but not a', o_r is the number that uses resource r in joint action a' but not a, and x_r is the number that uses r in both a and a'. In the set of games $\mathcal{G}_n$, let $\mathcal{I} = \{(e, x, o) \in \mathbb{N}^3_{\geq 0} \mid 1 \leq e + x + o \leq n\}$ denote the set of possible labels, and $\theta(e, x, o) := \sum_{r \in \mathcal{R}_{(e,x,o)}} v_r$, where $\mathcal{R}_{(e,x,o)} = \{r \in \mathcal{R} \mid e_r = e, x_r = x, o_r = o\}$ denotes the set of resources with label (e, x, o). The parameter $\theta \in \mathbb{R}^{|\mathcal{I}|}_{\geq 0}$ is a vector with elements for each label.

We will now express the terms in (2.5) using this parameterization. Because $W(a) = \sum_{r \in \mathcal{R}} v_r w(|a|_r)$ depends only on the number of agents utilizing a resource, we can represent $|a|_r = e_r + x_r$ and write the system welfare as

$$\begin{aligned} W(a) &= \sum_{r \in \mathcal{R}} v_r w(e_r + x_r) \\ &= \sum_{(e,x,o) \in \mathcal{I}} \left(\sum_{r \in \mathcal{R}_{e,x,o}} v_r \right) w(e + x) \\ &= \sum_{e,x,o} \theta(e, x, o) w(e + x). \end{aligned}$$

When not stated, the sum over (e, x, o) is implied to be for each label in $\mathcal{I}$. Similar steps can be followed to show $W(a') = \sum_{e,x,o} \theta(e, x, o) w(o + x)$.

Finally, the term $\sum_{\Gamma \in \mathcal{C}_\zeta} W(a'_\Gamma, a_{-\Gamma})$ can similarly be transcribed by this parameterization:

$$
\begin{aligned}
&\sum_{\Gamma \in \mathcal{C}_\zeta} W(a'_\Gamma, a_{-\Gamma}) \\
&= \sum_{\Gamma \in \mathcal{C}_\zeta} \sum_{e,x,o} \sum_{r \in \mathcal{R}_{(e,x,o)}} v_r w(|a'_\Gamma, a_{-\Gamma}|_r) \\
&= \sum_{e,x,o} \sum_{r \in \mathcal{R}_{(e,x,o)}} v_r \sum_{\Gamma \in \mathcal{C}_\zeta} w(|a'_\Gamma, a_{-\Gamma}|_r) \\
&= \sum_{e,x,o} \sum_{r \in \mathcal{R}_{(e,x,o)}} v_r \sum_{\substack{0 \le \alpha \le e \\ 0 \le \beta \le o \\ \alpha+\beta \le \zeta}} \binom{e}{\alpha} \binom{o}{\beta} \binom{n-e-o}{\zeta-\alpha-\beta} w(e+x+\beta-\alpha) \\
&= \sum_{e,x,o} \theta(e,x,o) \sum_{\substack{0 \le \alpha \le e \\ 0 \le \beta \le o \\ \alpha+\beta \le \zeta}} \binom{e}{\alpha} \binom{o}{\beta} \binom{n-e-o}{\zeta-\alpha-\beta} w(e+x+\beta-\alpha)
\end{aligned}
$$

where the set of coalitions $\mathcal{C}_\zeta$ was partitioned according to the action profile of the agents in each coalition. We let α denote the number of agents in Γ that utilize resource r only in joint action a and β the number of agents in Γ that utilize r only in joint action a'. By simple counting arguments, there are exactly $\binom{e}{\alpha}\binom{o}{\beta}\binom{n-e-o}{\zeta-\alpha-\beta}$ coalitions grouped with the same α and β. This decomposition is possible as the number of agents utilizing resource r after a group Γ deviates is precisely $e + x + \beta - \alpha$.

The smoothness constraint (2.5) is satisfied only if

$$
\frac{1}{\binom{n}{\zeta}} \sum_{e,x,o} \theta(e,x,o) \sum_{\substack{0 \le \alpha \le e \\ 0 \le \beta \le o \\ \alpha+\beta \le \zeta}} \binom{e}{\alpha} \binom{o}{\beta} \binom{n-e-o}{\zeta-\alpha-\beta} w(e+x+\beta-\alpha)
$$
$$
\ge \lambda_\zeta \sum_{e,x,o} \theta(e,x,o) w(o+x) - \mu_\zeta \sum_{e,x,o} \theta(e,x,o) w(e+x).
$$

As $\theta(e,x,o) \geq 0$ for all $(e,x,o) \in \mathcal{I}$, it is sufficient to satisfy

$$\frac{1}{\binom{n}{\zeta}} \sum_{\substack{0\leq\alpha\leq e\\ 0\leq\beta\leq o\\ \alpha+\beta\leq\zeta}} \binom{e}{\alpha}\binom{o}{\beta}\binom{n-e-o}{\zeta-\alpha-\beta} w(e+x+\beta-\alpha)$$
$$\geq \lambda_\zeta w(o+x) - \mu_\zeta w(e+x), \quad \forall (e,x,o) \in \mathcal{I}. \quad (2.9)$$

Observe that (2.9) is independent of a, a', and G. As such, this set of constraints serves as a sufficient condition that any $G \in \mathcal{G}_n$ satisfies (2.5) for all respective $a, a' \in \mathcal{A}$.

To find parameters λ_ζ and μ_ζ that provide the best k-strong price of anarchy guarantee, we formulate the following optimization problem:

$$\max_{\lambda_\zeta,\mu_\zeta\geq 0} \quad \frac{\lambda_\zeta}{1+\mu_\zeta} \qquad \text{(P1}\zeta\text{)}$$
$$\text{s.t.} \quad (2.9)$$

We restrict λ_ζ to be non-negative, though this constraint is not active except in degenerate cases. Finally, we transform (P1ζ) by substituting new decision variables $\rho = (1+\mu_\zeta)/\lambda_\zeta$ and $\nu = 1/\left(\binom{n}{\zeta}\lambda_\zeta\right) \geq 0$. The new objective becomes $1/\rho$. Note that the constraint $(e,x,o) = (1,0,0)$ implies $\rho \geq 0$; we can thus invert the objective and change the minimization to a maximization, giving (Pζ). □

The smoothness parameters found via Proposition 2.3.2 can be used with Proposition 2.3.1 to generate lower bounds on the k-strong price of anarchy. However, these bounds need not be tight, i.e., there may be no system in the class $\mathcal{G}_n \times \{w\}$ that attains this inefficiency, and better bounds may be possible. To study what efficiency we can guarantee across a class of resource allocation problems, we define the k-strong price of

$$
\begin{aligned}
P^\star = \min_{\rho,\{\nu_\zeta \geq 0\}_{\zeta\in[k]}} \quad & \rho \\
\text{s.t.} \quad & 0 \geq w(o+x) - \rho w(e+x) \\
& + \sum_{\zeta\in[k]} \nu_\zeta \left(\binom{n}{\zeta} w(e+x) - \sum_{\substack{0\leq\alpha\leq e \\ 0\leq\beta\leq o \\ \alpha+\beta\leq\zeta}} \binom{e}{\alpha}\binom{o}{\beta}\binom{n-e-o}{\zeta-\alpha-\beta} w(e+x+\beta-\alpha) \right) \\
& \qquad \forall (e,x,o)\in\mathcal{I}
\end{aligned} \tag{P[k]}
$$

anarchy bound for $(\mathcal{G}_n, w)$ as

$$\text{SPoA}_k(\mathcal{G}_n, w) = \min_{G\in\mathcal{G}_n} \text{SPoA}_k(G, w). \tag{2.10}$$

This performance ratio is parameterized by our choice of welfare function w and the size of collaborative coalitions k. In Theorem 2.3.1, we provide a linear program whose value provides an exact value of $\text{SPoA}_k(\mathcal{G}_n, w)$. We do this by showing that the constraints of the k linear programs in Proposition 2.3.2 can be combined to give an exact quantification of the k-strong price of anarchy bound.

Theorem 2.3.1. *For the class of resource allocation problems $\mathcal{G}_n$ with welfare function w, when groups maximize the common interest welfare, then*

$$\text{SPoA}_k(\mathcal{G}_n, w) = 1/P^\star(n, w, k), \tag{2.11}$$

where $P^\star(n,w,k)$ is the solution to (P[k]).

The proof appears in Appendix A.1.

In Fig. 2.3, we consider four welfare functions and plot the tight bounds on the k-strong price of anarchy for $1 \leq k \leq n$. As expected, we observe that increased communication improves efficiency guarantees; the amount of this increase is useful in determining the benefits of inter-agent communication/collaboration. However, this collaboration comes at a cost; in Section 2.4, we will study the complexity of distributed dynamics reaching k-strong Nash equilibria.

2.4 Coalitional Dynamics

Section 2.3 provided several tools for quantifying the efficiency guarantees of k-strong Nash equilibria. In this section, we will study the qualities of group-based dynamics that reach these equilibria. Particularly, we will discuss the convergence rate and transient performance when agents follow the Coalitional round-robin and Asynchronous Best Response, respectively. We will denote a^t as the joint action occurring at time $t \in \mathbb{N}$ and $\Gamma^t \subseteq N$ as the group of agents updating their action at time t.

2.4.1 Round Robin

We first consider the k-coalitional round robin agent dynamics, in which each group of k agents updates their actions sequentially, following a set order $\sigma \in \Sigma_{\binom{n}{k}}$, where $\sigma(z)$ for $z \in \{1, \ldots, \binom{n}{k}\}$ is the index of a group $\Gamma \in \mathcal{C}_{[k]}$. We will call a *round* one pass through σ in which each group updates their action. At their turn, the group Γ^t selects their best response to the current action, i.e., $a_{\Gamma^t}^{t+1} \in \arg\max_{a_{\Gamma^t} \in \mathcal{A}_{\Gamma^t}} W(a_{\Gamma^t}, a_{-\Gamma^t})$, where ties are broken uniformly at random unless $a_{\Gamma^t}^t \in \arg\max_{a_{\Gamma^t} \in \mathcal{A}_{\Gamma^t}} W(a_{\Gamma^t}, a_{-\Gamma^t})$, in which case the group selects their current action $a_{\Gamma^t}^{t+1} = a_{\Gamma^t}^t$. The dynamics are more formally described in Algorithm 1.

These dynamics are synchronous (in that agents must follow a set order) but provide an understanding of how groups of agents can make decisions in a localized manner, and we can analyze the equilibrium hitting time. In the fully distributed setting ($k = 1$), it has been shown that these dynamics reach a Nash equilibrium in finite time and require

Algorithm 1 k-Round-Robin Dynamics

```
procedure kRoundRobin(W, 𝒜, N, σ, a)
    ā ← NULL
    while ā ≠ a do
        ā ← a
        for z ∈ {1, ..., (n choose k)} do
            Γ ← 𝒞_k(σ(z))                              ▷ Get group
            for a⁺_Γ ∈ 𝒜_Γ \ a_Γ do                    ▷ Group deviations
                if W(a⁺_Γ, a_{−Γ}) > W(a) then
                    a ← (a⁺_Γ, a_{−Γ})
```

$\mathcal{O}(n^m)$ welfare evaluations [79]. In Proposition 2.4.1, we find that in the coalitional settings, we maintain the finite convergence time and incur a small base exponential gain in the number of welfare comparisons required. Recent work has shown that the examples that realize these worst-case hitting times are fragile and that equilibria can be computed in polynomial-time under smoothed running-time analysis [80]. As a first step, we consider the worst-case run time, but the authors believe that similar findings on the added complexity of group decision-making will hold under smoothed running-time analysis, though this is the subject of ongoing work.

Proposition 2.4.1. *The k-Coalitional-Round-Robin dynamics converge in finite time and requires* $\mathcal{O}\left(m^n\left(\frac{1}{1-1/m}\right)^k\right)$ *welfare evaluations.*

Proof. First, we verify that the output of Algorithm 1 is a k-strong Nash equilibrium, then we consider how long it takes Algorithm 1 to converge. Algorithm 1 terminates after a round in which no group $\Gamma \in \mathcal{C}_k$ can select a new action in which the welfare increases, i.e., $W(a) \geq W(a_\Gamma, a_{-\Gamma})$ for all $a_\Gamma \in \mathcal{A}_\Gamma$ and $\Gamma \in \mathcal{C}_k$ where a is the output of Algorithm 1. A deviation for a any subgroup $\Gamma' \in \mathcal{C}_{[k]}$ is subsumed by the joint action $(a_{\Gamma'}, a_{\Gamma\setminus\Gamma'}) \in \mathcal{A}_\Gamma$. As such, a state a terminates Algorithm 1 if and only if it satisfies (2.3) and is a k-strong Nash equilibrium.

Without loss of generality, we assume each agent possesses m actions; for each agent,

i that has fewer actions, assign $m - |\mathcal{A}_i|$ dummy actions with minimum welfare. In one round of the k-Round-Robin dynamics, each group of agents is given the opportunity to deviate their action. First, we note that no group Γ will respond to the same complimentary group action $a_{-\Gamma}$ in two consecutive rounds unless a is a k-strong Nash equilibrium. If the group Γ rejects a group action a_Γ in response to $a_{-\Gamma}$, the joint action $(a_\Gamma, a_{-\Gamma})$ is eliminated from consideration as an output of Algorithm 1. Accounting for overlaps between the groups, in any round that does not start in a k-strong Nash equilibrium, at least $y = \sum_{\zeta=1}^{k} \binom{n}{\zeta}(m-1)^{\zeta}$, joint actions are eliminated as possible outputs of Algorithm 1. As there are m^n joint actions in total, there can be at most $r \leq \lfloor \frac{m^n}{y} \rfloor + 1$ rounds that do not start in a k-strong Nash equilibrium; this proves the finite convergence time. In each round, there are exactly $\binom{n}{k} m^k$ welfare checks; thus, the total number of welfare checks is no more than $(\frac{m^n}{y} + 1)\binom{n}{k} m^k$. Removing lower order terms from y gives the stated bound. □

From Proposition 2.4.1, we observe two things: 1) the coalitional dynamics do not require drastically more welfare evaluations than the fully distributed round robin, but 2) the convergence rate is slow regardless of k. In light of this, we turn our focus to understanding the transient performance of collaborative decision-making dynamics. Further, in many settings, it is desirable to allow agents or groups to update their actions asynchronously. In Section 2.4.2, we will consider both of these factors in the asynchronous best response dynamics.

2.4.2 Asynchronous Best-Response Dynamics

Motivated by settings where agents (or groups of agents) perform action revisions asynchronously or on their own time scales, we consider a dynamical system where the next group of agents to update is random.

We define the Asynchronous k-Coalitional Best-Response Dynamics as follows: let $t \geq 0$ denote the number of agent (or group) updates that have yet occurred[1]. The updating group Γ^t is selected at random, such that the size of the group ζ is picked with probability $p_\zeta = \mathbb{E}[|\Gamma^t| = \zeta]$ and the specific agents in the group are drawn uniformly at random. Once formed, the updating group Γ^t chooses their best response in the same manner as the coalitional round robin described in Section 2.4.1.

From their distributed decision-making and asynchronicity, these dynamics capture the behavior of real-time multi-agent systems components. In Theorem 2.4.1, we show these dynamics converge almost surely to a k-strong Nash equilibrium, and further, if the system is (λ, μ)-k-coalitionally smooth , we provide a bound on the cumulative welfare relative to the optimal.

Theorem 2.4.1. *The Asynchronous k-Coalitional Best-Response Dynamics converge almost surely to the set of k-strong Nash equilibrium. Further, if (G, W) is a (λ, μ)-k-coalitionally smooth system, then after $T \geq 1$ update steps, the cumulative expected welfare satisfies*

$$\mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T} W(a^t)\right] \geq \frac{T-1}{2T} \frac{\sum_{\zeta=1}^{k} p_\zeta \lambda_\zeta}{1 + \sum_{\zeta=1}^{k} p_\zeta \mu_\zeta} W(a^{\mathrm{opt}}), \tag{2.12}$$

where p_ζ is the probability a group of size ζ best responds.

Interestingly, the bound on the average transient welfare depends on how frequently groups of different sizes are sampled to perform their best response. When the agents are designed to more regularly collaborate in larger groups, the transient guarantee will often be better.

[1]Counting time steps in terms of the number of updates subsumes cases where agents (or groups) update with respect to individual and independent random clocks. The rate of each clock is analogous to the selection probability for different groups.

Proof. First, we show that the Asynchronous k-Coalitional Best-Response Dynamics converges in general. A group Γ revises their action only to one of strictly higher payoff if one exists. Consider the resulting Markov chain $\mathcal{M}$ with states $\mathcal{A}$. Any state $a \in \mathcal{A} \setminus k\text{SNE}$ has an outgoing edge with positive probability as there exists some group $\Gamma \in \mathcal{C}_{[k]}$ that is selected with probability $p_{|\Gamma|}/|\mathcal{C}_{|\Gamma|}| > 0$ which would revise their action. Any state $a \in k\text{SNE}$ has no outgoing edges with positive probability as no group $\Gamma \in \mathcal{C}_{[k]}$ can revise their action to strictly increase the welfare. Finally, there are no cycles (excluding self-loops) in $\mathcal{M}$, as every outgoing edge is directed from a joint action of lower welfare to one of strictly higher welfare. As such, the set kSNE is absorbing and $\mathbb{P}[\lim_{t\to\infty} a^t \in k\text{SNE}] = 1$.

Now, consider that the system (G, W) is (λ, μ)-k-coalitionally smooth . As the selection of the updating group is random, the welfare at time $t+1$ is a random variable, even when conditioned on a^t; the expectation of the succeeding welfare can be written

$$\begin{aligned}
\mathbb{E}[W(a^{t+1})\,|a^t = a] &= \sum_{\zeta=1}^{k} p_\zeta \sum_{\Gamma \in \mathcal{C}_\zeta} \frac{1}{\binom{n}{\zeta}} W(a_\Gamma^+, a_{-\Gamma}) \\
&\geq \sum_{\zeta=1}^{k} p_\zeta \sum_{\Gamma \in \mathcal{C}_\zeta} \frac{1}{\binom{n}{\zeta}} W(a_\Gamma^{\text{opt}}, a_{-\Gamma}) \\
&\geq \sum_{\zeta=1}^{k} p_\zeta \left(\lambda_\zeta W(a^{\text{opt}}) - \mu_\zeta W(a)\right) \\
&= \left(\sum_{\zeta=1}^{k} p_\zeta \lambda_\zeta\right) W(a^{\text{opt}}) - \left(\sum_{\zeta=1}^{k} p_\zeta \mu_\zeta\right) W(a),
\end{aligned}$$

where $a_\Gamma^+ \in \arg\max_{a_\Gamma \in \mathcal{A}_\Gamma} W(a_\Gamma, a_{-\Gamma})$ is the update state for the group Γ following the dynamics; the welfare for each possible updated joint action is the same, so determining which group action is selected is irrelevant. As a_Γ^+ is a best response, the welfare is no better for selecting a different action, namely a_Γ^{opt}. The final inequality holds from (2.5).

Taking the expectation of $\mathbb{E}[W(a^{t+1}) \mid a^t = a]$ over a^t gives

$$\mathbb{E}\left[W(a^{t+1})\right] \geq \left(\sum_{\zeta=1}^{k} p_\zeta \lambda_\zeta\right) W(a^{\text{opt}}) - \left(\sum_{\zeta=1}^{k} p_\zeta \mu_\zeta\right) \mathbb{E}\left[W(a^t)\right].$$

Rearranging terms shows

$$\begin{aligned}\mathbb{E}\left[W(a^{t+1})\right] - \frac{\sum_{\zeta=1}^{k} p_\zeta \lambda_\zeta}{1 + \sum_{\zeta=1}^{k} p_\zeta \mu_\zeta} W(a^{\text{opt}}) \\ \geq \left(\sum_{\zeta=1}^{k} p_\zeta \mu_\zeta\right) \left(\frac{\sum_{\zeta=1}^{k} p_\zeta \lambda_\zeta}{1 + \sum_{\zeta=1}^{k} p_\zeta \mu_\zeta} W(a^{\text{opt}}) - \mathbb{E}\left[W(a^t)\right]\right).\end{aligned}$$

Observe that either $\mathbb{E}\left[W(a^t)\right] \geq \frac{\sum_{\zeta=1}^{k} p_\zeta \lambda_\zeta}{1+\sum_{\zeta=1}^{k} p_\zeta \mu_\zeta} W(a^{\text{opt}})$ or $\mathbb{E}\left[W(a^{t+1})\right] - \frac{\sum_{\zeta=1}^{k} p_\zeta \lambda_\zeta}{1+\sum_{\zeta=1}^{k} p_\zeta \mu_\zeta} W(a^{\text{opt}}) \geq 0$. Accordingly, in expectation, every other update must satisfy the bound, giving the average cumulative welfare bound in (2.12). □

Theorem 2.4.1 shows that the transient efficiency changes with the frequency with which different group sizes perform best responses. To attain the best transient guarantee, we can select p carefully.

Corollary 1. *If a system (G, w) is a resource allocation problem in $\mathcal{G}_n \times \{w\}$, then selecting $p_\zeta \propto \frac{\nu_\zeta^\star}{\sum_{\psi \in [k]} \binom{n}{\psi} \nu_\psi^\star}$ for all $\zeta \in [k]$ gives*

$$\mathbb{E}\left[\frac{1}{T} \sum_{t=1}^{T} W(a^t)\right] \geq \frac{T-1}{2T} \text{SPoA}_k(\mathcal{G}_n, w) W(a^{\text{opt}}).$$

The proof is omitted as it is straightforward by rearranging terms in the constraints of (D).

Together, Theorem 2.4.1 and Corollary 1 provide insight into the transient performance of non-deterministic multi-agent dynamical systems with collaborative communication. Future work will study the traits of non-best-response dynamics, namely regret-based decision-making.

2.4.3 Numerical Example

We support the findings of Section 2.4.2 by numerical example. We randomly generate resource allocation problems and simulate the coalitional asynchronous best response dynamics when groups of size $k \in \{1, 2, 3, 4, 5\}$ update.

The resource allocation problems are generated by creating 100 resources with values independently drawn uniformly at random on $[0, 1]$. Each of the 25 agents is endowed with between 1 and 10 actions (also sampled uniformly at random). For each action of each player, each resource is included in that particular action with probability 0.25. This defines a tuple G. We use the local welfare function $w(x) = xe^{-x/5}$ to capture some added benefit from having multiple agents use the same resource and eventual diminishing returns and increased cost from over congestion.

We select a random initial condition and run the asynchronous best response dynamics with $p_k = 1$ for one value $k \in \{1, 2, 3, 4, 5\}$ (i.e., only groups of exactly size k are sampled, but the simulation is repeated for $1 \leq k \leq 5$. We ran this simulation 100 times.

At left in Fig. 2.4, we plot the average welfare across the simulations over the number of group action revisions. We observe that the larger coalitions provide superior transient and long-run performance. However, a single group action revision requires more computation for larger coalitions. At right in Fig. 2.4, for each coalition size $k \in \{1, \ldots, 5\}$, we show a scatter plot of the number of cumulative welfare evaluations and the attained system welfare, along with a trend line fit to the data within two standard deviations of the average number of welfare evaluations. Here, we observe that for lower values of welfare, the smaller coalitions can attain similar welfare with fewer welfare evaluations

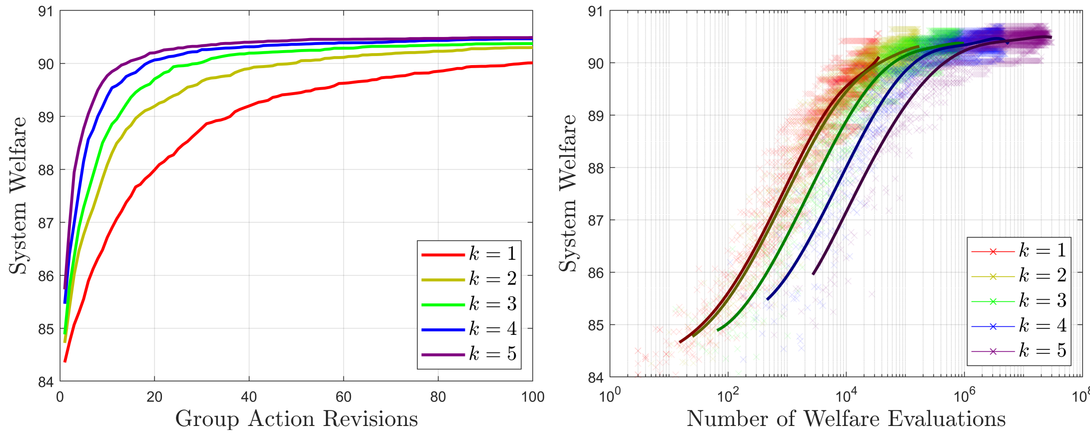

Figure 2.4: Numerical example of the coalitional asynchronous best response dynamics. At left, the system welfare is plotted over the number of group action revisions, and at right it is plotted over the number of welfare evaluations. From this data, we can observe that group revisions offer superior system transient and long-term performance but require more welfare evaluations to compute group actions.

but that the larger coalitions reach higher welfare much more regularly.

These conclusions help to identify the trade-off in designing systems with collaborative communication: better performance is attainable at the cost of greater computation.

2.5 Utility Design

Up until this point, agents and groups of agents have been set to optimize the system welfare W over their respective individual or group actions. Though this is a reasonable approach, the system designer may seek to further improve system performance by designing how a group of agents makes a decision. Consider that groups of agents instead maximize the objective function $U : \mathcal{A} \to \mathbb{R}_{\geq 0}$ (henceforth referred to as the *utility function*), i.e.,

$$a_\Gamma \in \arg\max_{a'_\Gamma \in \mathcal{A}_\Gamma} U(a'_\Gamma, a_{-\Gamma}), \tag{2.13}$$

where ties are still broken at random unless the current group action is in the argmax. By designing the utility function U, the system operator can alter how groups of agents make

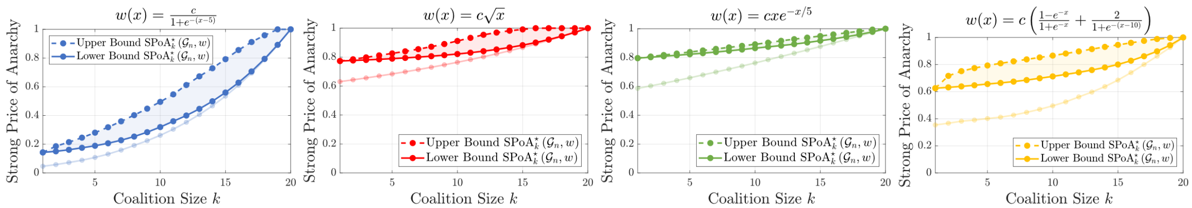

Figure 2.5: Bounds on the k-strong Price of Anarchy using the optimal utility function in the class of resource allocation games with welfare function w. Upper bound on $\mathrm{SPoA}_k^\star(\mathcal{G}_n, w)$ generated by Proposition 2.5.2 and lower bound and utility rule that attains it generated by Theorem 2.5.1. Compared with the k-strong price of anarchy when agents optimize the system welfare (lighter line), we demonstrate the possible and guaranteed gain in equilibrium performance attainable by designing group decision-making for collaborative multi-agent systems.

decisions and, ideally, improve the performance of the system. A multi-agent system is now captured by the tuple (G, W, U), where the previous results are the special case when $U = W$.

By redefining the objective functions groups of agents seek to maximize, we additionally alter the equilibria that emerge from collaborative decision-making. We alter the definition of k-strong Nash equilibria to hold with respect to the utility function, i.e.,

$$U(a^{k\mathrm{SNE}}) \geq U(a'_\Gamma, a^{k\mathrm{SNE}}_{-\Gamma}), \ \forall a'_\Gamma \in \mathcal{A}_\Gamma, \ \Gamma \in \mathcal{C}_{[k]}. \tag{2.14}$$

Let $k\mathrm{SNE}(G, U)$ denote the set of k-strong Nash equilibria when agents optimize the objective U. The new set of equilibria implies the equilibrium performance guarantee may also change. As such, we redefine the k-strong price of anarchy as the approximation of the optimal welfare provided the system equilibria under objective function U,

$$\mathrm{SPoA}_k(G, W, U) = \frac{\min_{a^{k\mathrm{SNE}} \in k\mathrm{SNE}(G,U)} W(a^{k\mathrm{SNE}})}{\max_{a^{\mathrm{opt}} \in \mathcal{A}} W(a^{\mathrm{opt}})}. \tag{2.15}$$

With this new design opportunity, we identify two goals in understanding the new attainable performance of collaborative decision-making: 1) quantifying the performance of a prescribed utility function, and 2) finding a utility function that provides the greatest k-strong price of anarchy guarantees. We address these two points in general in Section 2.5.1 and more thoroughly within resource allocation problems in Section 2.5.2.

2.5.1 Generalized Coalitionally Smooth Games

In this section, we consider the general setting and particularly focus on quantifying the k-strong price of anarchy of a system (G, W, U). As in Section 2.3.1, we introduce a notion of smooth systems now generalized to the setting where the agent objective U differs from the system objective W.

Definition 3. *A system (G, W, U) is (λ, μ)-k-generalized-coalitionally smooth , where $\lambda, \mu \in \mathbb{R}^k_{\geq 0}$, if for all $a, a' \in \mathcal{A}$*

$$\frac{1}{\binom{n}{\zeta}} \sum_{\Gamma \in \mathcal{C}_\zeta} U(a'_\Gamma, a_{-\Gamma}) - U(a) + W(a) \geq \lambda_\zeta W(a') - \mu_\zeta W(a), \ \forall \zeta \in [k]. \quad (2.16)$$

Like (2.5), (2.16) provides a bound on average deviation effect of a group of size ζ but on the utility function instead of the welfare. In Proposition 2.5.1, we show that (λ, μ)-k-generalized-coalitionally smooth system permits a bound on the k-strong price of anarchy.

Proposition 2.5.1. *A system (G, W, U) that is (λ, μ)-k-generalized-coalitionally smooth has k-strong price of anarchy satisfying*

$$\text{SPoA}_k(G, W, U) \geq \frac{\lambda_\zeta}{1 + \mu_\zeta}, \ \forall \zeta \in [k]. \quad (2.17)$$

Proof. Let $a^{k\mathrm{SNE}} \in \mathcal{A}$ denote a k-strong Nash equilibrium when agents follow objective function U, and let $a^{\mathrm{opt}} \in \arg\max_{a \in \mathcal{A}} W(a)$ denote an optimal joint action. For any $\zeta \in [k]$, we have

$$W(a^{k\mathrm{SNE}}) \geq \frac{1}{\binom{n}{\zeta}} \sum_{\Gamma \in \mathcal{C}_\zeta} U(a^{\mathrm{opt}}_\Gamma, a^{k\mathrm{SNE}}_{-\Gamma}) - U(a^{k\mathrm{SNE}}) + W(a^{k\mathrm{SNE}}) \tag{2.18a}$$

$$\geq \lambda_\zeta W(a^{\mathrm{opt}}) - \mu_\zeta W(a^{k\mathrm{SNE}}). \tag{2.18b}$$

Where (2.18a) holds from $\frac{1}{\binom{n}{\zeta}} \sum_{\Gamma \in \mathcal{C}_\zeta} U(a^{\mathrm{opt}}_\Gamma, a^{k\mathrm{SNE}}_{-\Gamma}) - U(a^{k\mathrm{SNE}}) \geq 0$ by $a^{k\mathrm{SNE}}$ being a k-strong Nash equilibrium and (2.7c) holds from Definition 3. Rearranging, we get $W(a^{k\mathrm{SNE}})/W(a^{\mathrm{opt}}) \geq \lambda_\zeta/(1+\mu_\zeta)$. □

Beyond quantifying the k-strong price of anarchy for a system (G, W, U), one may wish to find the utility function which provides the best efficiency guarantee, i.e.,

$$U \in \underset{U': \mathcal{A} \to \mathbb{R}_{\geq 0}}{\arg\max} \, \mathrm{SPoA}_k(G, W, U').$$

For a specific problem (G, W), it is possible to design a utility function which guarantees that a system optimal a^{opt} is a unique equilibrium and provides $\mathrm{SPoA}_k(G, W, U) = 1$ (e.g., $U(a) = \sum_{i \in N} \mathbb{1}_{[]} a_i = a_i^{\mathrm{opt}}]$). However, this would require knowing the optimal allocations a priori, which poses several problems, including: 1) computing an optimal allocation can be intractable, and 2) system parameters may be subject to modeling errors, noise, or changes over time, causing the optimal allocations to change. As such, we will consider the design of *utility rules*, which provide a set of instructions to construct a utility function across a class of systems and eliminate the computational burden of solving for a new utility function for each system while maintaining improved performance guarantees. Luckily, the approach in Proposition 2.5.1 is amenable to generating performance guarantees across a class of systems, and in Section 2.5.2, we will investigate optimal utility rules more thoroughly in resource allocation problems.

2.5.2 Resource Allocation Games

In this section, we consider the k-strong price of anarchy in classes of resource allocation problems when the agents' objective is derived from a utility rule $u \in \mathbb{R}^{n+1}_{\geq 0}$. In an agent environment $G = (N, \mathcal{A}, \mathcal{R}, \{v_r\}_{r \in \mathcal{R}})$, the utility rule u can be applied to derive the utility function

$$U(a) = \sum_{r \in \mathcal{R}} v_r u(|a|_r).$$

To normalize the utility function, we set $u(0) = 0$. We ultimately consider the performance of a utility rule u across all agent environments $G \in \mathcal{G}_n$ with welfare function w. We slightly abuse notation to refer to a system by the tuple (G, w, u). To quantify this performance, we generalize the k-strong price of anarchy bound defined in Section 2.4.1 to hold for cases where groups of agents optimize the utility function.

$$\text{SPoA}_k(\mathcal{G}_n, w, u) = \min_{G \in \mathcal{G}_n} \text{SPoA}_k(G, w, u). \tag{2.19}$$

The performance ratio is parameterized by the pair (w, u); as such, we will discuss the effectiveness of a utility rule u with respect to a given welfare function w.

Taking the utility rule approach completely eliminates the computational cost of deriving a utility function for each problem instance; now we seek to understand the capabilities of this approach in two ways: 1) in Theorem 2.5.1 we demonstrate how we can construct utility rules with good performance guarantees, and 2) in Proposition 2.5.1 we provide an upper bound on the best attainable performance a utility rule can provide. In Corollary 2, we provide a formal condition on when the constructed utility rule is optimal.

Theorem 2.5.1. *Any resource allocation problem* $(G, W) \in \mathcal{G}_n \times \{w\}$ *with the utility rule* $\widetilde{u}_\zeta$ *is* $(1, \widetilde{\rho}_\zeta - 1)$*-*$k$*-generalized-coallitionally smooth, where* $\widetilde{u}_\zeta$ *and* $\widetilde{\rho}_\zeta$ *are solutions to the linear program,*

$$
\begin{aligned}
(\widetilde{\rho}_\zeta, \widetilde{u}_\zeta) \in \; & \underset{\rho \geq 0, u \in \mathbb{R}^{n+1}_{\geq 0}}{\arg\min} \quad \rho \\
\text{s.t.} \quad & 0 \geq w(o{+}x) - \rho w(e{+}x) + \\
& \left(\binom{n}{\zeta} u(e{+}x) - \sum_{\substack{0 \leq \alpha \leq e \\ 0 \leq \beta \leq o \\ \alpha + \beta \leq \zeta}} \binom{e}{\alpha} \binom{o}{\beta} \binom{n{-}e{-}o}{\zeta{-}\alpha{-}\beta} u(e{+}x{+}\beta{-}\alpha) \right) \\
& \qquad \forall (e, x, o) \in \mathcal{I}.
\end{aligned} \tag{Qζ}
$$

Proof. Consider the parameterization described in the proof of Proposition 2.3.2, where for any two actions $a, a' \in \mathcal{A}$, we can rewrite $W(a) = \sum_{e,x,o} \theta(e, x, o) w(e{+}x)$ and $W(a') = \sum_{e,x,o} \theta(e, x, o) w(o{+}x)$. Now, we can additionally rewrite $U(a) = \sum_{e,x,o} \theta(e, x, o) u(e{+}x)$ and

$$
\begin{aligned}
& \sum_{\Gamma \in \mathcal{C}_\zeta} W(a'_\Gamma, a_{-\Gamma}) \\
& \qquad = \sum_{e,x,o} \theta(e, x, o) \sum_{\substack{0 \leq \alpha \leq e \\ 0 \leq \beta \leq o \\ \alpha + \beta \leq \zeta}} \binom{e}{\alpha} \binom{o}{\beta} \binom{n{-}e{-}o}{\zeta{-}\alpha{-}\beta} w(e{+}x{+}\beta{-}\alpha).
\end{aligned}
$$

We can now write out (2.16), the (λ, μ)-k-generalized-coalitionally smooth constraint,

as

$$\sum_{e,x,o}\theta(e,x,o)\left(\frac{1}{\binom{n}{\zeta}}\sum_{\substack{0\le\alpha\le e\\0\le\beta\le o\\\alpha+\beta\le\zeta}}\binom{e}{\alpha}\binom{o}{\beta}\binom{n-e-o}{\zeta-\alpha-\beta}u(e+x+\beta-\alpha) - u(e+x)\right) \geq \sum_{e,x,o}\theta(e,x,o)\left(\lambda_\zeta w(o+x)-(\mu_\zeta+1)w(e+x)\right).$$

As before, we can observe that this constraint is sufficiently satisfied when

$$\frac{1}{\binom{n}{\zeta}}\sum_{\substack{0\le\alpha\le e\\0\le\beta\le o\\\alpha+\beta\le\zeta}}\binom{e}{\alpha}\binom{o}{\beta}\binom{n-e-o}{\zeta-\alpha-\beta}u(e+x+\beta-\alpha) - u(e+x) \geq \lambda_\zeta w(o+x)-(\mu_\zeta+1)w(e+x), \quad \forall(e,x,o)\in\mathcal{I}. \quad (2.20)$$

The task of finding smoothness parameters that give the best price of anarchy guarantee becomes the same problem as (P1ζ) but now with constraint set (2.20). By substituting the decision variables $\rho = (1+\mu_\zeta)/\lambda_\zeta$ and $\nu = 1/\left(\binom{n}{\zeta}\lambda_\zeta\right) \geq 0$, we attain the new constraint set

$$0 \geq w(o+x) - \rho w(e+x) + \nu\left(\binom{n}{\zeta}u(e+x) - \sum_{\substack{0\le\alpha\le e\\0\le\beta\le o\\\alpha+\beta\le\zeta}}\binom{e}{\alpha}\binom{o}{\beta}\binom{n-e-o}{\zeta-\alpha-\beta}u(e+x+\beta-\alpha)\right) \quad \forall(e,x,o)\in\mathcal{I}. \quad (2.21)$$

The new objective[2] becomes $1/\rho$.

[2]As an aside, the transformed program up to this point can be used to evaluate the performance of a specified utility rule.

Finally, we let $u \in \mathbb{R}^n_{\geq 0}$ become a decision variable in the program. Observe that every occurrence of u is multiplied by ν and every occurrence of ν multiplies u. As such, we can define the new decision variable $u' = \nu u$ and retrieve the linear program (Qζ). □

The utility rule $\hat{u}_\zeta$ that (Qζ) provides gives us some guarantee on attainable performance from designing group decision-making in collaborative systems. However, it is not yet clear if these are the best possible utility rules. To understand what the best possible performance is of a collaborative system, we define the optimal k-strong price of anarchy as

$$\mathrm{SPoA}^\star_k(\mathcal{G}_n, w) = \sup_{u:[n]\to R_{\geq 0}} \mathrm{SPoA}_k(\mathcal{G}_n, w, u). \tag{2.22}$$

This upper bound informs us of what efficiency is possible to hope for out of a collaborative system. In Proposition 2.5.2, we bound this quantity.

Proposition 2.5.2. *For the class of resource allocation problems $\mathcal{G}_n \times \{w\}$, when agents maximize the optimal utility design objective $u^\star$,*

$$\mathrm{SPoA}^\star_k(\mathcal{G}_n, w) \leq 1/Q^\star(n, w, k), \tag{2.23}$$

where $Q^\star(n, w, k)$ is value of the linear program

$$\begin{aligned} Q^\star(n,w,k) = &\min_{\rho\geq 0, \{u_\zeta \in \mathbb{R}^{n+1}_{\geq 0}\}_{\zeta\in[k]}} \quad \rho \\ \text{s.t. } \; & 0 \geq w(o{+}x) - \rho w(e{+}x) + \\ & \sum_{\zeta\in[k]} \left(\binom{n}{\zeta} u_\zeta(e{+}x) - \sum_{\substack{0\leq\alpha\leq e\\ 0\leq\beta\leq o\\ \alpha+\beta\leq\zeta}} \binom{e}{\alpha}\binom{o}{\beta}\binom{n-e-o}{\zeta-\alpha-\beta} u_\zeta(e{+}x{+}\beta{-}\alpha) \right) \\ & \qquad \forall (e,x,o) \in \mathcal{I}. \end{aligned} \tag{Q[k]}$$

The proof appears in Appendix A.1.

Note that Theorem 2.5.1 provides a utility rule with associated performance guarantee which lower bounds $\mathrm{SPoA}_k^\star(\mathcal{G}_n, w)$, and Proposition 2.5.2 provides an upper bound. In Corollary 2, we note that when these two bounds match, we have a tight bound on $\mathrm{SPoA}_k^\star(\mathcal{G}_n, w)$ as well as an optimal utility rule.

Corollary 2. *For the class of resource allocation problems $\mathcal{G}_n \times \{w\}$, if the value of* (Q$\zeta$) *satisfies $\rho_\zeta^\star = Q^\star(n, w, k)$, then $\mathrm{SPoA}_k^\star(\mathcal{G}_n, w) = 1/Q^\star(n, w, k)$ is a tight bound and a solution $\widetilde{u}_\zeta$ to* (Qζ) *is an optimal utility rule.*

Proof. This follows immediately from $1/\rho_\zeta^\star = \frac{\lambda_\zeta}{1+\mu_\zeta}$ being a lower bound on $\mathrm{SPoA}_k^\star(\mathcal{G}_n, w)$ and the reciprocal of the value of (Q[k]), $1/Q^\star$ being an upper bound. When the two match, the bound must be tight. □

The two bounds coinciding is not guaranteed but does occur at the extremes ($k = 1$ and $k = n$); further, the gap between the two bounds is often small, and the lower bound attained by the utility rule constructed in Theorem 2.5.1 often demonstrates a significant improvement over the setting where agents simply optimize the system objective. Consider the four welfare functions from Fig. 2.3 again; for each, we find that the utility rule computed using Theorem 2.5.1 and the upper bound on $\mathrm{SPoA}_k^\star(\mathcal{G}_n, w)$ using Proposition 2.5.2. In Fig. 2.5 we plot these lower and upper bounds on $\mathrm{SPoA}_k^\star(\mathcal{G}_{20}, w)$ for each utility function and for each value of $1 \le k \le n$; these values are juxtaposed with the k-strong price of anarchy when agents optimize the system objective w to demonstrate the possible gain in performance from designing the agents' objective in collaborative systems.

2.6 Conclusion

In this section, we provided a variety of tools for evaluating the benefits and costs of collaborative communication in multi-agent systems. A collaborative multi-agent system was modeled by a common interest game where groups of players collaboratively perform their best responses simultaneously. We specifically considered the k-strong Nash equilibrium as a relevant equilibrium concept to gain insights into system behavior between the fully centralized and fully distributed settings. We introduced the notion of (λ, μ)-k-coalitionally smooth systems and derived bounds on how well the k-strong Nash equilibrium approximates the optimum in such systems. Further analysis studied the running time of collaborative multi-agent decision dynamics and their transient performance, as well as the possible performance gains from designing agents' objectives separately from the system objective. Finally, we underwent a more thorough study in the class of resource allocation games, in which we provided tractable linear programs whose solutions give tight bounds on the k-strong price of anarchy in resource allocation games. Future work will study less extensive communication paradigms and dynamical systems that emerge when agents learn together.

Chapter 3

The Cost of Informed Decision Making in Multi-Agent Maximum Coverage Problems

The emergent behavior of a distributed system is conditioned by the information available to the local decision-makers. Therefore, one may expect that providing decision-makers with more information will improve system performance; in this work, we find that this is not necessarily the case. In multi-agent maximum coverage problems, we find that even when agents' objectives are aligned with the global welfare, informing agents about the realization of the resource's random values can reduce equilibrium performance by a factor of 1/2. This affirms an important aspect of designing distributed systems: information need be shared carefully. We further this understanding by providing lower and upper bounds on the ratio of system welfare when information is (fully or partially) revealed and when it is not, termed the value-of-informing. We then identify a trade-off that emerges when optimizing the performance of the best-case and worst-case equilibrium.

3.1 Introduction

In large-scale systems, the prospect of distributing decision-making to local entities is becoming increasingly enticing as a method to reduce complexity while maintaining some level of performance. This can take the form of swarm control for robotic fleets [81], autonomous driving decisions in mobility services [40], local task assignment decisions [82], and many more. Taking a distributed approach entails assigning each agent a decision-making algorithm, such as maximizing an assigned local objective function [83], then analyzing the system's equilibria [84]. As agents need not possess full knowledge of the overall system; the local decisions (and, ultimately, global behavior) are dependent on the information communicated with and between the agents [85, 4] In this work, we address how available information affects system performance.

We focus on maximum coverage problems: a class of models in which each agent selects a set of resources from a ground set, with the objective of maximizing the value of covered resources. To solve this in a distributed fashion, each agent is given a utility function to evaluate what set of resources to cover; this forms a game played by the agents with their resource selection as their action and the evaluation of their assigned utility function as their payoff. Existing work has focused on how to design these utility rules and how well the resulting equilibria of the emergent game approximate the optimal welfare [72, 36, 86]. In this work, we generalize this model to consider the case where agents have uncertainty about the resources' values. In this setting, we ask how revealing information to local decision-makers affects equilibrium performance. Interestingly, we find that *revealing truthful information about the system state can worsen system performance*. While this phenomenon has been observed before in social systems [87], here we find that similar conclusions hold even when the local decision-makers' objectives are *aligned* with the global welfare.

To study this, we consider a *Bayesian persuasion* framework, in which a well-informed system operator can strategically reveal information to agents using messages (or signals) which contain partial information [88]. In this work, we study how this information revealing affects equilibrium welfare. To this end, we introduce a new performance metric termed the *value-of-informing*, which measures the ratio between the equilibrium welfare under an information-revealing policy and when no information is revealed. This measures the gain or loss in welfare from revealing information.

The framework of Bayesian persuasion has gained traction in the areas of economics, operations research, and engineering, but typically for settings concerned with the behavior and beliefs of human users (e.g., traffic routing apps [89, 13, 90], pricing/investing decisions [91], hybrid work policies [92], etc.). Results are typically restricted to a binary classification on whether revealing full information helps or not [93, 88], or methods to compute optimal information revealing policies in limited settings [94, 95, 90], often with no guarantee on the magnitude of improvement. Here, we adapt the ideas of information provisioning to the setting of engineered systems, where designed decision-making components can improve their estimate of the system state by receiving relevant messages (e.g., a fleet of surveillance drones receiving live map updates).

Revealing information to local decision-makers can obviously improve the welfare of emergent system equilibria; however, in this work, we find that this is not always the case. In fact, system performance may degrade by a factor of $1/2$ when revealing truthful information to local decision-makers. The possible loss in system welfare from informing decision-makers comes from two sources (1) the multiplicity of equilibria and (2) the local objective assigned to each decision-maker. We study the aforementioned value-of-informing when considering the best-case and worst-case equilibria for different local objectives. Ultimately, we highlight a trade-off between the possible loss from revealing information to best-case and worst-case equilibrium guarantees when agents' local objectives are designed.

3.2 Problem Formulation

3.2.1 Maximum Coverage Problem

Maximum coverage problems have been used to model resource allocation, sensor coverage, job scheduling, and more [96]. Consider the multi-agent maximum coverage problem, in which $\mathcal{R} = \{1, \dots, R\}$ is a finite set of resources. For each resource $r \in \mathcal{R}$, let $v_r \geq 0$ be the value of that resource; further, let $v \in \mathbb{R}_{\geq 0}^{|\mathcal{R}|}$ be the vector containing each resource value. Let $N = \{1, \dots, n\}$ be a set of agents, where each agent $i \in N$ can be assigned to cover a subset of resources $a_i \subseteq \mathcal{R}$. The set of allowable assignments for each agent is defined by $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$. When each agent is assigned, an allocation of agents is denoted $a = (a_1, \dots, a_n) \in \mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$. Let (G, v) define a maximum coverage problem where $G = (N, \mathcal{R}, \mathcal{A})$.

In an allocation a, the system welfare is equal to the total value of resources covered by at least one agent, i.e., $W(a; v) = \sum_{r \in \cup_{i \in N} a_i} v_r$. However, finding an optimal allocation $a^{\text{opt}} \in \arg\max_{a \in \mathcal{A}} W(a; v)$ is NP-hard [36]. It is for this reason, we consider a distributed solution technique to approximate this optimal solution.

3.2.2 Distributed Decision Making

Let each agent $i \in N$ possess a local objective function

$$U_i(a; v) = \sum_{r \in a_i} v_r f(|a|_r),$$

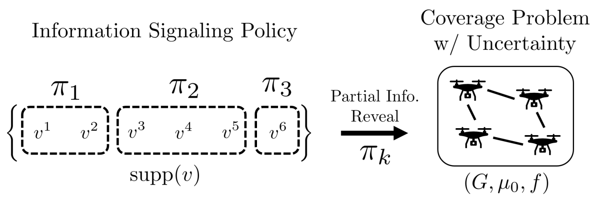

Figure 3.1: Depiction of information signaling in maximum coverage problems. On the left is the support of a random state variable v. At right is a maximum coverage problem, to which we have assigned each agent a local objective. The agents in the coverage problem possess the prior distribution of the unknown state variable and receive some partial information π_k about the realization. The manner in which information is revealed will alter how agents evaluate their objectives and change the emergent behavior.

which depends on their own action and the actions of every other agent. This local objective, or *utility function*, is parameterized by a local utility rule $f : \mathbb{N} \to \mathbb{R}$ that takes as argument $|a|_r$, the number of agents covering resource r in allocation a. The system operator can adopt a utility rule f without exact knowledge of the problem instance if needed.

When agents sequentially and repeatedly update their assignment to one that currently maximizes their utility function, these best-response dynamics have fixed points that are the Nash equilibria of the underlying game. The convergence of best response dynamics in potential games to pure Nash equilibria are addressed in [36, 84]. Nash equilibrium allocations can be defined by

$$U_i(a^{\text{Ne}}; v) \geq U_i(a_i', a_{-i}^{\text{Ne}}; v) \ \forall a_i' \in \mathcal{A}_i, \ i \in N, \tag{3.1}$$

where a_{-i} denotes the allocation of all agents but player i. Let $\text{NE}(G, v, f)$ denote the set of states satisfying (3.1). These states represent the possible solutions of the distributed dynamics, and we will consider their welfare as an approximation of the maximum coverage problem.

3.2.3 Uncertainty and Information Signaling

In this work, we consider how uncertainty and information can affect the efficacy of distributed decision-making. We consider this uncertainty in the form of randomness for the resources values. Let $v \in \mathbb{R}^{|\mathcal{R}|}_{\geq 0}$ (the vector containing the value of each resource $r \in \mathcal{R}$) be a discrete random variable with prior distribution μ_0. A realization of v determines each resource's value $[v_1, \ldots, v_{|\mathcal{R}|}] = v$, i.e., the resource values may be correlated. Let the support of v be $\mathcal{V} := \mathrm{supp}(v)$.

First, we consider the case where agents are *uninformed* about the system state, i.e., they know the prior distribution μ_0 but not the exact realization of v. In this setting, the agents optimize their expected utility, which we will denote $\overline{U}_i(a; \mu_0) = \mathbb{E}_{v \sim \mu_0}\left[\sum_{r \in a_i} v_r f(|a|_r)\right]$. Let $\mathrm{NE}(G, \mu_0, f) = \mathrm{NE}(G, \mathbb{E}_{v \sim \mu_0}[v], f)$. Additionally, the objective of the maximum coverage problem is to maximize the expected welfare, i.e., $\overline{W}(a; \mu_0) = \mathbb{E}_{\mu_0}\left[\sum_{r \in \cup_{i \in N} a_i} v_r\right]$.

As a means to try and improve the performance of the distributed decision-making agents, we may consider revealing information about the realization of the system state. One option is to reveal full information (or let agents know the realization exactly); however, either due to communication constraints or by design choice, it is often meaningful to reveal only partial information as well. In line with the broader Bayesian Persuasion framework, consider revealing information with an *information signaling policy* $\Pi = \{\pi_1, \ldots, \pi_m\}$, where $\pi_k \subseteq \mathcal{V}$, $\pi_j \cap \pi_k = \emptyset$, and $\bigcup_{k=1}^m \pi_k = \mathcal{V}$. This signaling policy Π forms a partition over the support of our random state variable v. The signal $\pi \in \Pi$ is revealed to the agents when $v \in \pi$.[1] The new system operates as follows: A system operator adopts a signaling policy Π and utility rule f. A state v is drawn from

[1]In this work, we consider signaling policies that are deterministic mappings from state to signal. In general, signal π could be drawn randomly based on state v. Many of the results easily generalize to this setting, but for ease of exposition and relevance to our problem setting of informing designed decision-makers, we present the results for deterministic signaling by treating Π as a partition of $\mathcal{V}$.

μ_0, and all agents are informed of which element of Π the realization belongs to. The corresponding signal π is sent, and each player $i \in N$ computes the posterior belief on the realization $\mu_\pi(x) = \mathbb{P}[v = x | v \in \pi] = \mu_0(x) / \left(\sum_{v' \in \pi} \mu_0(v')\right)$ if $v \in \pi$ and zero otherwise. The agents then seek to maximize their expected utility with the posterior belief.

Agents may now condition their action on the received signal. Let $\alpha \in \mathcal{A}^\Pi$ denote a joint strategy, where an element $\alpha_i(\pi) \in \mathcal{A}$ captures the action agent i takes when they receive signal π. In a strategy profile α, agent i has an expected payoff of $\overline{U}_i(\alpha; \mu_0, \Pi) = \mathbb{E}_{v \sim \mu_0}\left[\sum_{r \in \alpha_i} v_r f(|\alpha|_r)\right]$. Note that α_i is implicitly a function of the received signal π, which itself is determined by the state variable v; as such, each of α_i, π, and v are random variables. The expected welfare becomes $\overline{W}(\alpha; \mu_0, \Pi) = \mathbb{E}_{v \sim \mu_0}\left[\sum_{r \in \cup_{i \in N} \alpha_i} v_r\right]$.

When agents follow best-response dynamics, the set of fixed points becomes the set of *Bayes-Nash equilibria*, $\mathrm{BNE}(G, \mu_0, \Pi, f)$. A strategy α^{BNe} in this set satisfies

$$\overline{U}_i(\alpha^{\mathrm{BNe}}; \mu_0, \Pi) \geq \overline{U}_i(\alpha_i', \alpha_{-i}^{\mathrm{BNe}}; \mu_0, \Pi) \; \forall \alpha_i' \in \mathcal{A}_i^\Pi. \tag{3.2}$$

The signaling policy Π will alter this equilibria set and thus the guarantees of our distributed solution to the maximum coverage problem.

Our motivation for this work is understanding how giving agents information can affect system welfare. Due to the multiplicity of equilibria, we consider two perspectives: the *optimistic perspective* in which the system designer cares about the best attainable system performance, and the *pessimistic perspective* in which the system designer cares about the worst possible performance. For the optimistic perspective, the system designer evaluates a signaling policy Π by its effect on the system welfare in the best-case equilibrium; as such, let the optimistic *value-of-informing* with signaling policy Π be

$$\mathrm{VoI}^+(G, \mu_0, \Pi, f) = \frac{\max_{\alpha \in \mathrm{BNE}(G, \mu_0, \Pi, f)} W(\alpha; \mu_0, \Pi)}{\max_{a \in \mathrm{NE}(G, \mu_0, f)} W(a; \mu_0)},$$

which measures the gain in optimistic welfare by using policy Π. Similarly, for the pessimistic perspective, the system operator evaluates a signaling policy Π by its effect on the system welfare in the worst-case equilibrium; let the pessimistic value-of-informing be

$$\mathrm{VoI}^{-}(G,\mu_0,\Pi,f)=\frac{\min_{\alpha\in \mathrm{BNE}(G,\mu_0,\Pi,f)} W(\alpha;\mu_0,\Pi)}{\min_{a\in \mathrm{NE}(G,\mu_0,f)} W(a;\mu_0)},$$

which is the same ratio but now with the worst-case equilibrium strategy and allocation. These values inform a system operator of how revealing information will affect their equilibrium guarantees. They differ from the well-known price-of-anarchy/stability measures in that they relate two equilibrium performances rather than equilibrium to optimal.

3.3 Main Results

The main contribution of this work is in lower and upper bounding the value-of-informing for best- and worst-case equilibria. These bounds depend on agents' local decision-making process. In Section 3.3.1, we focus on the case where agents' payoffs are aligned with the system objective and find that revealing information improves the best-case equilibrium ($\mathrm{VoI}^{+}\geq 1$) but can worsen the worst-case equilibrium ($\mathrm{VoI}^{-}\leq 1$). In Section 3.3.2, we generalize these bounds to any local utility rule f. In Section 3.3.3, we observe a trade-off in the lower bounds on VoI^{+} and VoI^{-}; Fig. 3.2 characterizes this trade-off and highlights the fact that altering the local objectives of agents affects the efficacy of revealing information.

3.3.1 Marginal Contribution

The first utility design we will consider is the marginal contribution, where each agent makes decisions that maximize their contribution to the system welfare, i.e., $U_i(a) =$

$W(a) - W(a_{-i})$. This can be expressed by the local utility rule $f^{\mathrm{mc}}(x) := \mathbb{1}_{[x=1]}$. When agents follow this utility rule, their preferences are aligned with global welfare. This utility rule has the property that it maximizes the best-case equilibrium guarantee, known as the price-of-stability ratio [86]. However, as of yet, the effect of revealing information to decision-makers using this utility function has not been addressed. In Theorem 3.3.1, we address this question by providing lower and upper bounds on the value-of-informing for the best- and worst-case equilibria.

Theorem 3.3.1. *In a Bayesian Maximum Coverage problem* (G, μ_0, Π)*, with utility rule* f^{mc}*, the value-of-informing for the best-case equilibrium satisfies*

$$1 \leq \mathrm{VoI}^{+}(G, \mu_0, \Pi, f^{\mathrm{mc}}) \leq |\Pi|, \tag{3.3a}$$

and for the worst-case equilibrium satisfies

$$1/2 \leq \mathrm{VoI}^{-}(G, \mu_0, \Pi, f^{\mathrm{mc}}) \leq 2|\Pi|. \tag{3.3b}$$

All of these bounds are tight, but the upper bounds on VoI^{-}*.*

Before proving the statement, we discuss the consequences of these results. We first see that revealing more information (increasing the cardinality of Π) provides significant opportunities for improvement in either perspective[2]. However, revealing information need not always have such a positive effect. In the optimistic perspective, revealing information can only improve performance ($\mathrm{VoI}^{+} \geq 1$); however, doing so does not come without consequence, as revealing information can reduce the quality of the worst-case equilibrium by a factor of $1/2$ ($\mathrm{VoI}^{-} = 1/2$). This fact affirms an important property of information in a multi-agent system: revealing information must be done carefully.

[2]A more refined understanding of this improvement can be attained by considering the difference between the realizations and the prior mean.

The proof relies on the following lemma characterizing Bayes-Nash equilibria and the expected welfare.

Lemma 1. *A joint strategy α is a Bayes-Nash equilibrium if and only if*

$$(\alpha_1(\pi), \ldots, \alpha_n(\pi)) \in \mathrm{NE}(G, \mathbb{E}[v \mid \pi_k], f)$$

for each $\pi \in \Pi$. Additionally, the expected welfare of a joint strategy α is a weighted average of the welfare of the joint actions $\alpha(\pi)$ in the respective deterministic games, i.e.,

$$\overline{W}(\alpha; \mu_0, \Pi) = \sum_{k=1}^{m} p_k W\left(\alpha(\pi_k); \mathbb{E}[v|\pi_k]\right),$$

where $p_k = \sum_{v \in \pi_k} \mu_0(v)$.

Proof. Let $\alpha \in \mathcal{A}^{\Pi}$ denote a joint strategy. We show the first claim by observing the following transformation for any α,

$$\overline{U}_i(\alpha; \mu_0, \Pi) = \mathbb{E}\left[\sum_{r \in \mathcal{R}} \mathbb{E}[v_r|\pi] f(|\alpha(\pi)|_r)\right] \tag{3.4a}$$

$$= \sum_{k=1}^{m} p_k U_i\left(\alpha(\pi_k); \mathbb{E}[v|\pi_k]\right), \tag{3.4b}$$

where (3.4a) holds from the law of total expectation. Because α can be any m-tuple of joint actions, (3.2) is satisfied if and only if $U_i\left(\alpha(\pi); \mathbb{E}[v|\pi]\right) \geq U_i\left(a_i', \alpha_{-i}(\pi); \mathbb{E}[v|\pi]\right)$, for all $a_i' \in \mathcal{A}_i$, $\pi \in \Pi$; or, that $\alpha(\pi)$ is a Nash equilibrium for the deterministic game G with values $\mathbb{E}[v|\pi]$ for each $\pi \in \Pi$. The second claim follows (3.4) with welfare in place of utility. □

Proof of Theorem 3.3.1: Best-case equilibrium - We will make use of the function $W^\star(v)$, which denotes the welfare of an optimal allocation in G when the values are v. We note that with the marginal contribution utility rule, each optimal allocation $a^{\rm opt}$ is an equilibrium; thus, the welfare of the best Nash equilibrium is the optimal welfare [86], or $W^\star(v) = \max_{a\in\mathcal{A}} W(a;v) = \max_{a^{\rm Ne}\in{\rm NE}(G,v,f^{\rm mc})} W(a^{\rm Ne})$. We first make several observations about the function $W^\star$. Observe that $W^\star(v) = \max_{a\in\mathcal{A}} \sum_{r\in\mathcal{R}} v_r \mathbb{1}_{[|a|_r>0]}$, is the point-wise maximum of a set of affine (and thus convex) functions of v, which is itself convex. Further, $W^\star$ is positively homogeneous, i.e., $W^\star(\lambda v) = \lambda W^\star(v)$ for all $\lambda \geq 0$ and $v \geq 0$, and $W^\star$ is monotone in v, i.e., $v \succeq v' \Rightarrow W^\star(v) \geq W^\star(v')$ where "$\succeq$" denotes the element-wise inequality.

Using the properties of $W^\star$, we will prove the bounds on ${\rm VoI}^+$. First, the lower bound. Consider the Bayesian covering game (G, μ_0, Π). Observe that

$$\max_{a\in{\rm NE}(G,\mu_0,f^{\rm mc})} \overline{W}(a;\mu_0) = W^\star\left(\sum_{k=1}^{m} p_k \mathbb{E}[v|\pi_k]\right) \tag{3.5a}$$

$$\leq \sum_{k=1}^{m} p_k W^\star(\mathbb{E}[v|\pi_k]) \tag{3.5b}$$

$$= \max_{\alpha\in{\rm BNE}(G,\mu_0,\Pi,f^{\rm mc})} \overline{W}(\alpha;\mu_0,\Pi), \tag{3.5c}$$

where (3.5a) holds from the fact the maximum-welfare Nash equilibrium is a system optimum with values $\mathbb{E}_{\mu_0}[v]$ and the law of total probability, (3.5b) holds from $W^\star$ convex and Jensen's inequality, and (3.5c) holds from the second claim of Lemma 1 and the first claim of Lemma 1 with the fact the maximum-Nash is a system optimum. Rearranging terms gives the first inequality in (3.3a). It is tight when $|\mathcal{V}| = 1$.

Now, we consider the upper bound on VoI^+.

$$\max_{\alpha \in \mathrm{BNE}(G,\mu_0,\Pi,f^{\mathrm{mc}})} \overline{W}(\alpha;\mu_0,\Pi) = \sum_{k=1}^{m} p_k W^\star(\mathbb{E}[v|\pi_k]) \tag{3.6a}$$

$$= \sum_{k=1}^{m} W^\star\left(p_k \mathbb{E}[v|\pi_k]\right) \leq \sum_{k=1}^{m} W^\star(\mathbb{E}_{\mu_0}[v]) \tag{3.6b}$$

$$= |\Pi| \left(\max_{a \in \mathrm{NE}(G,\mu_0,f^{\mathrm{mc}})} \overline{W}(a;\mu_0) \right), \tag{3.6c}$$

where (3.6a) holds from the fact a Bayes-Nash joint strategy α is an m-tuple of Nash equilibria, the maximum-welfare Nash equilibrium is optimal in $W(\cdot\,;v)$ and the second claim in Lemma 1. (3.6b) holds from $W^\star$ positive homogeneous and the monotonicity of $W^\star$; more specifically,

$$\mathbb{E}_{\mu_0}[v_r] = \sum_{k \in [m]} p_k \mathbb{E}[v_r|\pi_k]$$

$$= p_k \mathbb{E}[v_r|\pi_k] + \sum_{k' \in [m] \setminus k} p_{k'} \mathbb{E}[v_r|\pi_{k'}] \geq p_k \mathbb{E}[v_r|\pi_k],$$

holds $\forall\, r \in \mathcal{R}$, $\pi_k \in \Pi$. (3.6c) holds from the definition $W^\star$.

To see this is tight, consider a problem with R resources and one agent who can select a single one of them, i.e., $\mathcal{A}_1 = \{1, \ldots, R\}$. Each resource can take on one of two values: 0 or 1. The prior μ_0 is that exactly one resource is ever of value 1 with equal probability; so the support of v has R elements, each of which occurs with probability $1/R$. When uninformed, the single agent is indifferent over their actions and cannot attain a payoff greater than $1/R$. When informed, the single agent can always select the resource of value 1. Thus $\mathrm{VoI}^+ = R = |\mathcal{V}| = |\Pi|$.

Worst-case equilibrium - To focus on equilibrium strategies, let $a^{\mathrm{Ne}}(v)$ denote a Nash equilibrium joint action in the game G when the values are v. The following steps will hold for any Nash equilibrium, and the proof will be completed by considering $a^{\rm Ne}(v)$ as the welfare minimizing Nash equilibrium. Observe that the uninformed Nash welfare satisfies

$$W(a^{\rm Ne}(\mathbb{E}_{\mu_0}[v]);\mathbb{E}_{\mu_0}[v]) \leq W^\star(\mathbb{E}_{\mu_0}[v]) \leq \sum_{k=1}^{m} p_k W^\star(\mathbb{E}[v|\pi_k]) \leq \sum_{k=1}^{m} p_k 2W(a^{\rm Ne}(\mathbb{E}[v|\pi_k]);\mathbb{E}[v|\pi_k]),$$

where the first inequality holds from the definition of $W^\star$, the second holds from properties of $W^\star$ shown in the first part of the proof, and the third holds from the price-of-anarchy bound of $1/2$ [78]. Letting $a^{\rm Ne}(v)$ be the worst-case Nash equilibrium when the values are v in each occurrence, the rightmost expression is the worst-case welfare in a Bayes-Nash joint strategy via Lemma 1. This gives the first inequality in (3.3a).

To see that this bound is tight, consider a resource allocation game with three resources, $\mathcal{R} = \{1,2,3\}$, and two players, $N = \{1,2\}$ with two actions each: $\mathcal{A}_1 = \{r_1, r_2\}$ and $\mathcal{A}_2 = \{r_2, r_3\}$. Let resource r_1 have value $v_1 = 1$ and r_3 have value $v_3 = 0$. Let the value of resource r_2 be a random variable, where $v_2 = 1-\varepsilon$ with probability $1-p$ and $v_2 = 1+\varepsilon(1-p)$ with probability p. When the agents are not informed of the realization of v_2, its expected value is $\mathbb{E}[v_2] = 1-\varepsilon(1-p)^2 < 1$ and there is a unique Nash equilibrium of $a = (r_1, r_2)$, providing a welfare of $W(a;\mu_0) = 2-\varepsilon(1-p)^2$. When the full revelation signaling policy $\Pi = \{\{r_1\},\{r_2\}\}$ is used, then when $v_2 = 1-\varepsilon$, there is a unique equilibrium of $\alpha^1 = (r_1, r_2)$, and when $v_2 = 1+\varepsilon(1-p)$ there are two equilibria, the worse of which is $\alpha^2 = (r_2, r_3)$. This gives an expected welfare of $W(\alpha;\mu_0) = (1-p)(2-\varepsilon)+p(1+\varepsilon(1-p))$ and a gain of informing agents of $\mathrm{VoI}^-(\Pi, f^{\rm mc}) = \frac{(1-p)(2-\varepsilon)+p(1+\varepsilon(1-p))}{2-\varepsilon(1-p)^2}$. Letting $p \to 1$ and $\varepsilon \to 0$ we get $\mathrm{VoI}^-(\Pi, f^{\rm mc}) \to 1/2$.

Finally, we show the upper bound on VoI^-. Observe that the welfare of a Bayes-Nash

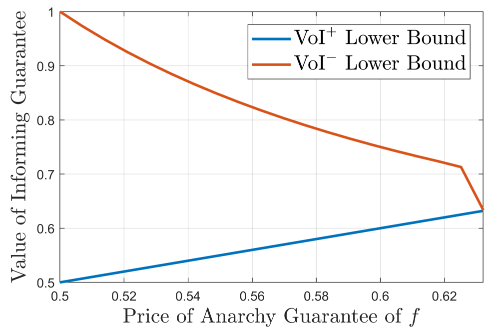

Figure 3.2: Lower bounds on the value of revealing information for the best- and worst-case equilibria. If the utility rule f is designed to lessen the loss to worst–case equilibria (increasing VoI$^-$), then there is a greater possible loss to the worst–case equilibrium (decreasing VoI$^+$). This trade-off matches the tight bounds from Theorem 3.3.1 and Proposition 3.3.1, which appear as the endpoints of this plot. The bounds are generated by comparing the value-of-informing to the price-of-anarchy/stability via Theorem 3.3.2 and the price of stability to the price-of-anarchy via [86, Theorem 4.1].

strategy $\alpha = \{a^{\rm Ne}(\mathbb{E}[v|\pi])\}_{\pi\in\Pi}$ under the signaling policy Π satisfies

$$\sum_{k=1}^{m} p_k W(a^{\rm Ne}(\mathbb{E}[v|\pi_k]);\mathbb{E}[v|\pi_k]) \leq \sum_{k=1}^{m} p_k W^\star(\mathbb{E}[v|\pi_k])$$
$$\leq |\Pi|\cdot W^\star(\mathbb{E}_{\mu_0}[v]) \leq 2|\Pi|\cdot W(a^{\rm Ne}(\mathbb{E}_{\mu_0}[v]);\mathbb{E}_{\mu_0}[v]),$$

where the first term is the expected payoff of α, and the first inequality holds from $W^\star$ being the optimal welfare, the second holds from the upper bound on VoI$^+$, and the third holds from the price-of-anarchy bound. □

3.3.2 Utility Design & Informing Efficacy

In Section 3.3.1, we considered the special case in which agents' local objectives were aligned with the global objective using the marginal cost utility rule $f^{\rm mc}$. In this section, we generalize this result to any utility rule by leveraging a connection to the *price-of-stability* and *price-of-anarchy*.

In the deterministic setting, the price-of-stability/anarchy are used to quantify how the best-case and worst-case equilibria approximate the system optimal. These metrics can be generalized to the Bayesian setting but need not be informative or insightful on the effects of revealing information within a single problem instance (i.e., failing to capture the benefits and consequences of comparing bounds derived from different problem instances). Let $\mathrm{PoA}(G, v, f) = \frac{\min_{a^{\mathrm{Ne}} \in \mathrm{NE}(G,v,f)} W(a^{\mathrm{Ne}};v)}{\max_{a^{\mathrm{opt}} \in \mathcal{A}} W(a^{\mathrm{opt}};v)}$ denote the price-of-anarchy for a deterministic maximum coverage problem, and let $\mathrm{PoS}(G, v, f) = \frac{\max_{a^{\mathrm{Ne}} \in \mathrm{NE}(G,v,f)} W(a^{\mathrm{Ne}};v)}{\max_{a^{\mathrm{opt}} \in \mathcal{A}} W(a^{\mathrm{opt}};v)}$ denote the price-of-stability.

Though not immediately apparent, we establish a connection between the price-of-stability/anarchy in deterministic covering games and the value-of-informing in Bayesian covering games. In Theorem 3.3.2, we leverage this connection to generate bounds on VoI^+ and VoI^- for any utility design.

Theorem 3.3.2. *Let* $\psi := \inf_{v \in \mathrm{conv}(\mathcal{V})} \mathrm{PoS}(G, v, f)$ *and* $\rho := \inf_{v \in \mathrm{conv}(\mathcal{V})} \mathrm{PoA}(G, v, f)$, *then the value of informing for the best-case equilibrium satisfies*

$$\psi \leq \mathrm{VoI}^+(G, \mu_0, \Pi, f) \leq \psi^{-1}|\Pi|, \tag{3.7a}$$

and for the worst-case equilibrium satisfies

$$\rho \leq \mathrm{VoI}^-(G, \mu_0, \Pi, f) \leq \rho^{-1}|\Pi|. \tag{3.7b}$$

Proof of Theorem 3.3.2: The proof will rely on the function $W^\star(v) = \max_{a \in \mathcal{A}} W(a; v)$ and several of its properties shown in the proof of Theorem 3.3.1. First, we prove the bounds on VoI^+. Let $a \in \mathrm{NE}(G, \mathbb{E}_{\mu_0}[v], f)$ be an arbitrary Nash equilibrium in the deterministic game G with values $\mathbb{E}_{\mu_0}[v]$ and utility rule f, and let $\alpha \in \mathrm{BNE}(G, \mu_0, \Pi, f)$ be an arbitrary Bayes-Nash equilibrium in G with prior μ_0 on v and information signaling policy Π and utility function f. Observe that the expected welfare of α satisfies

$$\begin{aligned}\overline{W}(\alpha;\mu_0,\Pi) &= \sum_{k=1}^{m} p_k W(\alpha(\pi_k);\mathbb{E}[v|\pi_k]) && \text{(3.8a)}\\ &\leq \sum_{k=1}^{m} p_k W^\star(\mathbb{E}[v|\pi_k]) && \text{(3.8b)}\\ &\leq |\Pi| W^\star(\mathbb{E}_{\mu_0}[v]) && \text{(3.8c)}\\ &\leq \rho^{-1}|\Pi|\overline{W}(a;\mathbb{E}_{\mu_0}[v]), && \text{(3.8d)}\end{aligned}$$

where (3.8a) holds from Lemma 1, (3.8b) holds from the definition of $W^\star$, (3.8c) holds from the monotonicity and positive homogeneity of $W^\star$ (previously shown in (3.6b)-(3.6c)), and (3.8d) holds from the definition of ρ.

Similarly, we can show that the expected total welfare of the uninformed equilibrium a satisfies

$$\begin{aligned}\overline{W}(a;\mu_0) &\leq W^\star(v) && \text{(3.9a)}\\ &\leq \sum_{k=1}^{m} p_k W^\star(\mathbb{E}[v|\pi_k]) && \text{(3.9b)}\\ &\leq \sum_{k=1}^{m} p_k \rho^{-1} W(\alpha(\pi);\mathbb{E}[v|\pi_k]) && \text{(3.9c)}\\ &= \rho^{-1} W(\alpha;\mu_0,\Pi), && \text{(3.9d)}\end{aligned}$$

where (3.9b) holds from the convexity of $W^\star$, (3.9c) holds from the definition of ρ, and (3.9d) holds from Lemma 1. Because (3.8) and (3.9) hold for any $a \in \mathrm{NE}(G,\mathbb{E}_{\mu_0}[v],f)$ and $\alpha \in \mathrm{BNE}(G,\mu_0,\Pi,f)$, it holds for each being the respective welfare minimizing equilibria. If we consider only the case where a and α are the welfare maximizing equilibria, i.e., $a \in \arg\max_{a' \in \mathrm{NE}(G,\mathbb{E}_{\mu_0}[v],f)} W(a';\mathbb{E}_{\mu_0}[v])$ and $\alpha \in \arg\max_{\alpha' \in \mathrm{BNE}(G,\mu_0,v,f)} \overline{W}(\alpha';\mu_0,\Pi)$ □

3.3.3 Optimistic / Pessimistic Trade-Off

Theorem 3.3.2 highlighted the fact that altering the utility design will change the impact of information revealing; however, the given bounds need not be tight. It appears that using a utility design with a higher price-of-anarchy in the deterministic setting will lead to an improved lower bound on VoI^-. As such, we will more closely consider the price-of-anarchy maximizing rule

$$f^{\mathrm{g}}(x) := (x-1)!\frac{\frac{1}{(n-1)(n-1)!}+\sum_{i=x}^{n-1}\frac{1}{i!}}{\frac{1}{(n-1)(n-1)!}+\sum_{i=1}^{n-1}\frac{1}{i!}},$$

when $x > 0$, and $f^{\mathrm{g}}(0) = 0$, proven optimal in [36]. In Proposition 3.3.1, we show tight lower bounds on VoI with f^g.

Proposition 3.3.1. *While using the the price-of-anarchy maximizing rule* f^{g}*, the value-of-informing for the best-case equilibrium satisfies*

$$1-\frac{1}{e} \leq \mathrm{VoI}^+(G, \mu_0, \Pi, f^{\mathrm{g}}), \tag{3.10a}$$

and for the worst-case equilibrium satisfies

$$1-\frac{1}{e} \leq \mathrm{VoI}^-(G, \mu_0, \Pi, f^{\mathrm{g}}), \tag{3.10b}$$

Further, each of these lower bounds is tight.

Proof of Proposition 3.3.1: Theorem 3.3.2 can be used to show that each of (3.10a) and (3.10) are valid lower bounds. To see these bounds are tight, consider a maximum coverage problem with resource set $\mathcal{R} = \{r_{p_i}\}_{i=1}^{n-1} \bigcup \{r_{s_j}\}_{j=1}^{z}$ where $z = \lceil 1/(f^g(n) - \varepsilon) \rceil$. The first $n-1$ players can select a single resource from the public resources $\{r_{s_j}\}_{j=1}^{z}$ or their respective private resource r_{p_i} for player i, i.e., $\mathcal{A}_i = \{r_{s_1}, \ldots, r_{s_z}, r_{p_i}\}$, $\forall i \in N \setminus \{n\}$.

The final player has exactly one action to use all the shared resources simultaneously $\mathcal{A}_n = (r_{s_1}, \ldots, r_{s_z})$. Each private resource r_{p_i} has value $v_{r_{p_i}} = f^{\mathrm{g}}(n) - \varepsilon$ with probability one where $\varepsilon > 0$. The value of the shared resources are random; each takes on value 1 w.p. $1/z < f^{\mathrm{g}}(n) - \varepsilon$ and 0 otherwise, and follow distribution μ_0 such that exactly one is ever the high value. When uninformed, each of the first $n-1$ players strictly prefers their private resource, giving a unique equilibrium welfare of $W(a^{\mathrm{Ne}}; \mu_0) = 1+(n-1)(f^{\mathrm{g}}(n)-\varepsilon)$. Under the full information reveal policy Π, each of the $n-1$ agents strictly prefer to use the one shared resource of value 1, giving an expected welfare of $\overline{W}(\alpha^{\mathrm{BNe}}; \mu_0, \Pi) = 1$. Because each of the informed and uninformed equilibria are unique, this gives $\mathrm{VoI}^+(\mu_0, \Pi) = \mathrm{VoI}^-(\mu_0, \Pi) = \frac{1}{1+(n-1)(f^{\mathrm{g}}(n)-\varepsilon)} \to 1 - \frac{1}{e}$ as $\varepsilon \to 0$ and $n \to \infty$. □

Comparing the lower bounds of Theorem 3.3.1 and Proposition 3.3.1 highlights a trade-off between revealing information in the optimistic and pessimistic perspectives. We further examine this trade-off in Fig. 3.2. Using Theorem 3.3.2 and recent results of [86], we can characterize lower bounds on VoI^+ and VoI^- for different utility rules.

3.4 Conclusion and Future Work

We addressed the possible benefit and consequences of revealing information to local decision-makers in a distributed system. By lower and upper bounding the value-of-informing, this work (1) quantified the possible effects of information revealing and (2) identified a trade-off between the guarantees of revealing information in the optimistic and pessimistic perspective. Future work will answer the design question and develop methods to solve for information signaling policies that optimize expected system welfare.

Chapter 4

Robust Utility Design in Distributed Resource Allocation Problems with Defective Agents

The use of multi-agent systems to solve large-scale problems can be an effective method to reduce physical and computational burdens; however, these systems should be robust to sub-system failures. In this work, we consider the problem of designing utility functions, which agents seek to maximize, as a method of distributed optimization in resource allocation problems. Though recent work has shown that optimal utility design can bring system operation into a reasonable approximation of optimal, our results extend the existing literature by investigating how robust the system's operation is to defective agents and by quantifying the achievable performance guarantees in this setting. Interestingly, we find that there is a trade-off between improving the robustness of the utility design and offering good nominal performance. We characterize this trade-off in the set of resource covering problems and find that there are considerable gains in robustness that can be made by sacrificing some nominal performance.

4.1 Introduction

Multi-agents systems have emerged as a viable method of implementing distributed system operation. In teams of robots [97, 98], resource allocation problems [99, 100], autonomous mobility and delivery [101, 102], and many other large-scale systems, distributing certain decision making processes to individual agents can help reduce computational and communication burdens. Designing local control laws for agents that guarantee good system performance overall can be difficult; a promising method of solving this problem is by using tools from game theory to describe the system operation [103]. A system designer can assign each agent a local objective or utility function which they seek to maximize. By carefully designing these utility functions, the system designer can guarantee that the emergent system behavior is within a good approximation of the system optimal [104, 105]. Though the use of game theoretic techniques in distributed control is encouraging, the robustness of these utility designs is not well understood, and certain sub-system failures may cause significant degradation in the system performance.

In this work, we consider the problem of utility design in resource allocation problems in which each agent selects a set of resources (or tasks) with the objective of maximizing the system welfare through the agents' collective actions. To determine a preference over actions, each agent is assigned a local objective (or utility) function which they seek to maximize by repeatedly updating their action; the equilibria of this process are the Nash equilibria of the game formed between the agents with their payoffs described by their local objectives. By designing the utility rules of the agents, a system operator can alter the Nash equilibria and improve the equilibrium performance guarantees. The problem of utility design has been studied in many settings of resource allocation problem [36, 72, 78] and has been proven to be effective at offering reasonable approximations of the optimal system welfare. In this work, we seek to understand the robustness of utility rules to sub-system failures by investigating the impact of defective agents on performance guarantees.

Previous work has considered how sub-system failures can affect the operation of a distributed control system. In [106], the authors look at how performance of greedy submodular-maximization degrades as communication channels between agents are removed and find that certain design modifications can be beneficial. The authors of [107] introduce a framework to discuss the robustness properties of log-linear learning: a process in which users' action updates are noisy best responses. They find that the presence of a single heterogeneous agent can alter the long-term group behavior but do not discuss the impact this agent has on system welfare. In [108], the authors consider the effect of heterogeneous agents on opinion dynamics in networked systems and their impact on group consensus. In this work, we not only consider how these types of agents affect the system operation and equilibria but also how they impact the system welfare and achievable performance guarantees of a utility design. We further consider how existing design techniques can be modified to offer improved robustness to defective agents.

The results of this work give insight on the robustness of game theoretic techniques of distributed control in resource allocation problems with defective agents. We consider two types of defective agents: (1) stubborn agents which do not update their actions and do not contribute to the system welfare but do alter the agents' perceived utilities, and (2) failure-prone agents that have a probability of failing to contribute to the system objective. In either setting, the presence of these defective agents will alter the intended system operation and worsen system performance guarantees. In Section 4.3, we first leverage existing results on utility design for resource allocation to provide a tractable linear program which can be used to compute the optimal, robust utility design in the presence of either type of defective agent. This program not only tells us the structure of the optimal utility design, but also the associated performance guarantees. In

Section 4.4, to better understand the affect these defective agents have on system performance, we focus on the class of covering problems. We find that significant performance improvements are attainable by designing utility rules more robustly; however, these design modifications necessarily reduce the system performance in the nominal setting. In Theorem 4.4.1, we characterize the trade-off frontier between offering good nominal and robust performance in the presence of a finite number of stubborn agents. Similarly, in Section 4.4.2 we offer a numerical analysis to highlight a similar trade-off in the setting of failure-prone agents.

4.2 Preliminaries

Consider a multi-agent system comprised of a finite set of agents $N = \{1, \dots, n\}$ whose interactions comprise the operation of the system. Each agent $i \in N$ has a set of allowable actions $\mathcal{A}_i$ that they can employ. When each agent $i \in N$ has selected an action $a_i \in \mathcal{A}_i$, we denote the joint action or allocation of users decisions by the tuple $a = (a_1, \dots, a_n) \in \mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$. To quantify the efficacy of system operation, each joint action is mapped to a global objective value via a system welfare function $W : \mathcal{A} \to \mathbb{R}$.

The objective of a system operator is to design a distributed decision-making process that enables system operation at a joint action $a \in \mathcal{A}$ that maximizes the system welfare $a^{\text{opt}} \in \arg\max_{a \in \mathcal{A}} W(a)$. One popular approach to the design of such processes is to design a utility function $U_i : \mathcal{A} \to \mathbb{R}$ for each agent $i \in N$ in the system that will influence and guide the agent's local decision making [103]. By employing suitable distributed learning algorithms, e.g., fictitious play [109] or regret matching [110], a system operator can ensure that the resulting distributed process converges to an equilibrium of the designed games. This work will primarily focus on the notion of pure Nash equilibrium, which is characterized by a joint action profile $a^{\mathrm{Ne}} \in \mathcal{A}$ that satisfies

$$U_i(a_i^{\mathrm{Ne}}, a_{-i}^{\mathrm{Ne}}) \geq U_i(a_i', a_{-i}^{\mathrm{Ne}}), \quad \forall a_i' \in \mathcal{A}_i,\ i \in N, \tag{4.1}$$

where a_{-i} denotes the action of all players excluding player i. Here, it is important to highlight that the design of utility function shapes the underlying equilibria of the systems which in turn influences the performance of the underlying system at such equilibria.

One of the fundamental goals of a system operator is to design utility function that lead to highly efficient equilibria as measured with regard to the system welfare. We measure the efficiency of a Nash equilibrium by the ratio between the equilibrium system welfare and the optimal system welfare. This performance guarantee is termed the *price of anarchy*, and can formally be defined as

$$\mathrm{PoA}(G) = \frac{\min_{a \in \mathrm{NE}(G)} W(a)}{\max_{a \in \mathcal{A}} W(a)}, \tag{4.2}$$

where $G = (N, \mathcal{A}, \{U_i\}_{i \in N}, W)$ is a given game equipped with a performance metric and $\mathrm{NE}(G)$ is the set of Nash equilibria in the game. The effectiveness of utility design has been studied in many settings, and has been shown to be effective at providing a reasonable approximation for the system optimal [36, 100]. In fact, in many cases such an approach provides the same approximation ratio as the best centralized algorithms [111].

The focus of this paper is on class of multi-agent systems pertaining to resource allocation problems. Here, there is ground set of resources (or tasks) $\mathcal{R}$ and each agent $i \in N$ is associated with an admissible choice set $\mathcal{A}_i \subset 2^{\mathcal{R}}$. Further, each resource r is associated with a non-decreasing welfare function $w_r : \{0, 1, \ldots, n\} \to \mathbb{R}_{\geq}$, where $w_r(k)$ encodes the welfare accrued at resource r when $k \geq 0$ agents are at resource r. The system welfare can thus be written as

$$W(a) = \sum_{r \in \mathcal{R}} w_r(|a|_r), \tag{4.3}$$

where $|a|_r$ denotes the number of agents utilizing resource r in allocation a. As above, the goal of a system operator is to design agent utility functions to optimize the efficiency of the resulting equilibria. To that end, here we focus on the design of agent utility functions where each resource r is associated with a a *local utility rule* $f_r : \{0, \ldots, n\} \to \mathbb{R}_{\geq 0}$ that describes how agents should assess the benefit of selecting resource r given the behavior of the other agents in the system. For these utility designs to be robust to different problem instances, we assume that the local utility function is the same for each agent[1]. When a system designer chooses a set of local utility rules $f = \{f_r\}_{r \in \mathcal{R}}$, the utility of agent i given an action profile $a \in \mathcal{A}$ is of the form

$$U_i(a) = \sum_{r \in a_i} f_r(|a|_r). \tag{4.4}$$

We will express a resource allocation game by the tuple $G = (N, \mathcal{A}, \mathcal{R}, \{w_r, f_r\}_{r \in \mathcal{R}})$, as the utility functions and welfare functions are derived from the pairs $\{w_r, f_r\}_{r \in \mathcal{R}}$. A resource allocation game, with non-decreasing local utility functions permits a potential function and thus at least one Nash equilibrium exists [72].

When designing the local utility rules f, a system operator may have minimal information about the specific game instance G, such as uncertainty about the agents available

[1]Because we consider a worst-case analysis, we cannot differentiate between the role of two agents, as if we designed the utility rule of agent i and j differently, a problem may be realized where their roles are reversed. If the system designer had full knowledge of the problem structure, designing agents' utilities heterogeneously can certainly help; however, it is currently unknown as to whether player specific utility functions can help in improving worst-case performance guarantees across a class of problem instances. Additionally, adopting a local utility rule that is consistent for each player lets us maintain the potential game structure.

actions or the full set of resources. For example, in the context of content distribution networks and data caching, the system operator has little knowledge of the available servers and the paths over the internet that connect them [112], in ride-sharing, the requests of passengers are not known until they are placed [113], or in team formations, the tasks may be changing over time [114]. A common assumption in the literature is that the system operator has knowledge of the type of different resource welfare functions, which we express by $\mathcal{W}$, but is unsure of the specific game instance or the specific welfare functions employed. The goal of a system operator is to design a local utility rule f_r for each welfare function $w_r \in \mathcal{W}$. We define this association by the map $\mathcal{F} : \mathcal{W} \to \mathbb{R}^n_{\geq 0}$, where $f_r = \mathcal{F}(w_r)$. Accordingly, given a set of potential welfare functions $\mathcal{W}$ and a local utility rule $\mathcal{F}$, we denote the set of possible games as

$$\mathcal{G}_{\mathcal{W},\mathcal{F}} = \{(N, \mathcal{A}, \mathcal{R}, \{w_r, \mathcal{F}(w_r)\}_{r\in\mathcal{R}}) : w_r \in \mathcal{W}, \mathcal{A}_1, \dots, \mathcal{A}_n \subseteq 2^{\mathcal{R}}\}.$$

Note that $\mathcal{G}_{\mathcal{W},\mathcal{F}}$ consists of any resource allocation game with n agents where each resource $r \in \mathcal{R}$ employs the pair $\{w_r, f_r = \mathcal{F}(w_r)\}$ where $w_r \in \mathcal{W}$. Given the set of possible welfare functions $\mathcal{W}$, the goal of a system operator is to design the local utility rules $\mathcal{F}$ to optimize the price of anarchy over the family of games $\mathcal{G}_{\mathcal{W},\mathcal{F}}$, i.e.,

$$\mathrm{PoA}(\mathcal{G}_{\mathcal{W},\mathcal{F}}) = \inf_{G\in\mathcal{G}_{\mathcal{W},\mathcal{F}}} \mathrm{PoA}(G). \tag{4.5}$$

There are several recent works that highlight how to compute the local utility rule $\mathcal{F}$ that optimizes the price of anarchy given in (4.5) [72, 105].

4.3 Robust Utility Design

This work focuses on understanding the robustness of these game theoretic methods of distributed control to sub-system failures in the form of defective agents. As such, we seek to understand the impact of defective agents on the price of anarchy and what improvements are possible when agents' utilities are designed with the knowledge that these hazards may exist. Specifically, we consider two different modes of agent failures, stubborn agents and failure prone agents, and illustrate how one can design robust utility function to optimize performance guarantees in these non-ideal settings.

4.3.1 Stubborn Agents

We consider a setting where there exist a finite number of stubborn agents in the system. While these stubborn agents do not contribute to the system welfare, they do impact the behavior of the agents in the systems and can potentially diminish the quality of the resulting equilibria.

To that end, we consider a scenario where there are at most m defective agents that can each occupy a subset of the resources in $\mathcal{R}$. Because each stubborn agent can occupy any subset of the resources, we instead represent the allocation of stubborn agents by the number of stubborn agents on each resource. Accordingly, we will express the behavior of the stubborn agents by the tuple $d = \{d_r\}_{r \in \mathcal{R}}$, where $d_r \in \{0, 1, \ldots, m\}$ for each resource $r \in \mathcal{R}$. Further, we will express the collective behavior in our system by the tuple $(a; d)$ to reflect the behavior of both the nominal and stubborn agents. Here, it is important to highlight that the stubborn agents do not impact the system welfare, which stays of the form

$$W(a; d) = \sum_{r \in \mathcal{R}} w_r(|a|_r).$$

However, the presence of stubborn agents does affect the utility functions of the nominal agents who are unable to differentiate between the behavior of stubborn agents and other nominal agents. Here, the utility function for any agent $i \in N$, action profile $a \in \mathcal{A}$, and stubborn agent profile d is of the form

$$U_i(a; d) = \sum_{r \in a_i} f_r(|a|_r + d_r), \tag{4.6}$$

where $f_r \in \mathbb{R}^{n+m}_{\geq 0}$. We will now measure the price of anarchy associated with a given form of stubborn behavior $d \in \{0, \ldots, m\}^{|\mathcal{R}|}$ as

$$\text{PoA}(G, d) = \frac{\min_{a \in \text{NE}(G,d)} W(a)}{\max_{a^{\text{opt}} \in \mathcal{A}} W(a^{\text{opt}})}, \tag{4.7}$$

where $\text{NE}(G, d)$ defines the equilibria associated with the game defined by d given in (4.6). As stated before, the goal of a system operator is to design a local utility rule $\mathcal{F}$ that optimizes the worst-case performance measure over both game instances $G \in \mathcal{G}_{\mathcal{W},\mathcal{F}}$ and stubborn-agent behavior $d \in \{0, \ldots, m\}^{|\mathcal{R}|}$, i.e.,

$$\text{PoA}(\mathcal{G}^m_{\mathcal{W},\mathcal{F}}) = \inf_{G \in \mathcal{G}_{\mathcal{W},\mathcal{F}}} \min_{d \in \{0,\ldots,m\}^{|\mathcal{R}|}} \text{PoA}(G, d). \tag{4.8}$$

The following proposition demonstrates that one can find the local utility rule $\mathcal{F}$ that optimizes the price of anarchy via the solution of a tractable linear program.

Proposition 4.3.1. *Let* $\mathcal{W} = \{\sum_{b=1}^{B} \alpha_b w_b | \alpha_b \geq 0 \ \forall b \in [B]\}$ *be a set of resource value functions where* $w_b \in \mathbb{R}^n_{\geq 0}$ *is a basis welfare function. For each* $b \in [B]$*, let* $(f_b^\star, \mu_b^\star)$ *be* *the solution of the following linear program*

$$(f_b^\star, \mu_b^\star) \in \underset{f\in\mathbb{R}^{n+m},\ \mu\in\mathbb{R}}{\arg\max}\ \mu \tag{4.9}$$
$$\text{s.t. } w_b(z+y) - \mu w_b(x+y) + xf(x+y+d) - zf(x+y+d+1) \leq 0$$
$$\forall (x,y,z) \in \mathcal{I}_n,\ d \in \{0,\dots,m\},$$

where $\mathcal{I}_n = \{(x,y,z) \in \mathbb{N}^3_{\geq 0} \mid 1 \leq x+y+z \leq n\}$*. The following statements hold true for the family of resource allocation games with at most* n *agents and* m *stubborn agents:*

(i) There exists a linear local utility rule $\mathcal{F}^*$ *that optimizes the price of anarchy given in* (4.8)*. Furthermore, if* $w = \sum_{b=1}^B \alpha_b w_b$*, then* $\mathcal{F}^*(w) = \sum_{b=1}^B \alpha_b f_b^*$*.*

(ii) The optimal price of anarchy satisfies $\mathrm{PoA}(\mathcal{G}^m_{\mathcal{W},\mathcal{F}^*}) = \max_{b\in[B]} 1/\mu_b^\star$*.*

The linear program (4.9) is a generalization of [105, Theorem 3.7] which solves the local utility design problem in the nominal setting and is recovered when $m = 0$. The proof appears in Appendix A.2.1. Here, the local utility function f is treated as a vector in $\mathbb{R}^{n+m}$ where $f(i)$ denotes what the utility rule evaluates to when i agents utilize a resource; the linear program thus has $n + m + 1$ decision variables and $|\mathcal{I}_n|$ constraints. It is interesting to note that the optimal local utility design in a class of games of the described form decomposes into finding an optimal basis local utility function for each basis value function (a phenomenon first observed in [7]), making the computation of these utility rules more efficient. The parameterization using tuples of the form (x, y, z, d) is described in the appendix along with the proof. This approach not only solves for the optimal local utility rule, but also gives the associated price of anarchy guarantee. By comparing the solution of (4.9) when $m = 0$ and when $m > 0$, we can observe the impact stubborn agents have on the capabilities of a system designer. The magnitude of this performance degradation is discussed in Section 4.4.1, along with an investigation of the impact the design modifications that promote robustness have on the nominal performance of the system.

4.3.2 Failure Prone Agents

The second type of defective agent we consider are failure prone agents, where each agent operates normally but has a probability of failing and not contributing to the global objective. Every agent will follow the designed utility rule but has a chance of failing and no longer contributing to the system welfare. In contrast to the defective agents in the previous section which were ineffective and stubborn in their action selection, here an agent that fails can be thought of as ineffective but still updating their actions as a normal agent would. Additionally, each agent independently fails to contribute to the welfare with probability p. In a resource allocation problem G, each agent (failed or not) will follow their best response dynamic until the system reaches a Nash equilibrium $a^{\text{Ne}} \in \text{NE}(G)$. In an allocation a, a resource $r \in \mathcal{R}$, utilized by $|a|_r$ agents, has $X_r \leq |a|_r$ non-failed agents remaining with probability

$$\mathbb{P}[X_r = x] = \binom{|a|_r}{x} (1-p)^x \; p^{|a|_r - x}.$$

In a game G, the expected system welfare in an allocation a is

$$W(a) = \mathbb{E}\left[\sum_{\{r \in \mathcal{R} : X_r > 0\}} w_r(X_r) \right]. \tag{4.10}$$

The price of anarchy when agents are failure-prone $\text{PoA}(G;p)$ is thus the worst-case ratio between the expected system welfare in a Nash equilibrium and the optimal expected system welfare. The worst-case performance guarantee is a lower bound on the price of anarchy over resource covering games with a probability of failure p, i.e., $\text{PoA}(\mathcal{G}_{\mathcal{W},\mathcal{F}}; p)$. As described in Section 4.3.1, the optimal local utility design problem can be described as finding a mapping $\mathcal{F}$ that maximizes $\text{PoA}(\mathcal{G}_{\mathcal{W},\mathcal{F}}; p)$.

Proposition 4.3.1 provides a tool for computing local utility rules that are robust to stubborn agents for general local welfare functions. Additionally, we can amend the local welfare function and use the same linear program to compute the optimal local utility rule in the face of failure prone agents.

Corollary 3. *Let* $\mathcal{W} = \{\sum_{b=1}^{B} \alpha_b w_b | \alpha_b \geq 0 \ \forall b \in [B]\}$ *be a set of resource value functions where* $w_b \in \mathbb{R}^n_{\geq 0}$ *is a basis welfare function. If agents fail with probability* p*, then for each* $b \in [B]$*, let* $(f_b^\star, \mu_b^\star)$ *be the solution of the* (4.9) *with* $m = 0$ *and the amended value functions*

$$\overline{w}_b(x) = \sum_{k=0}^{x} w_b(k) \binom{x}{k} (1-p)^k \ p^{x-k} \quad \forall \ b \in [B]. \tag{4.11}$$

The following statements hold true for the family of resource allocation games with at most n *agents and probability of failure* p*:*

(i) There exists a linear local utility rule $\mathcal{F}^*$ *that optimizes the price of anarchy given in* (4.8)*. Furthermore, if* $w = \sum_{b=1}^{B} \alpha_b w_b$*, then* $\mathcal{F}^*(w) = \sum_{b=1}^{B} \alpha_b f_b^*$*.*

(ii) The optimal price of anarchy satisfies $\text{PoA}(\mathcal{G}_{\mathcal{W},\mathcal{F}^*}; p) = \max_{b \in [B]} 1/\mu_b^\star$*.*

4.4 Efficacy of Robust Design

In Section 4.3, a linear program was introduced that can be used to solve for the optimal local utility design and associated performance guarantee. In this section, we seek to better understand the impact of these defective agents and what opportunities are available by utilizing a more robust design. To answer this, we look at a specific setting of resource allocation problems called *covering problems*. Here, each resource r

contributes a fixed value to the system if it is utilized (or covered) by at least one agent. The welfare function for a covering problem can thus be written as

$$W(a) = \sum_{r \in \mathcal{R}} v_r \cdot w(|a|_r),$$

where $w(x)$ is an indicator function that takes value 1 if $x > 0$ and 0 otherwise. These problems are a specification of the more general problem description in Section 4.3 where there is now only a single basis function and $\mathcal{W} = \{\alpha \cdot w \mid \alpha > 0\}$. The focus to this setting allows for a more detailed analysis of robust performance as closed form solutions to general resource allocation problems with defective agents is still an open problem.

As discussed in Proposition 4.3.1, the optimal local utility rules satisfy a linearity property, therefore in the set of covering problems, we need only find a single local utility function $f : [n] \to \mathbb{R}_{\geq 0}$ that is scaled by the resources value v_r to determine the local utility, i.e.,

$$U_i(a) = \sum_{r \in a_i} v_r \cdot f(|a|_r).$$

As the welfare functions are described by a single basis, and we need only consider a single vector f to define a utility rule, the set of covering games can be denoted $\mathcal{G}_{w,f}$. For brevity, and because the welfare basis function will not change, we will simply denote this set of games $\mathcal{G}_f$ and note the covering welfare function is implied. Altering f will constitute changes to the local utility design. Finding the form of the optimal local utility rule f in closed form is not trivial. However, in [36] the authors find that the local utility rule

$$f^0(j) := (j-1)! \frac{\frac{1}{(n-1)(n-1)!} + \sum_{i=j}^{n-1} \frac{1}{i!}}{\frac{1}{(n-1)(n-1)!} + \sum_{i=1}^{n-1} \frac{1}{i!}}, \tag{4.12}$$

is optimal in the nominal covering setting and provides a price of anarchy guarantee of $1 - \frac{1}{e}$. In this work, we investigate how performance guarantees like this are impacted by the presence of defective agents. Additionally, we ask if appropriate design modifications to the local utility rule can make the system more robust to defective agents and what impact these design modifications have on the systems nominal performance.

4.4.1 Stubborn Agents in Covering

In Proposition 4.3.1, a linear program was introduced whose solution gives the optimal local utility rule for a class of games with stubborn agents. In this section, we seek to further understand the impact of stubborn agents on performance guarantees by finding the optimal, robust local utility rule and associated performance guarantee in covering problems. Interestingly, we find that though robust design modifications can improve robust performance, these changes necessarily reduce the utility rules performance in the nominal setting, thus highlighting a trade-off between guaranteeing good robust and nominal performance. We let $\mathcal{G}_f^m$ denote the set of covering games with local utility rule f and at most m stubborn agents. By comparing the price of anarchy guarantees in the nominal setting $\mathcal{G}_f^0$ and in the presence of stubborn agents $\mathcal{G}_f^m$, we can discuss not only the newly found robustness of a utility rule f, but also what performance guarantees it maintains in the nominal setting.

Theorem 4.4.1 quantifies the trade-off between utility rules that are robust to stubborn agents and utility rules with good nominal performance in terms of price of anarchy guarantees.

Theorem 4.4.1. *Let $t \in [0,1]$ be a chosen tuning parameter. In the class of covering games, if a local utility rule f achieves a robust price of anarchy guarantee of*

$$\mathrm{PoA}(\mathcal{G}_f^m) \geq \frac{\Gamma_m + \frac{e}{e-1}}{1 + t\Gamma_m}, \tag{4.13}$$

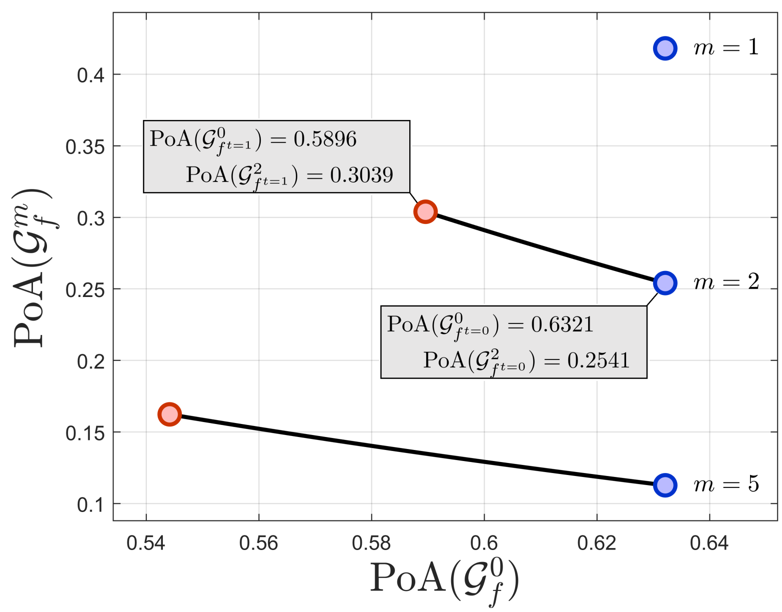

Figure 4.1: Achievable price of anarchy guarantees in the nominal setting $\mathcal{G}^0_f$ and in the presence of m stubborn agents $\mathcal{G}^m_f$. Each line represents a Pareto-optimal frontier for the achievable performance guarantee in each setting for a specific m. The left (red) endpoints represent the price of anarchy guarantees of the optimal, robust utility rule $f^{t=1}$ and the right (blue) endpoints represent the price of anarchy guarantees of the optimal, nominal utility rule $f^{t=0}$. A system designer is only capable of offering joint performance guarantees that are on the line connecting the endpoints or lower.

where $\Gamma_m = m!\frac{e-\sum_{i=0}^{m-1}\frac{1}{i!}}{e-1} - 1$*, then its nominal price of anarchy guarantee will be no better than*

$$\mathrm{PoA}(\mathcal{G}_f^0) \leq \frac{(e-1)(1+t\Gamma_m)}{1+(e-1)(1+t\Gamma_m)}. \tag{4.14}$$

Further, the price of anarchy guarantees in (4.13) *and* (4.14) *can be jointly realized by a local utility rule*

$$f^t(j) = f^0(j) - \max\left\{t\left(f^0(j) - \frac{m}{j}f^0(m)\right), 0\right\}. \tag{4.15}$$

The proof of Theorem 4.4.1 appears in Appendix A.2.2.

The trade-off described in (4.13) and (4.14) is depicted in Fig. 4.1 for several values of m. The horizontal axis measures the nominal price of anarchy and the vertical axis measures the price of anarchy when there are at most m stubborn agents. By choosing $t = 0$ the local utility rule optimizes the nominal price of anarchy guarantee and choosing $t = 1$ optimizes the robust price of anarchy guarantee. The line drawn by varying the parameter $t \in [0, 1]$ constitutes a Pareto-optimal frontier on the multi-objective problem of maximizing the nominal and robust performance guarantees.

By letting $t = 0$, we can evaluate the performance of the optimal, nominal utility rule $f^{t=0} = f^0$ defined in (4.12). Clearly, the performance degrades as more stubborn agents are introduced into the problem: the presence of two stubborn agents reduces the performance of the nominal utility rule $f^{t=0}$ by almost 60% down to $\mathrm{PoA}(\mathcal{G}_{f^{t=0}}^{m=2}) = 0.2541$. By designing the utility rule more robustly, the price of anarchy guarantee in $\mathcal{G}_f^{m=2}$ can be improved by almost 20% by using $f^{t=1}$; however, this increase in robustness comes at the cost of nominal performance, as the local utility rule $f^{t=1}$ is approximately 7% less efficient than the optimal in the nominal setting. A system designer who would like to optimize both performance metrics can provide guarantees only up to the Pareto-optimal frontier

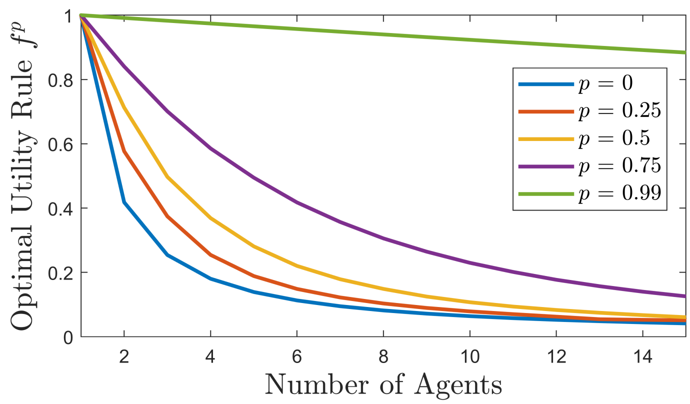

Figure 4.2: Optimal local utility rules that are robust to failure-prone agents for several probabilities of failure p. As the probability of agent failure increases, the optimal utility rule is larger for values of $k > 1$, thus incentivizing more overlap in agents resource usage.

described by (4.13) and (4.14) and shown in Fig. 4.1; these Pareto-optimal performance guarantees can be achieved by using f^t for $t \in [0, 1]$.

4.4.2 Failure-Prone in Covering

Optimal Utility Rules

Utilizing our results from Section 4.3.2, we perform a numerical analysis with (4.9) to understand the necessary design modifications and attainable performance guarantees with failure-prone agents in the setting of recource covering, i.e., w is the indicator function. From Corollary 3, we can compute optimal utility rules and price of anarchy guarantees using the augmented value function

$$\overline{w}(x) = \sum_{k=0}^{x} w(k) \binom{x}{k} (1-p)^k \, p^{x-k} = 1 - p^x,$$

and the linear program provided in (4.9) with $m = 0$. Via the solution to the linear program, we not only investigate how the optimal utility rule and price of anarchy change

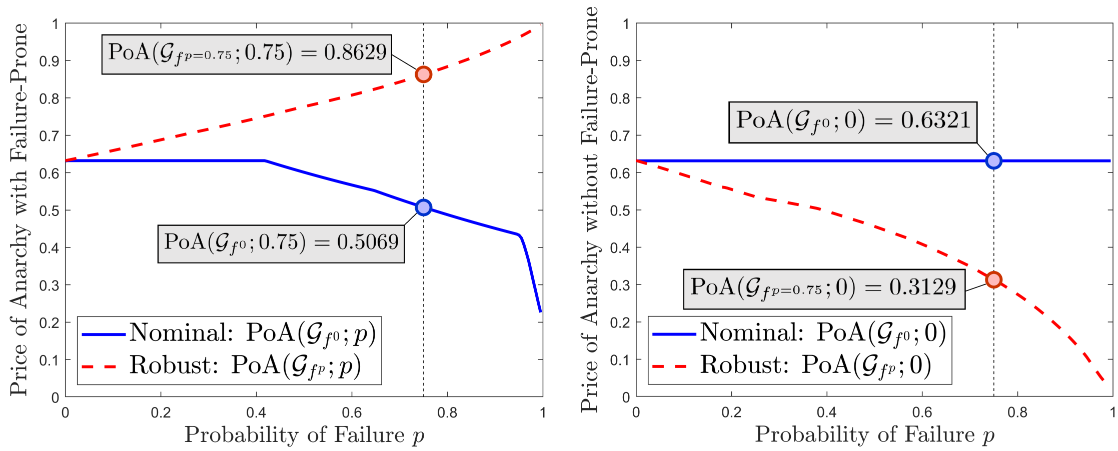

Figure 4.3: (a) Price of anarchy in covering problems with failure prone agents under nominal utility rule f^0 and robust utility rule f^p. The robust utility rule offers significant performance improvements in the presence of failure prone agents. (b) Price of anarchy in covering problems under nominal utility rule f^0 and robust utility rule f^p designed for agents failing with probability p. The robust utility rule sacrifices notable performance in the original setting.

with the probability of failure, we further investigate how the design modifications affect the nominal performance, had the agents not been failure prone. Additionally, we consider these performance metrics when the agents are not only failure-prone, but when there exist stubborn agents as well.

First, we utilize (4.9) to compute the optimal utility rule in the presence of failure-prone agents for several values of p, the probability failure. In Fig. 4.2, we see that the optimal local utility rule $f^p(k)$ increases with p for values of $k > 1$. Intuitively, this implies that it is optimal to design agents utilities as to promote more overlap in the resources they utilize. The larger the number of agents utilizing a resource will lead to a better chance that the resource will be covered by at least one non-failed agent. As p approaches 1, it is optimal for all agents to greedily, and without consideration of one another, to choose the most valuable set of resources.

Though these design modifications may make the system more robust to failure-prone agents, they may not be effective in the nominal setting.

Nominal & Robust Performance

In Fig. 4.3, we analyze the performance trade-off between designing utility rules for the nominal and failure-prone settings. For $p \in [0, 1]$, we compute the optimal utility rule f^p that is robust to agent failure with probability p using (4.9) with $m = 0$. We then compare the performance of this new, robust utility rule with the performance of the nominal utility rule f^0, defined in (4.12), when agents are failure-prone and when agents are not failure-prone. Fig. 4.3 (left) shows the price of anarchy guarantees of the nominal and robust utility rules in presence of failure-prone agents. When the probability of failure is large, the robust utility rule offers large improvements to the expected system welfare; when $p = 0.75$, the robust utility rule offers a price of anarchy guarantee of $\text{PoA}(\mathcal{G}_{f^{p=0.75}}; p = 0.75) = 0.8629$ which is a 63% increase from the performance of the nominal utility rule f^0 in this setting. However, as seen in Fig. 4.3 (right), if the system designer is incorrect and the agents are not failure-prone, the use of the robust utility rule causes a loss in performance; when $p = 0.75$ the price of anarchy of the robust utility rule without failure prone agents is $\text{PoA}(\mathcal{G}_{f^{p=0.75}}; p = 0) = 0.3129$, which is a 49% decrease from if the nominal had been used. This difference in price of anarchy guarantees again highlights a trade-off between the achievable nominal and robust performance. As the probability of failure increases, the optimal local utility rules in the robust setting struggle to offer good, nominal performance guarantees.

Performance Trade-off

The previous numerical result highlights a trade-off between optimizing nominal performance and robust performance in the setting of failure-prone. However Fig. 4.3 only captures the fact that optimizing both objectives becomes more difficult as the probability of failure increases. Similar to Section 4.3.1, we would also like to understand the trade-off between the objectives when good (but not optimal) performance is desired in

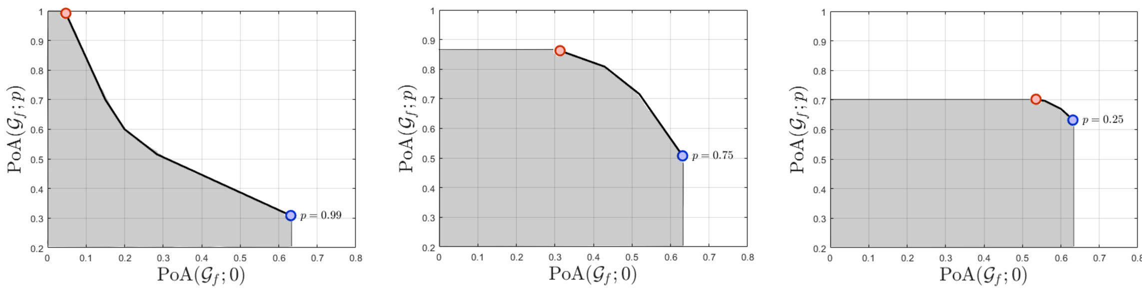

Figure 4.4: Trade-off between nominal and robust performance guarantees for agents that fail with probability $p = 0.99$ (left), $p = 0.75$ (middle), and $p = 0.25$ (right). Each line represents an empirical Pareto-optimal frontier for the achievable joint price of anarchy bound in each setting. The left (red) endpoints represent the price of anarchy guarantees of the optimal, robust utility rule f^p and the right (blue) endpoints represent the price of anarchy guarantees of the optimal, nominal utility rule f^0. By simulation, utility rules were generated that populated the grey region of each plot and up to the trade-off frontier.

both settings. In Fig. 4.4, we offer a bound on this trade-off by means of Monte Carlo simulation. By randomly generating many possible local utility rules f, and using (4.9) to compute the performance guarantees in the nominal and robust setting, we are able to generate a lower bound on the Pareto-optimal frontier. Fig. 4.4 demonstrates similar results to Fig. 4.1 in that the trade-off becomes more significant as the magnitude of the possible sub-system failures (here failure-prone agents) increases. Additionally, it is clear that reasonable concessions can be made to either the nominal or robust objective to guarantee better performance in both.

Failure-Prone and Stubborn Agents

Finally, we consider the case where agents are not only failure-prone (are ineffective but not stubborn), but that there also exists some defective, stubborn agents as introduced in Section 4.3.1. Here, we again utilize (4.9) with welfare function defined accordingly by Corollary 3 to represent the failure prone agent setting, but now with $m > 0$ stubborn agents. In Fig. 4.5, we see the price of anarchy guarantees unsurprisingly degrade as more stubborn agents are introduced into the system, however it is interesting to note that the relative amount of this degradation does not vary much with p. This essentially

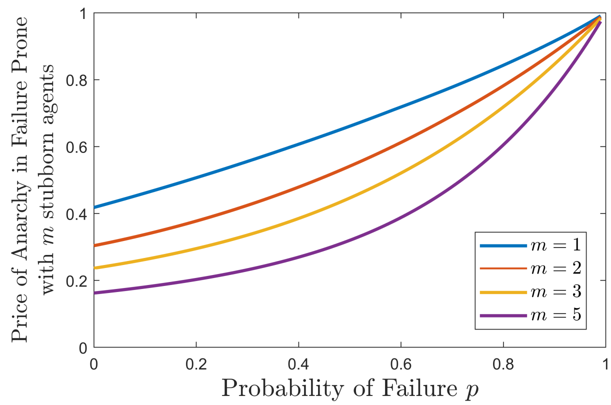

Figure 4.5: Price of anarchy guarantee in the setting of failure-prone agents where m stubborn agents are present. The presence of stubborn agents unsurprisingly lowers the performance guarantees, but the amount by which this happens appears to be relatively independent of p, the probability of failure.

shows that the two forms of failures do not cause cascading issues but do degrade system performance.

4.5 Conclusion

This work studies the robustness of local utility rules to sub-system failures in the form of stubborn and failure-prone agents. We provide linear programs that compute and evaluate the optimal local utility rules in the face of these defective agents. Our results show that there is a trade-off in designing utility rules that are robust and that give good nominal performance, which is characterized for the setting of covering problems.

Part II

Congestion Control and Social Influencing

Chapter 5

The Effectiveness of Subsidies and Tolls in Congestion Games

Are rewards or penalties more effective in influencing user behavior? This work compares the effectiveness of subsidies and tolls in incentivizing user behavior in congestion games. The predominantly studied method of influencing user behavior in network routing problems is to institute taxes which alter users' observed costs in a manner that causes their self-interested choices to more closely align with a system-level objective. Another conceivable method to accomplish the same goal is to subsidize the users' actions that are preferable from a system-level perspective. We show that, when users behave similarly and predictably, subsidies offer superior performance guarantees to tolls under similar budgetary constraints; however, in the presence of unknown player heterogeneity, subsidies fail to offer the same robustness as tolls.

5.1 Introduction

In systems governed by a collective of multiple decision making users, system performance is often dictated by the choices those users make. Though each user may make decisions rationally, the emergent behavior observed in the system need not align with the objective of the system designer. This phenomenon appears in many engineering settings including distributed control [103], resource allocation problems [115], electric power grids [116], and transportation networks [113], as well as many logistical problem settings such as marketing [117] and supply-chain management [118]. A prominent metric to quantify this emergent inefficiency is the *price of anarchy*, defined as the worst-case ratio between the social welfare experienced when users make self interested decisions and the optimal social welfare [119, 120].

A promising method of mitigating this inefficiency is by introducing incentives to the system's users, influencing their decisions to more closely align with the system optimal [121]. One example of such incentives is to levy *taxes*, eliciting monetary fees from users will affect their preferences over the available actions (e.g., tolls in transportation) [19, 122, 123]. Such taxes have been shown to be effective in reducing system inefficiency as measured by the price of anarchy ratio [71, 124, 7, 125]. Another method to influence user behavior is to *subsidize* the actions that are preferable from a system level perspective. Subsidies have been studied as a tool to influence users in transportation [126], supply chains [127], congestion [128], and emissions [129]. Though subsidies require the system operator to pay its users, it is possible that the savings obtained from efficient use of the infrastructure outweigh the cost incurred from the implemented incentives [130, 131]; additionally, one could consider implementing subsidies as rebates to a fixed, opt-in fee, to prevent a loss of revenue for the system operator. Though the use of subsidies is feasible in theory and in implementation, this method has been studied significantly less than the tax equivalent; the relative performance of each is thus unknown.

In this paper, we seek to understand the relative performance of subsidies and taxes in influencing user behavior in socio-technical systems. Specifically, we consider a network routing problem in which users must traverse a network with congestible edges with delays that grow as a function of the local mass of users. Finding a route for each user that minimizes the total latency in the system is straightforward if the system designer has full control in directing the users. However, when users select their own routes, the resulting network flow need not be optimal [33]. Modeling the selfish routing problem as a congestion game, we adopt the *Nash flow* as a solution concept of the emergent behavior in the system. From the users' selfish routing, the price of anarchy may be large [132]. To alleviate this emergent inefficiency, we introduce incentives to the users' which alter their observed costs and preferences. The objective of such incentives is to shape the users' preferences so the performance of the resulting Nash flow will improve.

A well studied method of incentivizing users in congestion games is to tax the users, i.e., introducing tolls to links in the network [19, 122, 123, 71, 124, 7, 125, 133, 134]. In each of these referenced works, the price of anarchy is used to measure the effectiveness of a tolling scheme. Indeed in the most elementary settings, tolls exist that influence users to self route in line with the system optimum [33]. However, when more nuance is introduced in the form of *player heterogeneity* (i.e., players differing in their response to incentives), the task of designing tolls becomes more involved. When the toll designer possesses sufficient knowledge of the network structure and user population, they may still compute and implement tolls which incentivize optimal routing [125]. However, in the case where the system designer has some uncertainty in the network parameters or behavior of the user population, it may not be possible to design tolls that give optimal system performance; thus, tolls are often designed to minimize inefficiency measured by the price of anarchy ratio [123, 135, 136], and again, encouraging results exist.

Though the study of tolling in congestion games is extensive, there are few results regarding subsidies as incentives in this context, especially in the presence of uncertain user heterogeneity. In [137], the authors investigate budget-balanced tolls in which the sum of all monetary transactions is zero, but the authors only consider homogeneous users. The authors of [138] give the first formal analysis of subsidies in congestion games and provide an algorithm that computes optimal rebates when users are homogeneous and the network structure is known. The authors of [139] consider more general incentives, but in an evolutionary setting. From a system designer's perspective, subsidies may be a feasible method of influencing user behavior; the performance guarantees of subsidies is thus of interest as well as how this performance compares to tax incentives.

Though there is a clear disparity in the breadth of results in the literature on tolls and subsidies, we bridge this gap by proving fundamental relationships between the performance and robustness of subsidies and tolls. Namely, subsidies offer better performance guarantees than tolls under budgetary constraints but are inherently less robust to user heterogeneity. The manuscript is outlined as follows:

Section 5.2.6: Performance of Incentives. In Theorem 5.2.1, it is shown in the nominal setting, where users behave similarly and predictably, that subsidies give better performance guarantees under similar budgetary constraints.

Section 5.3: Incentives with Heterogeneity. In Theorem 5.3.1 it is shown that tolls can effectively mitigate the negative effects of player heterogeneity while in Theorem 5.3.2 it is shown subsidies cannot.

Section 5.3.1: Robustness of Incentives. It is shown that tolls are more robust to uncertainty in the user population than subsidies. In the presence of a budgetary constraint, Theorem 5.3.3 shows that uncertainty degrades subsidy performance more rapidly than it degrades toll performance.

Section 5.4: Trade-off in Performance and Robustness. Given the contrast in the nominal performance of subsidies and the robustness of tolls to user heterogeneity, this fundamental relationship is analyzed between the two in parallel-affine congestion games by finding the level of uncertainty at which the robustness of tolls gives superior performance guarantees than subsidies.

In addition to finding general performance and robustness relationships between subsidies and tolls, we additionally find explicit price of anarchy bounds for optimal tolls and subsidies in several classes of congestion games to show that the differences in performance can be significant. We introduce tools to construct optimal incentives and corresponding performance guarantees.

5.2 Preliminaries

5.2.1 System Model

Consider a directed graph (V, E) with vertex set V, edge set $E \subseteq (V \times V)$, and k origin-destination pairs (o_i, d_i). Denote by $\mathcal{P}_i$ the set of all simple paths connecting origin o_i to destination d_i. Further, let $\mathcal{P} = \cup_{i=1}^{k} \mathcal{P}_i$ denote the set of all paths in the graph. A *flow* on the graph is a vector $f \in \mathbb{R}_{\geq 0}^{|\mathcal{P}|}$ that expresses the mass of traffic utilizing each path. The mass of traffic on an edge $e \in E$ is thus $f_e = \sum_{P:e \in P} f_P$, and we say $f = \{f_e\}_{e \in E}$. A flow f is *feasible* if it satisfies $\sum_{P \in \mathcal{P}_i} f_P = r_i$ for each source-destination pair, where r_i is the mass of traffic traveling from origin o_i to destination d_i.

Each edge $e \in E$ in the network is endowed with a non-negative, non-decreasing *latency* function $\ell_e : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ that maps the mass of traffic on an edge to the delay

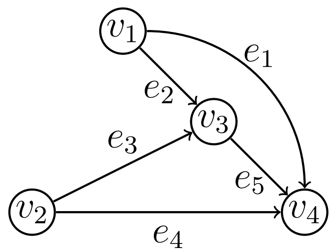

e_1	$\ell_1(f_1) = 4f_1^2$
e_2	$\ell_2(f_2) = 1/2$
e_3	$\ell_3(f_3) = 1/2$
e_4	$\ell_4(f_4) = 2f_4$
e_5	$\ell_5(f_5) = 1/2$

Figure 5.1: An example network routing problem G with two origin-destination pairs: $(o_1, d_1) = (v_1, v_4)$ with $r_1 = 1/2$, and $(o_2, d_2) = (v_2, v_4)$ with $r_2 = 1/2$.

users on that edge observe. The system cost of a flow f is the *total latency*,

$$\mathcal{L}(f) = \sum_{e \in E} f_e \cdot \ell_e(f_e). \tag{5.1}$$

A *routing problem* is specified by the tuple $G = (V, E, \{\ell_e\}_{e\in E}, \{r_i, (o_i, d_i)\}_{i=1}^k)$ as illustrated in Fig. 5.1, and we let $\mathcal{F}(G)$ denote the set of all feasible flows. We define the optimal flow f^{opt} as one that minimizes the total latency, i.e.,

$$f^{\text{opt}} \in \arg\min_{f \in \mathcal{F}(G)} \mathcal{L}(f). \tag{5.2}$$

We denote a family of routing problems by $\mathcal{G}$. A family of routing problems is any set of routing problems, often specified by a specific network topology (e.g., parallel networks) and/or edge latency function types (e.g., polynomial latency functions) but can also be a singleton.

5.2.2 Incentives

In this paper, we consider the problem of selfish routing, where each user in the system chooses a path as to minimize their own observed delay. Let N_i be the set of users traveling from origin o_i to destination d_i. Each non-atomic user $x \in N_i$ is thus free to choose between paths $P \in \mathcal{P}_i$. Let each N_i be a closed interval with Lebesgue measure $\mu(N_i) = r_i$ that is disjoint from each other set of users, i.e., $N_i \cap N_j = \emptyset \ \forall \ i, j \in \{1, \ldots, k\}, \ i \neq j$. The full set of agents is thus $N = \cup_{i=1}^k N_i$ whose mass is $\mu(N) = \sum_{i=1}^k r_i$.

It is well known that selfish routing can lead to sub-optimal system performance [132]. It is therefore up to a system designer to select a set of *incentive functions* $\tau_e : \mathbb{R}_{\geq 0} \to \mathbb{R} \ \forall e \in E$ to influence the behavior of the users in the system to more closely align with the system optimal flow. These incentives can be regarded as monetary transfers with the users dependent on the paths they choose.

A user $x \in N_i$ traveling on a path $P_x \in \mathcal{P}_i$ observes cost

$$J_x(P_x, f) = \sum_{e \in P_x} \ell_e(f_e) + \tau_e(f_e). \tag{5.3}$$

A flow f is a *Nash flow* if

$$J_x(P_x, f) = \min_{P \in \mathcal{P}_i} \left\{ \sum_{e \in P} \ell_e(f_e) + \tau_e(f_e) \right\} \qquad \forall x \in N_i, \ i \in \{1, \ldots, k\}. \tag{5.4}$$

A game is therefore characterized by a routing problem G and a set of incentive functions $\{\tau_e\}_{e \in E}$, denoted by the tuple $(G, \{\tau_e\}_{e \in E})$. It is shown in [140] that a Nash flow exists in a congestion game of this form if the latency and incentive functions are Lebesgue-integrable.

5.2.3 Incentive Mechanisms & Performance Metrics

To determine the manner in which incentive functions are applied to edges, we investigate *incentive mechanisms*. To formalize this notion, let

$$L(G) := \{(\ell_e, e, G)\}_{e \in E(G)}$$

be the set of identifiers for each link or edge in the routing problem G. Further, for a family of problems, let $L(\mathcal{G}) = \cup_{G \in \mathcal{G}} L(G)$ be the set of links that occur in the family of games $\mathcal{G}$. An element in $L(\mathcal{G})$ is a tuple of the latency function ℓ_e, edge index e, and routing problem G which it exists in; with this information, a specific edge can be identified to which incentives can be assigned.

For each edge e in the routing problem G with latency function ℓ_e, an incentive mechanism T assigns an incentive $T(\ell_e; e, G)$, i.e. $\tau_e(f_e) = T(\ell_e; e, G)[f_e]$, where $T(\ell_e; e, G)[f_e]$ is the incentive evaluated at f_e. This mapping is denoted by $T : L(\mathcal{G}) \to \mathcal{T}$ where $\mathcal{T}$ is some set of allowable incentive functions. For brevity, an incentive mechanism will be written simply as $T(\ell_e)$, but it is assumed that, unless otherwise stated, the incentive designer has knowledge of the exact edge and full network structure when assigning an incentive $T(\ell_e)$; these are termed *network-aware* incentive mechanisms [125, 71], and are the focus of Theorem 5.2.1 and Theorem 5.3.3.

In the case where the incentive mechanism must be designed for a family of routing problems and without knowledge of the full network structure, we add the implied constraint that two edges with the same latency function are indistinguishable and must have the same assigned incentive; we highlight such cases by terming the mechanism *network-agnostic*. The use of network-agnostic incentive mechanisms has been studied in [136, 7]; these incentives are useful because of their robustness in settings with frequent changes to the system structure (i.e., commerce, supply-chain-management, and even traffic when considering accidents and emergencies), where partial changes to the network structure or edge latencies need not require global redesign of the incentive mechanism. One such incentive that fits this framework is the classic Pigouvian or marginal cost tax,

$$T^{\text{mc}}(\ell)[f] = f \cdot \frac{d}{df}\ell(f), \tag{5.5}$$

which is known to incentivize users to route optimally in many classes of congestion games [33]. This is only true however, when there is no bound on the incentive and users are homogeneous [136].

We use the *price of anarchy* to evaluate the performance of a taxation mechanism, defined as the worst case ratio between total latency in a Nash flow and an optimal flow, exemplified in Fig. 5.1. Let $\mathcal{L}^{\text{nf}}(G,T)$ be the highest total latency in a Nash flow of the game $(G, T(L(G)))$. Additionally, let $\mathcal{L}^{\text{opt}}(G)$ be the total latency under the optimal flow f^{opt}. The inefficiency can be characterized by

$$\text{PoA}(G,T) = \frac{\mathcal{L}^{\text{nf}}(G,T)}{\mathcal{L}^{\text{opt}}(G)}. \tag{5.6}$$

We extend this definition to a family of instances

$$\text{PoA}(\mathcal{G},T) = \sup_{G\in\mathcal{G}} \frac{\mathcal{L}^{\text{nf}}(G,T)}{\mathcal{L}^{\text{opt}}(G)}, \tag{5.7}$$

where T is used in each routing problem. The price of anarchy is now the worst case inefficiency over all such routing problems while using incentive mechanism T. The objective of such incentive mechanisms is to minimize this worst case inefficiency, thus the optimal incentive mechanism is defined as,

$$T^{\text{opt}} \in \underset{T:L(\mathcal{G})\to\mathcal{T}}{\arg\inf}\ \text{PoA}(\mathcal{G},T), \tag{5.8}$$

such that it minimizes the price of anarchy for a class of routing problems $\mathcal{G}$.

5.2.4 Tolls & Subsidies

We differentiate between two forms of incentives, *tolls* $\tau_e^+ : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ and *subsidies* $\tau_e^- : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\leq 0}$. With tolls, the player's observed cost is strictly increased, i.e., the system designer levies taxes for the users to pay depending on their choice of edges. With subsidies, the players cost is strictly reduced, i.e., the system designer offers some payments to users for their choice of action. The main focus of this work is to assess which is more effective in influencing user behavior, tolls or subsidies.

A *tolling mechanism* is one which only assigns tolling functions, defined as $T^+ : L(\mathcal{G}) \to \mathcal{T}^+$ where $\mathcal{T}^+$ is the set of all non-negative, integrable functions on $\mathbb{R}^+$. An optimal tolling mechanism is one that minimizes the price of anarchy ratio, i.e.,

$$T^{\text{opt}+} \in \operatorname*{arg\,inf}_{T:L(\mathcal{G})\to\mathcal{T}^+} \text{PoA}(\mathcal{G}, T). \tag{5.9}$$

An optimal *subsidy mechanism* is defined analogously with non-positive subsidy functions. In the following sections, we compare the price of anarchy ratio associated with the optimal toll and optimal subsidy.

The following example, illustrated in Fig. 5.1, highlights the notation and the difference between tolls and subsidies.

Example 1. Consider the network G in Fig. 5.1 with two origin destination pairs: (o_1, d_1) $= (v_1, v_4)$ with $r_1 = 1/2$, and $(o_2, d_2) = (v_2, v_4)$ with $r_2 = 1/2$. The optimal flow in G, that minimizes (5.1), is $f^{\text{opt}} \approx \{0.289, 0.211, 0.25, 0.25, 0.461\}$ with a total latency of $\mathcal{L}(f^{\text{opt}}) \approx 0.683$. With no tolling, the Nash flow is $f^{\text{Ne}} = \{1/2, 0, 0, 1/2, 0\}$ with total latency $\mathcal{L}(f^{\text{Ne}}) = 1$ producing a price of anarchy of $\text{PoA}(G, \emptyset) \approx 1.465$. Under a scaled marginal-cost toll, the cost incurred by a user for utilizing edge e is $\ell_e(f_e) + f_e \cdot \frac{d}{df_e}\ell_e(f_e)$ and the Nash flow becomes the same as f^{opt}, leading to a price of anarchy of $\text{PoA}(G, T^{\text{mc}}) = 1$. Similarly, under a subsidy mechanism $T^{-}(\ell_e) = \frac{1}{3} f_e \cdot \frac{d}{df_e}\ell_e(f_e) - \frac{2}{3}\ell_e$, the Nash flow is again the optimal, and $\text{PoA}(G, T^{-}) = 1$.

This example highlights that subsidies and tolls are both effective at reducing the inefficiencies associated with selfish routing. In this work, we study how this performance changes under budgetary constraints and user price-heterogeneity.

5.2.5 Summary of Our Contributions

We start by addressing the nominal homogeneous setting, in which all users react to incentives identically. In Theorem 5.2.1, in any congestion game, we show that under a similar budgetary constraint, the optimal subsidy offers better performance than the optimal toll; the magnitude of this difference is exemplified in Proposition 5.2.1 by deriving explicit price of anarchy bounds for optimal tolls and subsidies in affine congestion games.

Next, we look at the efficacy of each incentive in mitigating the effect of user heterogeneity as the budgetary constraint is lifted. In Theorem 5.3.1, we show that tolls can effectively eliminate the effect of user heterogeneity when the bound on incentives is lifted. However, in Theorem 5.3.2, it is shown that even in congestion games with convex, non-decreasing, continuously-differentiable latency functions, it is impossible for subsidies to mitigate the effect of user heterogeneity, even with the ability to give arbitrarily large payments.

When budgetary constraints do exist and users are heterogeneous in their response to incentives, we show in Theorem 5.3.3 that, for tolls and subsidies bounded to give similar performance in the homogeneous setting, the performance of subsides is worse than tolls when users become heterogeneous, i.e., the performance of subsidies degrades more significantly from player heterogeneity than tolls. This is exemplified in Proposition 5.3.1, giving price of anarchy bounds for robust incentives in affine congestion games.

Finally, because subsides offer better performance under similar budgetary constraints in the homogeneous setting and tolls offer better robustness in the face of user heterogeneity, we investigate what level of user heterogeneity allows tolls to outperform similarly bounded subsides. In Theorem 5.4.1, a relationship between the incentive bound and level of heterogeneity is derived in a class of parrallel-affine congestion games that leads to similar performance guarantees of the optimal toll and subsidy.

5.2.6 Bounded Incentives

We first look at the case where users are homogeneous in their response to incentives. This setting has been the focus of study for many incentive related works [33, 7, 71, 133, 134]. For these reasons, we start by comparing the effectiveness of subsidies and tolls in this setting when additional budgetary constraints are added. Subsidies and tolls both serve as mechanisms for influencing user behavior and can be implemented by similar methods. The act of applying constraints on either is of little difference to the system designer, however the reasoning for these constraints may differ. For instance, budgetary constraints on subsidies can serve to limit the monetary obligation of the system operator, while bounding tolls can prevent scenarios where users may avoid using the network entirely. Though any specific budgetary constraint on either incentive is heavily influenced by the problem setting, here we seek to understand more generally how limits on the magnitude of incentives comparatively affect subsidies and tolls.

To explore this, we introduce bounded tolls and subsidies. A bounded toll satisfies $\tau_e^+(f_e) \in [0, \beta \cdot \ell_e(f_e)]$ for $f_e \geq 0$ and each $e \in E$, where β is a bounding factor. A bounded tolling mechanism is denoted by $T^+(\ell_e; \beta)$. Similarly, a bounded subsidy satisfies

$\tau_e^-(f_e) \in [-\beta \cdot \ell_e(f_e), 0]$ for $f_e \geq 0$ and each $e \in E$, and a bounded subsidy mechanism is denoted by $T^-(\ell_e; \beta)$. This form of bounded incentive functions resembles the bounded path deviations studied in [120]. Though many forms of bounding constraint can be considered, this form is chosen as it can be applied to network-aware and-network agnostic incentive mechanisms, captures the idea that larger delays can be incentivized more significantly, and avoids trivialities caused by arbitrarily large delays. Additionally, these constraints can be represented as the total incentive in a routing problem being within a multiplicative factor β of the total latency, i.e., $\sum_{e \in E} f_e \tau_e(f_e) \leq \beta \mathcal{L}(f)$.

For some bounding factor β, let $\mathcal{T}_\beta^+$ denote the set of taxation mechanisms appropriately bounded by β. More formally, $\mathcal{T}_\beta^+ = \{T | T : L(\mathcal{G}) \to \mathcal{T}^+(\beta)\}$, where

$$\mathcal{T}^+(\beta) = \{\tau_e^+ \in \mathcal{T}^+ \mid \tau_e^+(f_e) \in [0, \beta \cdot \ell_e(f_e)] \; \forall f_e \geq 0\}$$

is the set of all tolling functions bounded by β. To compare the efficacy of bounded tolls and subsidies, we define an optimal bounded tolling mechanism as

$$T^{\text{opt}+}(\beta) \in \operatorname*{arg\,inf}_{T^+ \in \mathcal{T}_\beta^+} \text{PoA}(\mathcal{G}, T^+). \tag{5.10}$$

The optimal bounded subsidy mechanism $T^{\text{opt}-}(\beta)$ is defined analogously. For brevity, bounded mechanisms are often written $T(\beta)$ when being discussed without reference to their use on a specific edge and $T(\ell_e; \beta)$ when they are referenced to a specific edge latency function.

Though we consider any toll bound $\beta \geq 0$, we offer the following definition to differentiate from cases where the bound is very large or trivially zero.

Definition 4. *A toll (subsidy) is* tightly bounded *if $\tau(f) = \beta \ell(f)$, (if $\tau(f) = -\beta \ell(f)$) for some $f \geq 0$.*

When an optimal incentive is tightly bounded, the budgetary constraint is active.

5.2.7 General Relation of Performance

We first consider the relationship between bounded subsidies and tolls in general for congestion games (i.e., arbitrary latency functions and network topologies). Theorem 5.2.1 states that bounded subsidies outperform similarly bounded tolls with respect to the price of anarchy, and strictly outperform when the budgetary constraint is active.

Theorem 5.2.1. *For a congestion games G, under a bounding factor $\beta \geq 0$ the optimal subsidy mechanism $T^{\text{opt}-}(\beta)$ has no greater price of anarchy than the optimal tolling mechanism $T^{\text{opt}+}(\beta)$, i.e.,*

$$\text{PoA}\left(G, T^{\text{opt}+}(\beta)\right) \geq \text{PoA}\left(G, T^{\text{opt}-}(\beta)\right) \geq 1. \tag{5.11}$$

Additionally, if every optimal subsidy is tightly bounded[1]*, then the first inequality in* (5.11) *is strict.*

The proof of Theorem 5.2.1 appears at the end of this section; we first discuss the implications of this result. Theorem 5.2.1 implies that when limiting the size of monetary transactions with homogeneous users, subsidies are more effective than tolls at influencing user behavior. This result holds for any congestion game. Though (5.11) need not be strict in general, there does exist a gap between the performance of tolls and subsidies in many non-trivial settings. To illustrate this, we offer the following example to highlight that bounded subsidies may strictly outperform bounded tolls and outline the proof structure.

Example 2. *Polynomial Congestion Network.* Consider a congestion game, depicted in

[1] $T^{\text{opt}-}(\ell_e; \beta)$ satisfies Definition 4 with bounding factor β for each $\ell_e \in L(\mathcal{G})$.

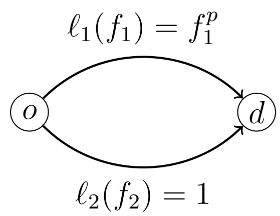

Figure 5.2: Two link parallel congestion game. One edge possesses a polynomial latency function, the other a constant latency function. This routing problem realizes the worst case price of anarchy for polynomial congestion games [141].

Fig. 5.2, possessing two nodes forming a source destination pair with unit mass of traffic and two parallel edges between them, one with latency function $\ell_1(f_1) = f_1^p$, where p is a positive integer, and the other $\ell_2(f_2) = 1$. This example has been shown to demonstrate the worst case inefficiency among polynomial congestion games [141].

Step 1: Identify an Optimal Incentive. When users are homogeneous in their sensitivity to incentives, an optimal toll for this class of games is the marginal cost toll in (5.5), proven to incentivize optimal routing [33]. Notice that the marginal-cost toll will manifest in this network as

$$\tau_1^{\mathrm{mc}}(f_1) = p f_1^p, \quad \tau_2^{\mathrm{mc}}(f_2) = 0, \tag{5.12}$$

and indeed incentivize the Nash flow to be the system optimal of $f_1 = 1/\sqrt[p]{p+1}$.

Step 2: Find Incentives with similar performance. It can be shown that any incentive mechanism in the set

$$\left\{T(\ell) = \lambda T^{\mathrm{mc}}(\ell) + (\lambda - 1)\ell \mid \lambda > 0\right\}, \tag{5.13}$$

has the same performance as the marginal cost taxation mechanism. This observation can be proven from the later Lemma 2.

Step 3: Identify Bounded Subsidies and Tolls. For a bounding factor $\beta \geq p$ the marginal cost taxation mechanism gives a price of anarchy of one; however, for $\beta \in [0, p)$, there exists no taxation mechanism in the set defined in (5.13) which possesses all optimal mechanisms. The similar subsidy mechanism

$$T^{-}(\ell) = (1/(p+1) - 1)\ell + (1/(p+1))T^{\mathrm{mc}}(\ell), \tag{5.14}$$

which manifests in the network as

$$\tau_1^{-}(f_1) = 0, \quad \tau_2^{-}(f_2) = \tfrac{-p}{p+1}, \tag{5.15}$$

is in the set of optimal incentive mechanisms and is valid under bounding factors $\beta \geq \frac{p}{p+1}$. Thus, for bounding factors $\beta \in [\frac{p}{p+1}, p)$, there exists a subsidy mechanism that gives price of anarchy one, but there does not exist a tolling mechanism that does the same. For other bounding factors, the same principles can be followed. In Section 5.2.8, the magnitude of the difference of performance between subsidies and tolls is further explored in the context of affine congestion games.

Having concluded Example 2, in Lemma 2 we show a transformation on incentive mechanisms that does not affect the price of anarchy under homogeneous user sensitivities. This transformation gives us the important relationship between incentive mechanisms that their performance is not unique and similar performance can be garnered with different magnitudes of transactions.

Lemma 2. *Let $T : L(\mathcal{G}) \to \mathcal{T}$ be an incentive mechanism over the family of congestion games $\mathcal{G}$. If another influencing mechanism is defined as $T_\lambda(\ell_e) = \lambda T(\ell_e) + (\lambda - 1)\ell_e$ for any $\lambda > 0$, then*

$$\mathrm{PoA}(\mathcal{G}, T) = \mathrm{PoA}(\mathcal{G}, T_\lambda). \tag{5.16}$$

The proof of Lemma 2 appears in Appendix B.1.

Proof of Theorem 5.2.1: First, observe that if $\beta = 0$ the only permissible incentive function for tolls and subsidies is $\tau_e^+(f_e) = \tau_e^-(f_e) = 0$, i.e., there is no incentive. Therefore, the left and right hand side of (5.11) equate to the unincentivized case and (5.11) holds with equality.

Let $j_e(f_e) = \ell_e(f_e) + \tau_e(f_e)$ denote the cost a player observes for utilizing an edge e when a mass of f_e users are utilizing it. The observed cost of a player $x \in N$ can be rewritten as $J_x(P_x, f) = \sum_{e \in P_x} j_e(f_e)$. In the case where $\beta > 0$, a bounded tolling function on an edge must exist between $\tau_e^+(f_e) \in [0, \beta \cdot \ell(f_e)]$, and the edges observed cost satisfies $j_e^+(f_e) \in [\ell_e(f_e), (1+\beta) \cdot \ell_e(f_e)]$. Similarly, a subsidy function on an edge must exist between $\tau_e^-(f_e) \in [-\beta \cdot \ell(f_e), 0]$, and the edges observed cost satisfies $j_e^-(f_e) \in [(1-\beta) \cdot \ell_e(f_e), \ell_e(f_e)]$.

Let $T^+(\ell_e; \beta)$ be a bounded tolling mechanism with edge costs of $j_e^+(f_e)$. Now, define $T_\lambda(\ell_e) = \lambda T^+(\ell_e; \beta) + (\lambda - 1)\ell_e$; from Lemma 2, T^+ and T_λ have the same price of anarchy for any $\lambda > 0$. Let $\hat{j}_e$ be the edge cost under influencing mechanism T_λ, from the construction of T_λ

$$\hat{j}_e = \ell_e + T_\lambda(\ell_e) = \ell_e + \lambda T^+(\ell_e; \beta) + (\lambda - 1)\ell_e = \lambda j_e^+. \tag{5.17}$$

We now look at the cases where $\beta \in (0,1)$ and $\beta \geq 1$ respectively. When $\beta \in (0,1)$, let $\lambda = (1-\beta)$. Now,

$$\begin{aligned}\hat{j}_e(f_e) = (1-\beta) j_e^+(f_e) &\in [(1-\beta)\ell_e(f_e), (1-\beta^2)\ell_e(f_e)] \\ &\subset [(1-\beta)\ell_e(f_e), \ell_e(f_e)],\end{aligned}$$

thus the edge costs are sufficiently bounded such that T_λ is a permissible subsidy mechanism bounded by β with the same price of anarchy as T^+. If $\beta \geq 1$ let $\lambda = 1/(1+\beta)$

and get

$$\begin{aligned}\hat{j}_e(f_e) = \frac{1}{(1+\beta)} j_e^{+}(f_e) &\in \left[\frac{1}{(1+\beta)}\ell_e(f_e), \ell_e(f_e)\right] \\ &\subset [(1-\beta)\ell_e(f_e), \ell_e(f_e)],\end{aligned}$$

and again T_λ is a permissible subsidy mechanism bounded by β. By letting $T^{+} = T^{\text{opt}+}$ we obtain (5.11).

We have proven that, for $\beta > 0$, if $\text{PoA}(\mathcal{G}, T^{\text{opt}-}(\beta)) = \text{PoA}(\mathcal{G}, T^{\text{opt}+}(\beta))$, then there exists a $T^{\text{opt}-}(\beta)$ that does not achieve the bound. The contrapositive of this is that if every optimal subsidy achieves the bound, the price of anarchy guarantees are not equal. In this case, the optimal subsidies are each tightly bounded and $\text{PoA}(\mathcal{G}, T^{\text{opt}-}(\beta)) < \text{PoA}(\mathcal{G}, T^{\text{opt}+}(\beta))$, proving the final part of Theorem 5.2.1. □

5.2.8 Bounded Incentives in Affine Congestion Games

In Proposition 5.2.1, we explicitly give the price of anarchy bounds of optimal bounded tolls and subsidies in affine congestion games with homogeneous users, again demonstrating the strictly superior performance of subsidies as well as illustrating the magnitude of this difference in performance. Observe that the optimal subsidy outperforms the optimal toll for each incentive bound, matching the results from Theorem 5.2.1.

As a means of illustrating Theorem 5.2.1, we look at the well studied class of affine congestion games, denoted by

$$\mathcal{G}^{\text{aff}} := \{G | \ell_e(f_e) = a_e f_e + b_e, a_e, b_e \geq 0, \forall e \in E(G)\}.$$

We include this result to highlight the appreciable gap in performance between subsidies and tolls in this setting.

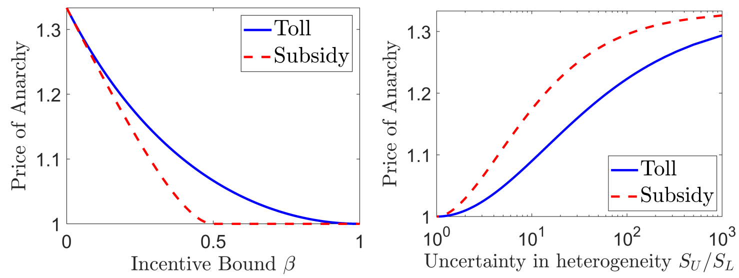

Figure 5.3: Price of Anarchy bounds for comparable tolls and subsidies in affine congestion games. (Left) Price of Anarchy under optimal toll and subsidy respectively bounded by a factor β from Proposition 5.2.1. (Right) Price of Anarchy of a nominally equivalent toll and subsidy with heterogeneity of user sensitivity introduced from Proposition 5.3.1; $S_{\mathrm{U}}/S_{\mathrm{L}}$ expresses the amount of possible heterogeneity in the population.

Proposition 5.2.1. *The optimal bounded network-agnostic tolling mechanism in $\mathcal{G}^{\mathrm{aff}}$ is*

$$T^{\mathrm{opt}+}(af+b;\beta) = \begin{cases} \beta a x & \beta \in [0,1), \\ a x & \beta \geq 1, \end{cases} \tag{5.18}$$

with a price of anarchy bound of

$$\mathrm{PoA}(\mathcal{G}^{\mathrm{aff}}, T^{\mathrm{opt}+}(\beta)) = \begin{cases} \frac{4}{3+2\beta-\beta^2} & \beta \in [0,1), \\ 1 & \beta \geq 1. \end{cases} \tag{5.19}$$

Additionally, the optimal bounded network-agnostic subsidy mechanism in $\mathcal{G}^{\mathrm{aff}}$ is

$$T^{\mathrm{opt}-}(af+b;\beta) = \begin{cases} -\beta b & \beta \in [0,1/2), \\ -b/2 & \beta \geq 1/2, \end{cases} \tag{5.20}$$

with a price of anarchy bound of

$$\mathrm{PoA}(\mathcal{G}^{\mathrm{aff}}, T^{\mathrm{opt}-}(\beta)) = \begin{cases} \frac{4}{3+2\hat{\beta}-\hat{\beta}^2} & \beta \in [0, 1/2), \\ 1 & \beta \geq 1/2, \end{cases} \tag{5.21}$$

where $\hat{\beta} = 1/(1-\beta) - 1$. *Accordingly, for any* $\beta \in (0, 1)$,

$$\mathrm{PoA}(\mathcal{G}^{\mathrm{aff}}, T^{\mathrm{opt}+}(\beta)) > \mathrm{PoA}(\mathcal{G}^{\mathrm{aff}}, T^{\mathrm{opt}-}(\beta)). \tag{5.22}$$

The proof of Proposition 5.2.1 appears in Appendix B.1. Fig. 5.3 (left) illustrates the price of anarchy for tolls and subsidies respectively over various incentive bounds. Though this result is only for a specific class of games, it helps to quantify the broader notion of Theorem 5.2.1: *when users are homogeneous in their response to incentives, a subsidy can consistently give price of anarchy closer to one and often by a significant margin.* In the following sections, we further inspect this relationship when user heterogeneity is introduced.

5.3 Incentives with Heterogeneity

Section 5.2.6 showed that, when users are homogeneous in their response to incentives, subsidies offer better performance guarantees than tolls under budgetary constraints. We now seek to understand how each type of incentive performs when users differ in their price sensitivity.

Specifically, each user $x \in N$ is associated with a sensitivity $s_x > 0$ to incentives. We call $s : N \to \mathbb{R}_{>0}$ a *sensitivity distribution*. We highlight the case where $s_x = c \ \forall x \in N$ for some known constant c as a *homogeneous* distribution of user sensitivities[2], in which

[2]Without loss of generality, we use $s_x = 1$ for a homogeneous population, as was the case in Sec-

each user behaves similarly; any other distribution is referred to as a population of *heterogeneous* users.

A user $x \in N_i$ traveling on a path $P_x \in \mathcal{P}_i$ observes cost

$$J_x(P_x, f) = \sum_{e \in P_x} \ell_e(f_e) + s_x \tau_e(f_e). \tag{5.23}$$

A flow f is a Nash flow if

$$J_x(P_x, f) = \min_{P \in \mathcal{P}_i} \left\{ \sum_{e \in P} \ell_e(f_e) + s_x \tau_e(f_e) \right\} \qquad \forall x \in N_i,\ i \in \{1, \ldots, k\}. \tag{5.24}$$

A game is now denoted by the tuple $(G, s, \{\tau_e\}_{e \in E})$.

To quantify the robustness of an incentive mechanism, we also consider that the system designer may be unaware of users' response to incentives. We denote a set of sensitivity distributions by $\mathcal{S} = \{s : N \to [S_{\mathrm{L}}, S_{\mathrm{U}}]\}$, where $S_{\mathrm{L}} > 0$ is a lower bound on users' sensitivity to incentives and $S_{\mathrm{U}} \geq S_{\mathrm{L}}$ is an upper bound; we include these bounds to quantify the range of users responses, signifying the amount of possible user heterogeneity.

We extend the prior definition of the price of anarchy to include the heterogeneity of users. Let $\mathcal{L}^{\mathrm{nf}}(G, s, T)$ be the highest total latency in a Nash flow of the game $(G, s, T(L(G)))$. Now we define,

$$\mathrm{PoA}(\mathcal{G}, \mathcal{S}, T) = \sup_{G \in \mathcal{G}} \sup_{s \in \mathcal{S}} \frac{\mathcal{L}^{\mathrm{nf}}(G, s, T)}{\mathcal{L}^{\mathrm{opt}}(G)}, \tag{5.25}$$

where the price of anarchy ratio is now the worst case inefficiency over all routing problem,

tion 5.2.6

sensitivity distribution pairs using the incentive mechanism T.

To illustrate this notation, we revisit Example 1, also depicted in Fig. 5.1, but now with user heterogeneity.

Example 3. In the routing problem G, depicted in Fig. 5.1, consider the user sensitivity distribution $s = \{s_x = 2 \forall x \in N_1,\ s_x = 1/2 \forall x \in N_2\}$. As a reminder, the optimal flow in G is $f^{\text{opt}} \approx \{0.289, 0.211, 0.25, 0.25, 0.461\}$ with a total latency of $\mathcal{L}(f^{\text{opt}}) \approx 0.683$, and with no tolling, the Nash flow is $f^{\text{Ne}} = \{1/2, 0, 0, 1/2, 0\}$ with total latency $\mathcal{L}(f^{\text{Ne}}) = 1$ producing a price of anarchy of $\text{PoA}(G, s, \emptyset) \approx 1.465$. With a marginal cost toll T^{mc} as defined in (5.5), the Nash flow becomes $f^{\text{Ne}} \approx \{0.224, 0.276, 0.167, 0.333, 0.443\}$ producing a price of anarchy of $\text{PoA}(G, s, T^{\text{mc}}) \approx 1.04$. With a subsidy mechanism $T^-(\ell_e) = \frac{1}{3} f_e \cdot \frac{d}{df_e} \ell_e(f_e) - \frac{2}{3} \ell_e$ as defined in (5.14) with $p = 2$, the Nash flow becomes $f^{\text{Ne}} \approx \{0, 0.5, 0.137, 0.363, 0.637\}$ producing a price of anarchy of $\text{PoA}(G, s, T^{\text{sub}}) \approx 1.32$.

This example shows that user heterogeneity can have a notable impact on the effectiveness of incentives and can affect their relative performance. In the remainder of this paper, we consider the setting where users are heterogeneous in their price sensitivity when discussing the relative performance of subsidies and tolls. We start by looking at tolls and subsidies independently and investigate their performance in the limit of allowable incentives, i.e., as the budgetary constraint is lifted, how does each type of incentive fare?

In Theorem 5.3.1 we look at the performance of tolls first and find that, when the budgetary constraint is lifted, tolls can eliminate the negative effect of user heterogeneity.

Theorem 5.3.1. *For a class of congestion games $\mathcal{G}$, let $T^* \in \arg\inf_T \text{PoA}(\mathcal{G}, T)$ be an optimal incentive mechanism for homogeneous populations, then*

$$\lim_{\beta \to \infty} \inf_{T^+ \in \mathcal{T}^+_\beta} \text{PoA}(\mathcal{G}, \mathcal{S}, T^+) = \text{PoA}(\mathcal{G}, T^*). \tag{5.26}$$

Furthermore, if $\overline{\mathcal{G}}$ is any class of non-atomic congestion games that has convex, non-decreasing, and continuously differentiable latency functions, then

$$\lim_{\beta\to\infty} \inf_{T^+\in\mathcal{T}^+_\beta} \mathrm{PoA}(\overline{\mathcal{G}}, \mathcal{S}, T^+) = 1. \tag{5.27}$$

The proof of Theorem 5.3.1 appears in Appendix B.1[3]. The proof of Theorem 5.3.1 follows closely from Lemma 3 and the notion of responsiveness to heterogeneity presented in the following section. The result follows from the idea that larger incentives are less impacted by user heterogeneity.

After observing positive results for the use of tolls with user heterogeneity, we next seek to understand the effectiveness of subsidies in the same situation. In Theorem 5.3.2, we show that, even in a restricted class of congestion games, subsidies cannot effectively mitigate the effect of player heterogeneity in the same way tolls do.

Theorem 5.3.2. *Let $\overline{\mathcal{G}}$ be any class of non-atomic congestion games that has convex, non-decreasing, and continuously differentiable latency functions, the set of latency functions is closed under nonnegative scalar multiplication, and has at least one network where the untolled price of anarchy is greater than one. There exists no network-agnostic subsidy mechanism T that gives price of anarchy of 1, i.e.,*

$$\lim_{\beta\to\infty} \inf_{T^-\in\mathcal{T}^-_\beta} \mathrm{PoA}(\overline{\mathcal{G}}, \mathcal{S}, T^-) > 1. \tag{5.28}$$

[3]This result is reminiscent of [33] stating that there exist tolls that influence optimal selfish routing in some settings. In this paper, we extend the result from [33] to cases where users are heterogeneous and classes of games where a price of anarchy of one may not be achievable. We note that Theorem 5.3.1 is more general than of [136, Theorem 1], as this result is given for general incentives and is not reliant on marginal cost taxes nor is it limited to the family of congestion games in which they are optimal. Further, the results of [125] cover the case in which the system designer is fully aware of the users' price sensitivities (or value of time in their case) and applies fixed tolls. In contrast, in this paper the toll designer is unaware of the users' exact price sensitivities but is still able to provide a flow-varying tolling scheme that gives a price of anarchy of one as the bounding constraint is lifted.

The proof of Theorem 5.3.2 appears in Appendix B.1.

Though the class of routing problems has a more strict definition than in Theorem 5.3.1, the result is still very general and holds for most cases other than singleton networks and those where the price of anarchy is always 1. From Theorem 5.3.1 and Theorem 5.3.2 we conclude that *without the presence of budgetary constraints, tolls can mitigate the effect of player heterogeneity while subsidies cannot.* However, this relationship was shown only as the budgetary constraint was lifted; in the next section, we further investigate the effect of user heterogeneity on subsidies and tolls while budgetary constraints on the incentives remain.

5.3.1 Robustness of Incentives

In Section 5.3, user heterogeneity was discussed in the sense of whether incentives could or could not fully mitigate the effect of non-uniform user behavior. In many cases the very large incentives needed to completely eliminate the negative effects of user heterogeneity are not possible, particularly in the presence of budgetary constraints. It is thus of interest what the performance guarantees are when the effects of user heterogeneity cannot be entirely overcome and how this compares when using subsidies or tolls.

To compare the robustness of bounded tolls and subsidies, we define an optimal bounded tolling mechanism as

$$T^{\mathrm{opt}+}(\beta, \mathcal{S}) \in \operatorname*{arg\,inf}_{T^+ \in \mathcal{T}^+_\beta} \mathrm{PoA}(\mathcal{G}, \mathcal{S}, T^+). \tag{5.29}$$

The optimal bounded subsidy mechanism $T^{\mathrm{opt}-}(\beta, \mathcal{S})$ is defined analogously. For notational convenience, we will omit the dependence on $\mathcal{S}$ in the homogeneous setting.

Often, increased user heterogeneity causes performance of an incentive mechanism to diminish. We give the following definition for classes of congestion games with this property.

Definition 5. *A class of congestion games is* responsive to player heterogeneity *if* $\mathrm{PoA}(\mathcal{G},\mathcal{S},T^*)$ *is strictly increasing with* $S_{\mathrm{U}}/S_{\mathrm{L}} > 1$ *for an optimal bounded incentive mechanism* $T^* \in \arg\inf_T \mathrm{PoA}(\mathcal{G},\mathcal{S},T)$.

These classes of games are those that have a degradation in performance from increased player heterogeneity, even while the optimal incentive mechanism is in use; many classes of well studied congestion games possess this property [136].

5.3.2 General Relation of Robustness

In Theorem 5.3.3, we give a robustness result that shows the performance of subsidies degrades more quickly than tolls as player heterogeneity is introduced.

Theorem 5.3.3. *For a class of congestion games* $\mathcal{G}$*, define two incentive bounds* β^+ *and* β^- *such that*

$$\mathrm{PoA}\left(\mathcal{G},T^{\mathrm{opt}-}(\beta^-)\right) = \mathrm{PoA}(\mathcal{G},T^{\mathrm{opt}+}(\beta^+)\big), \tag{5.30}$$

then at the introduction of player heterogeneity,

$$\mathrm{PoA}\left(\mathcal{G},\mathcal{S},T^{\mathrm{opt}-}(\beta^-,\mathcal{S})\right) \geq \mathrm{PoA}(\mathcal{G},\mathcal{S},T^{\mathrm{opt}+}(\beta^+,\mathcal{S})\big) \geq 1. \tag{5.31}$$

Additionally, each inequality in (5.31) *is strict if* $\mathcal{G}$ *is responsive to player heterogeneity and* $S_{\mathrm{L}} < S_{\mathrm{U}}$.

Intuitively, this result stems from the fact that subsidies are more finely tuned to give performance guarantees, as guaranteed in Theorem 5.2.1. Essentially, applying a small, negative incentive to an edge's increasing latency function will have a more significant impact on the shape of the users' cost function than a larger, positive toll. This fact causes the same amount of player heterogeneity to have a larger effect on Nash flows caused by subsidies than with an equivalent toll. Thus, when increased player heterogeneity escalates the inefficiency, this relationship is strict. Though the relationship isn't strict for general classes of congestion games, it is for many well studied cases, including the aforementioned polynomial congestion games.

We show in Lemma 3 a relation between nominally equivalent incentives in the heterogeneous population setting; specifically, we show that the heterogeneous price of anarchy decreases as incentives increase costs to the users.

Lemma 3. *For a class of congestion games $\mathcal{G}$, let T be an incentive mechanism. If $T_\lambda(\ell) = (\lambda-1)\ell+\lambda T$, then $\mathrm{PoA}(\mathcal{G},\mathcal{S},T_\lambda)$ is non-increasing with λ and strictly decreasing if $\mathcal{G}$ is responsive to user heterogeneity and $S_\mathrm{L} < S_\mathrm{U}$.*

The proof of Lemma 3 appears in the appendix.

Proof of Theorem 5.3.3: First, we give the following definition for incentives that have the same performance in the homogeneous setting.

Definition 6. *For any incentive mechanism T and $\lambda > 0$, each incentive mechanism satisfying $T_\lambda(\ell_e) = (\lambda-1)\ell_e + \lambda T(\ell_e)$ is termed nominally equivalent. From Lemma 2, nominally equivalent incentives satisfy*

$$\mathrm{PoA}(\mathcal{G},T) = \mathrm{PoA}(\mathcal{G},T_\lambda). \tag{5.32}$$

The theorem follows closely from Lemma 2 and Lemma 3. First, suppose $T^{\mathrm{opt}+}(\beta^+)$ is an optimal tolling mechanism bounded by β^+. From Lemma 2 there exists a nominally equivalent subsidy T_λ^-. If $T_\lambda^- \notin \mathcal{T}_{\beta^-}^-$, then there must exist a $T_\lambda^+ \in \mathcal{T}_{\beta^+}^+$ that is nominally equivalent to $T^{\mathrm{opt}-}(\beta^-)$ from the monotonicity and invertability of the transformation in

Lemma 2. From (5.30), this implies there exists a nominally equivalent $T^{\text{opt}+}(\beta^+)$ and $T^{\text{opt}-}(\beta^-)$.

Now, let $T^{\text{opt}-}(\beta^-, \mathcal{S})$ be the optimal subsidy with player heterogeneity bounded by β^-. From the fact before, we know there exists a toll T^+ that is nominally equivalent to $T^{\text{opt}-}(\beta^-, \mathcal{S})$ and bounded by β^+. From Lemma 3, we obtain that

$$\text{PoA}(\mathcal{G}, \mathcal{S}, T^+) \leq \text{PoA}(\mathcal{G}, \mathcal{S}, T^{\text{opt}-}(\beta^-, \mathcal{S})), \tag{5.33}$$

and by the definition of $T^{\text{opt}+}(\beta^+, \mathcal{S})$, we get

$$\text{PoA}(\mathcal{G}, \mathcal{S}, T^{\text{opt}+}(\beta^+, \mathcal{S})) \leq \text{PoA}(\mathcal{G}, \mathcal{S}, T^+). \tag{5.34}$$

Combining (5.33) and (5.34) gives (5.31). If the class of games is responsive to player heterogeneity, then $\text{PoA}(\mathcal{G}, \mathcal{S}, T_\lambda)$ is strictly decreasing with λ and the relationship is strict. □

5.3.3 Robustness of Incentives in Affine Congestion Games

Theorem 5.3.3 states that the performance of subsidies degrades more quickly than tolls when users differ in their response to incentives. Further, if a subsidy and a toll perform the same in the homogeneous setting, the subsidy performs worse than the toll with any level of user heterogeneity. To illustrate this fact, we again look at the class of affine congestion games. In this section, we specifically look at $\mathcal{G}^{\text{pa}}$, defined as the class of parallel-network affine-latency congestion games in which each edge has positive traffic in the untolled Nash flow. We assign taxes using the *optimal scaled marginal cost toll with player heterogeneity*, $T^{\text{smc}}(af + b) := (\sqrt{S_{\text{L}} S_{\text{U}}})^{-1} af$. This tolling mechanism was first introduced in [142], and was shown to minimize the price of anarchy in parallel affine congestion games with sensitivity distributions in $\mathcal{S}$ bounded by $S_{\rm L}$ and $S_{\rm U}$. In Proposition 5.3.1, we give price of anarchy bounds on the optimal scaled marginal cost toll as well as a nominally equivalent subsidy $T^{\rm nes}$.

Proposition 5.3.1. *Let $\mathcal{G}^{\rm pa}$ be the set of fully-utilized parallel affine congestion games with sensitivity distributions in $\mathcal{S}$. The optimal scaled marginal cost tolling mechanism is $T^{\rm smc}(af+b) = \frac{af}{\sqrt{S_{\rm L}S_{\rm U}}}$ with price of anarchy*

$$\mathrm{PoA}(\mathcal{G}^{\rm pa}, \mathcal{S}, T^{\rm smc}) = \frac{4}{3}\left(1 - \frac{\sqrt{q}}{(1+\sqrt{q})^2}\right). \tag{5.35}$$

where $q := S_{\rm L}/S_{\rm U}$. Additionally, a nominally equivalent subsidy is $T^{\rm nes}(af+b) = -\frac{1}{1+\sqrt{S_{\rm L}S_{\rm U}}}b$, with price of anarchy

$$\mathrm{PoA}(\mathcal{G}^{\rm pa}, \mathcal{S}, T^{\rm nes}) = \frac{4}{3}\left(1 - \frac{\sqrt{\hat{q}}}{(1+\sqrt{\hat{q}})^2}\right), \tag{5.36}$$

where

$$\hat{q} = \frac{\lambda q}{1 - q + \lambda q} < q,$$

and $\lambda = \sqrt{S_{\rm L}S_{\rm U}}/(1+\sqrt{S_{\rm L}S_{\rm U}})$.

The proof of Proposition 5.3.1 appears in the appendix. Observe that, because $\hat{q} < q$ in (5.35) and (5.36) the nominally equivalent subsidy has greater price of anarchy when player heterogeneity is introduced. This can be seen in Fig. 5.3 (right). Intuitively, the same amount of player heterogeneity has a larger effect on the subsidy than the toll.

5.4 Bounded & Robust Incentives

In the previous sections, it was shown that when users are homogeneous in their response to incentives, subsidies offer better performance guarantees than tolls under similar budgetary constraints; however, as users become heterogeneous in their response to incentives, the performance of subsidies degrades more quickly than that of tolls. The logical next question we address is, how much heterogeneity causes bounded tolls to outperform bounded subsidies? In general, this question is difficult to answer. We therefore look at the case of affine congestion games on parallel networks while using network-agnostic affine incentive functions. In Theorem 5.4.1, we find the incentive bound β^* that causes the price of anarchy of the optimal bounded toll and subsidy with user heterogeneity to be equal. Without loss of generality (because we assume $S_{\rm L}$ and $S_{\rm U}$ are known to the system designer), we normalize to $S_{\rm L}S_{\rm U} = 1$.

Theorem 5.4.1. *Let $T^{\rm opt+}(\beta, \mathcal{S})$ and $T^{\rm opt-}(\beta, \mathcal{S})$ be an optimal, affine toll and subsidy mechanism for $\mathcal{G}^{\rm pa}$ with incentive bound β and player sensitivities between $S_{\rm L}$ and $S_{\rm U}$. An incentive bound of $\beta^* = 1/S_{\rm U} = S_{\rm L}$ gives*

$$\text{PoA}(\mathcal{G}^{\rm pa}, \mathcal{S}, T^{\rm opt-}(\beta^*, \mathcal{S})) = \text{PoA}(\mathcal{G}^{\rm pa}, \mathcal{S}, T^{\rm opt+}(\beta^*, \mathcal{S})). \tag{5.37}$$

As illustrated in Fig. 5.4, for lower levels of user heterogeneity (i.e., $\beta^* < 1/S_{\rm U}$), the optimal subsidy offers price of anarchy closer to one than the optimal toll. When there is a larger amount of user heterogeneity (i.e., $\beta^* > 1/S_{\rm U}$) the optimal toll has a lower price of anarchy bound than the optimal subsidy.

The proof of Theorem 5.4.1 appears at the end of this section and is supported by the following two propositions. Proposition 5.4.1 (originally introduced in [136]) gives the optimal affine tolling mechanism and the accompanying price of anarchy guarantee.

Proposition 5.4.1. (Brown & Marden [136]) *Let $T^+(k_1, k_2)$ denote an affine taxation mechanism that assigns tolling functions $\tau_e^+(f_e) = k_1 a_e f_e + k_2 b_e$. For any $\beta > 0$, the*

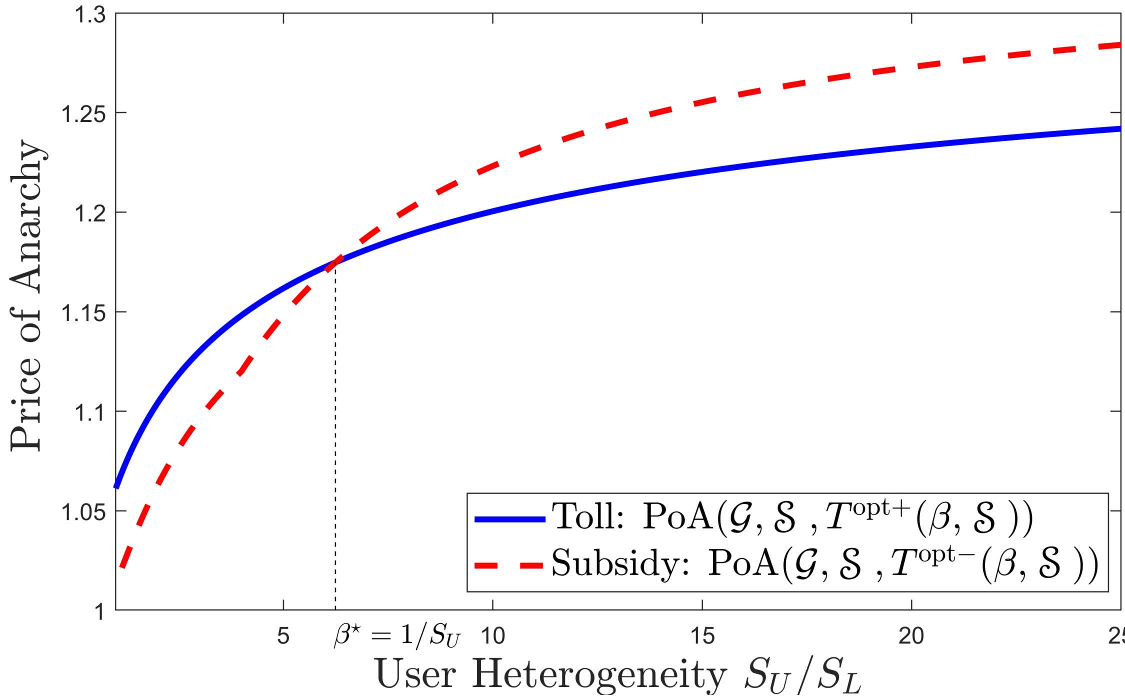

Figure 5.4: Price of anarchy under optimal bounded tolls and subsidies with heterogeneous users in parallel-affine congestion games with $\beta = 0.4$. When the amount of user heterogeneity is low (i.e. $S_{\mathrm{U}}/S_{\mathrm{L}}$ close to one), subsidies offer better performance guarantees than tolls as stated in Theorem 5.2.1; however, as the level of heterogeneity increases, the performance of subsidies degrade more quickly than tolls, stated in Theorem 5.3.3. When the incentive bound is $\beta = 1/S_{\mathrm{U}}$ the performance of subsidies and tolls is equal, as stated in Theorem 5.4.1.

optimal coefficients k_1^* *and* k_2^* *satisfying*

$$(k_1^*, k_2^*) \in \operatorname*{arg\,min}_{0 \le k_1, k_2 \le \beta} \mathrm{PoA}\left(\mathcal{G}^{\mathrm{pa}}, \mathcal{S}, T^+(k_1, k_2)\right), \tag{5.38}$$

are given by

$$k_1^* = \beta, \tag{5.39}$$

$$k_2^* = \max\left\{0, \frac{\beta^2 S_{\mathrm{L}} S_{\mathrm{U}} - 1}{S_{\mathrm{L}} + S_{\mathrm{U}} + 2\beta S_{\mathrm{L}} S_{\mathrm{U}}}\right\}. \tag{5.40}$$

Furthermore, for any $G \in \mathcal{G}^{\mathrm{pa}}$, $\mathrm{PoA}(G, \mathcal{S}, T^+(k_1^*, k_2^*))$ *is upper bounded by the following* *expression:*

$$\frac{4}{3}\left(1-\frac{\beta S_{\mathrm{L}}}{(1+\beta S_{\mathrm{L}})^2}\right) \quad \text{if} \quad \beta<\frac{1}{\sqrt{S_{\mathrm{L}}S_{\mathrm{U}}}} \tag{5.41}$$

$$\frac{4}{3}\left(1-\frac{(1+\beta S_{\mathrm{L}})(\frac{S_{\mathrm{L}}}{S_{\mathrm{U}}}+\beta S_{\mathrm{L}})}{(1+2\beta S_{\mathrm{L}}+\frac{S_{\mathrm{L}}}{S_{\mathrm{U}}})^2}\right) \quad \text{if} \quad \beta\geq\frac{1}{\sqrt{S_{\mathrm{L}}S_{\mathrm{U}}}}. \tag{5.42}$$

The proof of Proposition 5.4.1 appears in Appendix B.1. The price of anarchy bound is shown in Fig. 5.4. Similarly, in Proposition 5.4.2 the optimal affine subsidy is given along with its price of anarchy guarantee.

Proposition 5.4.2. *Let $T^-(k_1,k_2)$ denote an affine subsidy mechanism that assigns subsidy functions $\tau_e^-(f_e)=k_1a_ef_e+k_2b_e$. For any $\beta>0$, the optimal coefficients k_1^* and k_2^* satisfying*

$$(k_1^*,k_2^*)\in \underset{-\beta\leq k_1,k_2\leq 0}{\arg\min}\ \mathrm{PoA}\left(\mathcal{G}^{\mathrm{pa}},\mathcal{S},T^-(k_1,k_2)\right), \tag{5.43}$$

are given by

$$k_1^*=0, \tag{5.44}$$

$$k_2^*=-\min\left\{\beta,\frac{1}{S_{\mathrm{L}}+S_{\mathrm{U}}}\right\}. \tag{5.45}$$

Furthermore, for any $G\in\mathcal{G}^{\mathrm{pa}}$, $\mathrm{PoA}(G,\mathcal{S},T^-(k_1^*,k_2^*))$ is upper bounded by the following expression:

$$\frac{4}{3}(1-\beta S_{\mathrm{L}}(1-\beta S_{\mathrm{L}})) \quad \text{if} \quad \beta<\frac{1}{S_{\mathrm{L}}+S_{\mathrm{U}}} \tag{5.46}$$

$$\frac{4}{3}\left(1-\frac{S_{\mathrm{L}}/S_{\mathrm{U}}}{(1+S_{\mathrm{L}}/S_{\mathrm{U}})^2}\right) \quad \text{if} \quad \beta\geq\frac{1}{S_{\mathrm{L}}+S_{\mathrm{U}}}. \tag{5.47}$$

The proof of Proposition 5.4.2 appears in Appendix B.1. The price of anarchy bound is shown in Fig. 5.4. The price of anarchy bounds equate at $\beta=1/S_{\mathrm{U}}$, as substantiated by Theorem 5.4.1, and for $\beta < 1/S_{\mathrm{U}}$ the subsidy price of anarchy bound is lower, while for $\beta > 1/S_{\mathrm{U}}$ the toll price of anarchy bound is lower and converging to one.

Proof of Theorem 5.4.1: Proposition 5.4.1 and Proposition 5.4.2 give the price of anarchy bounds for the optimal affine incentives. By inspection, when $\beta \in [\frac{1}{S_{\mathrm{L}}+S_{\mathrm{U}}}, \frac{1}{\sqrt{S_{\mathrm{L}}S_{\mathrm{U}}}}]$, the optimal toll and subsidy price of anarchy bounds fall in the domain of (5.41) and (5.46) respectively. Additionally, when $\beta = 1/S_{\mathrm{U}}$, we can see that the optimal toll is $T^{+}(\frac{1}{S_{\mathrm{U}}}, 0)$ and the optimal subsidy is $T^{-}(0, \frac{-1}{S_{\mathrm{L}}+S_{\mathrm{U}}})$; furthermore, these incentives have the same price of anarchy bound, i.e.,

$$\mathrm{PoA}\left(\mathcal{G}^{\mathrm{pa}}, \mathcal{S}, T^{+}\left(\frac{1}{S_{\mathrm{U}}}, 0\right)\right) = \mathrm{PoA}\left(\mathcal{G}^{\mathrm{pa}}, \mathcal{S}, T^{-}\left(0, \frac{-1}{S_{\mathrm{L}}+S_{\mathrm{U}}}\right)\right). \quad (5.48)$$

It is easy to see from (5.41),(5.42),(5.46), and (5.47) that for $\beta > 1/S_{\mathrm{U}}$,

$$\mathrm{PoA}\left(\mathcal{G}^{\mathrm{pa}}, \mathcal{S}, T^{\mathrm{opt}+}\right) < \mathrm{PoA}\left(\mathcal{G}^{\mathrm{pa}}, \mathcal{S}, T^{\mathrm{opt}-}\right),$$

and for $\beta < 1/S_{\mathrm{U}}$,

$$\mathrm{PoA}\left(\mathcal{G}^{\mathrm{pa}}, \mathcal{S}, T^{\mathrm{opt}+}\right) > \mathrm{PoA}\left(\mathcal{G}^{\mathrm{pa}}, \mathcal{S}, T^{\mathrm{opt}-}\right).$$

Therefore, $\beta = 1/S_{\mathrm{U}}$ is the unique incentive bound that gives equal price of anarchy for subsidies and tolls with heterogeneous users in the class of parallel, affine congestion games. □

5.5 Conclusion

In this work, the effectiveness of subsidies and tolls in congestion games were compared in the presence of budgetary constraints on incentives and user heterogeneity. The results of this manuscript show that, in a nominal setting, smaller subsidies offer better performance guarantees than tolls; however, in the face of unknown user heterogeneity, tolls are more robust than subsidies. These results hold for general classes of non-atomic congestion games, and future work will investigate if the main conclusions hold for atomic congestion games as well. Future work may look at more general notions of user sensitivities as well as other realistic emergent behavior for the society of users.

Chapter 6

Information Signalling with Concurrent Monetary Incentives in Bayesian Congestion Games

The uncertainty held by a system's users can cause ineffective decision-making. Nowhere is this more apparent than in transportation networks, where drivers' uncertainty over current road/traffic conditions can negatively alter their routing choices. To alleviate this, an informed system operator may *signal* information to uninformed users to persuade them into taking more preferable actions (e.g., Google/Apple maps providing live traffic updates). In this work, we study public signalling mechanisms in the context of Bayesian congestion games. We observe the phenomenon that though revealing information can reduce system cost in some settings, in others, it can induce worse performance than not signalling at all. However, we find an important relationship between information signalling and monetary incentives: by utilizing both mechanisms concurrently, the system operator can guarantee that revealing information does not worsen performance. We prove these findings in a general class of Bayesian congestion games. To understand the magnitude at which information signalling can affect system performance, we put a deeper focus in the class of parallel networks with polynomial latency functions and analytically characterize bounds on the change in system cost from signalling. Finally, we consider the problem of solving for optimal signals with and without the concurrent use of monetary incentives. We construct solvable optimization problems whose solutions give optimal signalling policies even when the signalling policy is limited in its support; we then quantify the benefit of these and other signalling mechanisms in numerical examples.

6.1 Introduction

The degree of traffic congestion on highways and roads in busy city areas is inherently caused by the collective route choices of the drivers [33]. Though drivers often choose routes that minimize their own travel time, the system behavior that emerges from this decision-making need not be optimal [143]. This inefficiency can be further exacerbated by drivers' uncertainty over the state of the system [144, 87], e.g., uncertainty on current weather conditions, traffic rates, or on-road collisions. With the deployment of new sensing and communication technologies (e.g., vehicle-to-device and vehicle-to-infrastructure), the traffic engineers overseeing these systems gain the opportunity to learn these unknown system parameters; however, the effect of revealing this information to drivers is not well understood.

The emergence of new sensing and communication technologies opens the door to new methods for coordinating driving behavior and improving traffic patterns. One such method is that of *information signalling* by a well-informed central authority [88, 145]. By partially revealing their information about system parameters to uninformed users, the signaller allows the system users the opportunity to form new beliefs about their environment. If the signaller reveals this information strategically, they may alter user behavior in such a way that the overall system performance is improved; for example, Google/Apple maps can share the travel times of certain routes to guide driver decision making in a way that can alter aggregate driving patterns and improve performance for the user population [146]. One may initially think that all information should be shared with the users; however, it has been observed in several problem settings (and affirmed here) that this need not be optimal and could further degrade system performance [147, 148, 87, 149, 150]. The main focus of this work is determining what capabilities a system operator has in improving congestion via information signalling and identifying when and how this information should be shared.

We study the principles of information signalling in the context of *Bayesian congestion games*, where a group of users (syn. drivers) must route themselves through a congestible network while the exact congestion characteristics of each path are unknown. Deterministic congestion games have been used in the transportation literature to model driver decision-making and its effect on road traffic [33, 151, 152, 153, 154, 155, 156, 157, 158, 159]. Recently, Bayesian congestion games have emerged as a generalized model in which the edge latency functions are random variables. The users possess a common prior belief over the random latencies of each route (for example, the belief there was an accident on a road or the chance weather has affected driving conditions) but do not observe the realization. The informed system operator does observe the realized values of the random parameters and can strategically signal information about them to the system users. This model for uncertain driving conditions have been used to study how information signalling policies should be designed [89, 160, 95], what behavior is likely to emerge [161, 162], and the associated performance of specific signalling structures [90, 163, 164]. The results are typically limited to computational methods for finding signalling policies or identifying whether or not revealing the state exactly is optimal. Additionally, for ease of analysis, much of the work in this area often assumes the signals are private (sent to individual users) [149], which does not give relevant insights on public signals [165] (sent to all users) which we consider in this work.

Signalling mechanisms are becoming a topic of increasing research in their ability to influence user behavior; however, this is not the only influencing mechanism at a system operator's disposal. Incentive mechanisms, where users are assessed monetary penalties or rewards based on their actions, have long been studied as an effective means of coordinating system behavior [5, 122, 7]. In transportation settings, these incentives may manifest as road/bridge tolls or transit prices. The interplay between incentives and signalling is an emerging area of study, and has up until now been limited to studying simpler two-route networks [150, 166], consider only the situation where full information is revealed to users or limited models of uncertainty [167, 168, 169], focus on mechanisms where users must pay to acquire information [170, 171, 172], or provide numerical studies rather than theoretical guarantees [173, 174]. To the best of our knowledge, no existing work has analytically studied the relationship between monetary incentives and information signalling in general a model for network routing with stochastic delays.

In this work, we provide insights on the benefit information signalling can provide with and without the concurrent use of monetary incentives. Through example, we demonstrate two key observations: (1) signalling, on its own, can worsen system cost, and (2) co-designing signal and incentive mechanisms offer opportunities for improvement that were not present when using each separately. In Theorem 6.3.1, we formalize the benefit of co-designing these mechanisms: with appropriate monetary incentives, information signalling will not worsen system performance, essentially making signalling robust. To further understand the benefit of utilizing both mechanisms in tandem, we consider the special class of parallel networks with polynomial latency functions in Section 6.4 and derive bounds on the possible benefit a signalling policy can provide with and without concurrent incentives.

Finally, in Section 6.5, we address the problem of finding optimal signal-incentive pairs, including when the possible number of signals is bounded. We show how one can create solvable optimization problems to find the optimal signalling policies with and without concurrent incentives that may or may not be allowed to update with the sent signal.

We bolster the conclusions of this work with numerical examples in Section 6.4.3 and Section 6.5.4, in which we find that concurrent signal-incentive mechanisms offer notable improvements.

6.2 System Model

6.2.1 Congestion Games

Consider a directed graph (V, E) with vertex set V, edge set $E \subseteq (V \times V)$, and k origin-terminal pairs (o_i, t_i). Denote by $\mathcal{P}_i$ the set of all simple paths connecting origin o_i to destination t_i. Further, let $\mathcal{P} = \cup_{i=1}^{k} \mathcal{P}_i$ denote the set of all paths in the graph. A *flow* on the graph is a vector $f \in \mathbb{R}_{\geq 0}^{|\mathcal{P}|}$, where f_P expresses the mass of traffic utilizing path P. The mass of traffic on an edge $e \in E$ is thus $f_e = \sum_{P: e \in P} f_P$, and we say $f = \{f_e\}_{e \in E}$. A flow f is *feasible* if it satisfies $\sum_{P \in \mathcal{P}_i} f_P = r_i$ for each source-destination pair, where r_i is the mass of traffic traveling from origin o_i to terminal t_i.

When a larger number of users traverse the same path, the congestion (and thus transit delay) on that path increases. To characterize this, each edge e is endowed with a *latency function* $\ell_e : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ that maps the mass of traffic on an edge to the delay users on that edge observe. We assume each latency function is positive, convex, non-decreasing, and continuously differentiable. The system cost of a flow f is the *total* *latency*,

$$\mathcal{L}(f) = \sum_{e \in E} f_e \cdot \ell_e(f_e). \tag{6.1}$$

A *routing problem* is specified by the tuple $G = (V, E, \{\ell_e\}_{e\in E}, \{r_i, (o_i, t_i)\}_{i=1}^k)$, and we let $\mathcal{F}(G)$ denote the set of all feasible flows. We define the optimal flow f^{opt} as one that minimizes the total latency, i.e.,

$$f^{\text{opt}} \in \underset{f \in \mathcal{F}(G)}{\arg\min}\, \mathcal{L}(f). \tag{6.2}$$

Though this flow is desirable, it need not emerge from the self-interested decision-making of the users. To model the setting where users are free to choose their own paths (such as drivers selecting their own routes), let $x \in [0, r_i]$ denote the index of an infinitesimal agent who uses a path $P_x \in \mathcal{P}_i$; the flow f_e thus represents the mass of infinitesimal users sharing an edge. Agents that minimize their own observed cost (e.g., their individual travel delay) possess a cost function $J_x(P_x; f) = \sum_{e \in P_x} \ell_e(f_e)$; plausible behavior that can emerge in the system is that of a *Nash flow* f^{Ne} [33, 151, 175], which satisfies

$$J_x(P_x; f) \leq J_x(P'; f),\ \forall P' \in \mathcal{P}_i,\ x \in [0, r_i],\ i \in [k]. \tag{6.3}$$

These system states are those where no user has the incentive to change their action and need not be optimal [35]; additionally, the total latency in any Nash flow in a game of this form is the same [140].

6.2.2 Bayesian Congestion Games & Information Signalling

We consider a setting where the exact traffic conditions are unknown to drivers but known precisely by a central system operator (e.g., Google Maps, Waze, Apple Maps, etc.). Let each latency function take the form $\ell_e(f_e) = \sum_{k=1}^{|\mathcal{K}|} \alpha_{e,k} \cdot \ell_k(f_e)$, where $\mathcal{K} =$

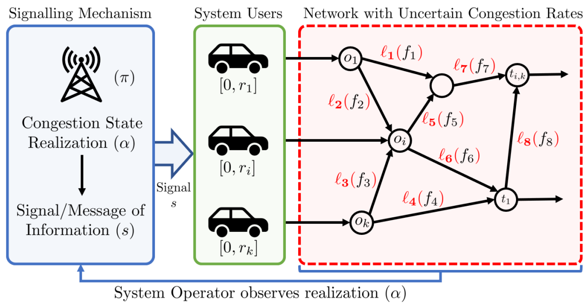

Figure 6.1: System model diagram. System users (drivers) must travel from their source to their destination through a congested network with uncertain congestion rates. The users do not know the current network state; however, the system operator does. Leveraging its greater information, the system operator can devise a signalling mechanism to send messages of partial information to the users to alter their beliefs and, ultimately, their actions.

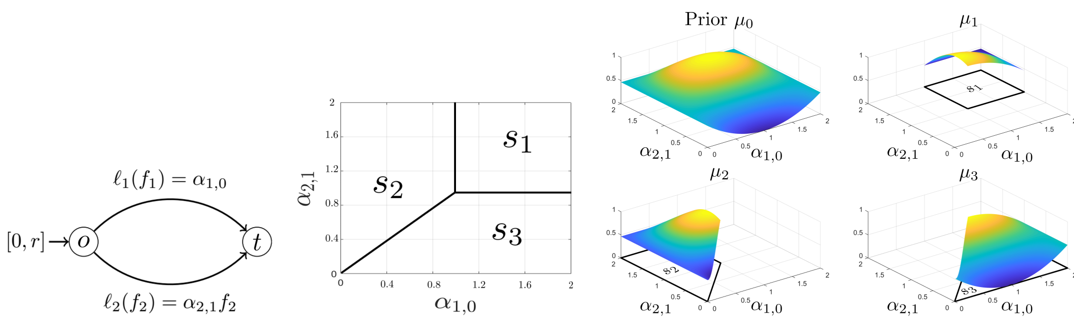

Figure 6.2: Two-link, parallel, Bayesian congestion game. One edge possesses a linear latency function, the other a constant latency function. The coefficients of each of these latency functions $\alpha_{1,0}, \alpha_{2,1}$ are unknown but distributed with prior μ_0 over $A = [0, 2]^2$. At right, is an illustration of a truthful signalling policy $\pi : A \to \{s_1, s_2, s_3\}$, which partitions A to map realizations to signals. After receiving a signal s, the agents compute their posterior μ_s, as illustrated by the posterior beliefs with support defined by the subset of A to which it is associated. In general, signals need not be deterministic/truthful, and we may choose a signalling policy $\pi : A \to \Delta(S)$, such that the signal s is drawn from a distribution conditioned on α; the posterior beliefs are computed the same, but now need not be of the partitioned form shown in this example.

$\{\ell_1, \ldots, \ell_K\}$ is a set of basis latency functions[1] and $\alpha_{e,k} \geq 0$ is the weight of the basis function $\ell_k(\cdot)$ on the edge $e \in E$. To capture the idea of uncertainty in this problem, let

[1]This formulation can capture many models of congestion including polynomial ($\mathcal{K} = \{x^0, x^1, x^2, \ldots, x^K\}$), exponential ($\mathcal{K} = \{e^{0x}, e^{0.5x}, e^{2x}, \ldots\}$), and the Bureau of Public Roads (BPR) latency functions ($\mathcal{K} = \{x^0, x^4\}$), commonly used to model the congestion characteristics of physical roads [176, 177].

$\alpha \in \mathbb{R}_{\geq 0}^{|E|\cdot|\mathcal{K}|}$ (whose elements contain the weight of each basis latency function on each edge) be a random variable with distribution $\mu_0(x) = \mathbb{P}[\alpha = x]$ and support A. We assume the system operator observes the realization of this parameter, but the system users do not. If the users reach a flow f, we extend (6.1) to be the *expected total latency* over a distribution μ,

$$\mathcal{L}(f;\mu) = \mathop{\mathbb{E}}_{\alpha\sim\mu}\left[\sum_{e\in E} f_e \cdot \ell_e(f_e)\right].$$

Because $\ell_e(\cdot)$ is determined by α, it is a random variable.

As a method to coordinate behavior and induce more desirable system states, the system operator may choose to signal relevant information to the users so they may update their beliefs. To do so, a system operator selects a *signalling policy* $\pi : A \to \Delta(S)$ that maps realizations of the system state $\alpha \in A$ to distribution, from which a signal $s \in S$ is sampled[2] which may reveal information to system users. We assume these signals are *public*, in that every user receives the same message but need not be truthful or reveal the exact realization. Fig. 6.1 illustrates the signalling model in the context of network congestion games where only partial information is provided to the users through signal s. At the reception of signal $s \in S$, users infer the posterior distribution over the system state α as

$$\mu_s(\alpha) = \frac{\pi(s|\alpha)\cdot\mu_0(\alpha)}{\int_A \pi(s|\beta)\cdot\mu_0(\beta)d\beta},$$

where $\pi(s|\alpha)$ is the probability of sending signal s when the system state realization is α. Fig. 6.2 demonstrates how agents beliefs may be shaped in a two-link network with two unknown parameters. By utilizing three signals, the system operator can induce three posteriors that differ from the prior and lead to different network flows.

Under a signalling policy π, agents may change their chosen path based on which

[2]For ease of notation, we will often treat the set of signals S as finite; however, the set of signals can be generalized to include a unique signal for each realization of the system state, i.e., $S = A$, which can be uncountable.

signal they receive. Let $\mathbf{f} = \{f(s)\}_{s\in S}$ denote the tuple containing the flow that occurs at the reception of each signal, and let $\sigma_x = \{\sigma_x(s) \in \mathcal{P}_i\}_{s\in S}$ denote the path user $x \in [0, r_i]$ selects after receiving each signal. When each agent adopts a strategy based on the information system's signals, the system designer now cares about the expected total latency, expressed as

$$\mathcal{L}(\mathbf{f}; \mu_0, \pi) = \sum_{s\in S} \psi(s) \cdot \mathcal{L}(f(s); \mu_s), \tag{6.4}$$

where $\psi(y) = \mathbb{P}[s = y] = \int_{\alpha' \in A} \pi(s|\alpha) \cdot \mu_0(\alpha) d\alpha$ denotes the distribution over signals. An agent's cost will now be their expected travel time,

$$J_x(\sigma_x; \mathbf{f}, \mu_0, \pi) = \sum_{s\in S} \psi(s) \cdot \mathop{\mathbb{E}}_{\alpha\sim\mu_s} \left[\sum_{e\in\sigma_x(s)} \ell_e\left(f_e(s)\right) \right]$$

We can now define a *Bayes-Nash Equilibrium* as a tuple $(\mathbf{f}^{\mathrm{BNe}}, \sigma^{\mathrm{BNe}})$ as a set of strategies where no agent elects to unilaterally change, i.e.,

$$J_x(\sigma_x^{\mathrm{BNe}}; \mathbf{f}^{\mathrm{BNe}}, \mu_0, \pi) \leq J_x(\sigma'; \mathbf{f}^{\mathrm{BNe}}, \mu_0, \pi), \\ \forall \sigma' \in (\mathcal{P}_i)^m,\ x \in [0, r_i], i \in \{1, \ldots, k\}. \tag{6.5}$$

Our main focus in this work is understanding what opportunities a system designer has in lowering the expected total latency by way of information provisioning, i.e., comparing $\mathcal{L}(\mathbf{f}^{\mathrm{BNe}}; \mu, \pi)$ with $\mathcal{L}(\mathbf{f}^{\mathrm{BNe}}; \mu, \emptyset)$ (where the use of $\emptyset$ denotes the case where no information is shared with users). To quantify this improvement in system performance, we define the *benefit to system cost* as

$$\mathbf{B}(\pi; \mu) = \mathcal{L}(\mathbf{f}^{\mathrm{BNe}}; \mu, \emptyset) - \mathcal{L}(\mathbf{f}^{\mathrm{BNe}}; \mu, \pi), \tag{6.6}$$

which measures the reduction in system cost from utilizing a signal policy π. The system operator's objective is to institute a signalling structure that reduces the system cost or, equivalently, has a positive benefit. Several works have shown encouraging results on the capabilities of information signalling and identified situations in which system cost can be significantly reduced [90, 161, 160]. However, the consequences of information signalling need not always be positive. In the following example, we identify that this may be the case even in simple settings.

Example 4 (Consequences From Signalling)*.* In this example, consider a population of drivers tasked with selecting one of two commute options. One of the routes is always delayed (either from natural hazards, uncertain demand, or irregular maintenance), while the other has free-flowing traffic but congests as the number of drivers on that route increases. However, the drivers are uncertain which route will be delayed.

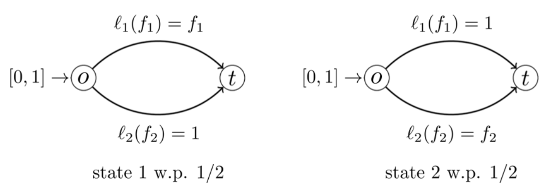

Figure 6.3: Bayesian congestion game where which route has a constant large delay and which has a linearly increasing delay is random.

To model this, consider a congestion game with two parallel edges $E = \{e_1, e_2\}$. One edge has a linear latency function, and the other a constant; each edge is the linear congestible edge with probability 1/2 (as depicted in Example 4). Let $\mu_0(\alpha^1) = 1/2$ be the prior belief that the first edge is the linear congestible edge (state α^1). When no information is revealed, users split over the two edges equally ($f_1 = f_2 = 1/2$), and the expected system cost is $\mathcal{L}(\mathbf{f}^{\mathrm{BNe}}; \mu_0, \emptyset) = 0.75$.

If an information signal s is sent to the users, let $q := \mathbb{P}[\alpha = \alpha^1 \mid s]$ be the posterior belief that the first edge is the linear congestible edge. With this posterior, $f_1 = q$ users utilize the first edge, and the expected cost is $\mathcal{L}(f;\mu_s) = q^2 - q + 1$. For any value of $q \neq 1/2$, the expected system cost is greater than not revealing information; as such, any signalling policy that causes users beliefs to differ from the prior will increase cost, i.e., $\mathbf{B}(\pi;\mu) = \mathcal{L}(\mathbf{f}^{\mathrm{BNe}};\mu,\emptyset) - \mathcal{L}(\mathbf{f}^{\mathrm{BNe}};\mu,\pi) < 0$ for all $\pi : A \to \Delta(S)$. This demonstrates our first observation:

Observation: Revealing information to users can have negative consequences and increase system cost.

Example 4 highlights that signalling, on its own, may not be capable of reducing system cost. However, this is not the only influencing mechanism at a traffic engineer's disposal. Another mechanism to influence user behavior is that of *monetary incentives*, which have been well studied in transportation and alleviating congestion [5, 122, 7], but, to the authors' best knowledge, the use of information signalling and monetary incentives in tandem has yet to be studied in the context of traffic networks.

6.2.3 Monetary Incentives

Consider a congestion game G; an incentive designer can apply an incentive $\tau_e \in \mathbb{R}$ to each edge $e \in E$ to change the cost experienced by users utilizing that edge, i.e.,

$$J_x(e;f) = \ell_e(f_e) + \tau_e.$$

When a signal s is sent to the users, the expected cost to a user x on path P_x in flow f becomes $J_x(P_x;f) = \mathbb{E}\left[\sum_{e\in P_x} \ell_e(f_e) + \tau_e \mid s\right]$. This change in cost affects the users' decision-making and ultimately leads to new Nash flows, ideally with lower total latency. Monetary incentives are a well-studied and highly utilized method of controlling congestion in transportation [155, 156]. However, the relationship between incentives and information signalling is not currently well understood; studying these two mechanisms concurrently is the main focus of this work.

To model the interplay of these two influencing mechanisms, note that at each signal, the selected tolls will alter the Bayes-Nash Equilibrium[3] by propagating this new cost into (6.5). With incentives τ and signalling policy π, the equilibrium system cost will be written $\mathcal{L}(\mathbf{f}^{\mathrm{BNe}}; \mu_0, \pi, \tau)$. One can identify scenarios where either mechanism is capable of reducing congestion; however, this is not true in general. For a given Bayesian congestion game, it is not immediately apparent if either influencing mechanism can independently reduce system cost at all. In the following example, we will see that even in a simple setting, quantifiable benefits exist to designing these mechanisms concurrently.

Example 5 (The Need for Co-Design). In this example, again, consider a population of drivers tasked with choosing between two commutes. The traffic rates on one route are always known, but the second sometimes contains a significant delay (perhaps caused by routine closures and detours).

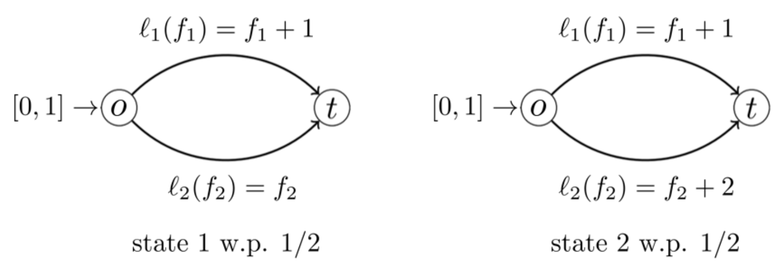

Figure 6.4: Bayesian congestion game where one edge has a constant congestion profile while the other has an additional large delay with probability 1/2.

To model this situation, consider a congestion game with two parallel edges $E = \{e_1, e_2\}$. The first edge has a deterministic latency function $\ell_1(f_1) = f_1 + 1$, while the second edge has a latency function $\ell_2(f_2) = f_2 + \zeta$ where $\zeta = 2$ with probability 1/2 and $\zeta = 0$ otherwise (as depicted in Fig. 6.4). First, we show that no toll can reduce system cost alone. When no information is revealed, each edge has the same expected

[3]Note that the Bayes-Nash equilibrium flow $\mathbf{f}^{\mathrm{BNe}}$ is now inherently dependent on the selected incentives.

cost, and the Bayes-Nash flow is $\mathbf{f}^{\mathrm{BNe}} = \{(1/2, 1/2)\}$; the optimal flow is the same, i.e., $f^{\mathrm{opt}} = (1/2, 1/2)$. As the unincentivized equilibrium is already optimal, clearly, no toll can reduce system cost.

Now, consider some information signalling policy π with signal set $S = \{s_1, \ldots, s_m\}$. Note that $\mathcal{L}^{\mathrm{Ne}}(\zeta) = 1 + \zeta/2$. Using forthcoming tools from this work (i.e., Lemma 4), it can be found that $\mathcal{L}(\mathbf{f}^{\mathrm{BNe}}; \mu_0, \emptyset) = 1 + \mathbb{E}[\zeta]/2$, thus, any signalling policy π does not reduce system cost.

Finally, let π be the full-reveal signalling policy and $\tau_1 = 0.5$ and $\tau_2 = 0$. The expected system cost with this signal/incentive pair is $1.4375 < 1.5 = 1 + \mathbb{E}[\zeta]/2$. This demonstrates our second observation:

Observation: There exist situations where signalling alone or tolling alone cannot provide the same opportunities in reducing system cost as signalling and tolling together.

Example 5 points to an important relationship between signals and incentives: there exist opportunities in designing the two together, but the benefits are not readily obvious. A co-design of the two mechanisms can be accomplished in two ways (1) by creating a larger optimization problem in which signals and incentives are both decision variables (see Section 6.5), and (2) by designing an incentive policy that can update with the sent signal. We call incentives that can update with the signal *signal-aware* and incentives that cannot *signal-agnostic*. This section highlighted the limitations of signal-agnostic incentives; in Section 6.3 and Section 6.4, we will largely focus on the benefit of signal-aware incentives.

6.2.4 Summary of Contributions

The main contributions of this work come in characterizing the interplay of two influencing mechanisms: information signalling and monetary incentives. Further, we describe how these mechanisms can be designed concurrently to provide increased benefits in reducing total latency. We propose two methods for this co-design. The first is utilizing signal-aware incentives designed for a given signalling policy. In Proposition 6.3.1, we characterize the optimal signal-aware toll for any signalling policy. One insight this work provides is that these incentives make the signalling policy robust; in Theorem 6.3.1, we show that while using the optimal signal-aware incentives, no information signalling policy can worsen system cost. To further illustrate the advantage of co-designing signals and incentives, we consider the sub-class of problems with parallel networks and polynomial latency functions in Section 6.4, in which we find analytical bounds for how much a signalling policy can change system cost. The insights from this section follow the more general results and show that signalling can still provide a significant reduction in system cost when incentives are also used.

The second method of co-design involves directly solving for the optimal signal-incentive pairs. To do so, we leverage existing results on Generalized Moment Problems to solve for optimal signalling mechanisms in the aforementioned class of parallel networks with polynomial latency functions. In Section 6.5, we survey the existing literature and show how the optimal signal-incentive pairs (with either signal-agnostic or signal-aware incentives) can be transcribed and solved as GMPs. Additionally, we amend the problem to handle the case where there is a limited number of signals that can be sent (i.e., $|S|$ is bounded).

Finally, in Section 6.5.4, we offer a numerical simulation to quantify the above results of signalling and incentive mechanisms concurrent use. This experiment demonstrates several of the insights from this work, including that co-designed incentive mechanisms offer notable performance improvements.

6.3 Advantage of Incentives

Example 4 and Example 5 highlighted an opportunity to design signals and incentives in tandem. In this section, we will focus on the qualities of incentives that can update with the sent signal. Consider a *signal-aware incentive mechanism* $T(s;\pi,\mu_0)$ that assigns tolls $\{\tau_e(s)\}_{s\in S}$ dependent on the signal broadcast by the information provider. A player $x \in [0, r_i]$ with the strategy σ_x now observes an expected cost of

$$J_x(\sigma_x;\mathbf{f},\mu_0,\pi,T) = \sum_{s\in S}\psi(s)\cdot \mathop{\mathbb{E}}_{\alpha\sim\mu_s}\left[\sum_{e\in\sigma_x(s)}\ell_e\left(f_e(s)\right)+\tau_e(s)\right].$$

The Bayes-Nash flow definition remains as shown in (6.5), but now with users' tolled cost. We now seek to understand the effectiveness of jointly implementing a signalling policy π and an incentive mechanism T. As such, we extend the definition of (6.6), which quantifies the gain in system performance to include the effect of an incentive mechanism T, i.e.,

$$\mathbf{B}(\pi;\mu,T) = \mathcal{L}(\mathbf{f}^{\mathrm{BNe}};\mu,\emptyset,T) - \mathcal{L}(\mathbf{f}^{\mathrm{BNe}};\mu,\pi,T). \tag{6.7}$$

We measure the benefit of a signalling policy by comparing the system cost with incentives and signalling and incentives alone. We do this because we largely want to focus on the value that information signalling can provide on its own.

First, we must decide what monetary incentives to use. In Proposition 6.3.1, we characterize an optimal signal-aware incentive mechanism for a given signalling policy.

Proposition 6.3.1. *Let μ_0 be a prior on the latency coefficients α in a Bayesian congestion game G with positive, convex, non-decreasing, and continuously differentiable latency functions that are of the form $\ell_e(f_e) = \sum_{k=1}^{|\mathcal{K}|}\alpha_{e,k}\cdot\ell_k(f_e)$ where $\ell_k\in\mathcal{K}$, and let*

$\pi : A \to \Delta(S)$ be a signalling policy. An optimal signal-aware incentive mechanism $T^\star$ (i.e., maximizes $\mathcal{L}(\mathbf{f}^{\mathrm{BNe}}; \mu_0, \pi, T)$) assigns tolls according to

$$\tau_e^\star(s) = \sum_{k=1}^{|\mathcal{K}|} \mathop{\mathbb{E}}_{\alpha_{e,k} \sim \mu_s} [\alpha_{e,k}] \xi_e \ell_k'(x_e), \tag{6.8}$$

where $\xi \in \arg\min_{f \in \mathcal{F}(G)} \mathcal{L}(f; \mathbb{E}_{\alpha \sim \mu_s}[\alpha])$.

The proof appears in Appendix A. Proposition 6.3.1 provides a mechanism for computing the optimal incentives for any signalling policy π. The use of these incentives in tandem with a signalling policy will alter the equilibrium flow and thus the system cost. Motivated by the observed negative consequences of information signalling shown in Example 4, Theorem 6.3.1 finds that the concurrent use of monetary incentives $T^\star$ makes it such that signalling can never worsen system performance, i.e., have no negative benefit. Under the use of any signaling policy π, at the reception of any signal $s \in \mathcal{S}$, $T^\star$ incentivizes the network flow that minimizes the posterior expected cost. We show that the total latency in an optimal flow is concave in α and apply Jensen's inequality to show the expected posterior total latency is no greater than the expected prior total latency.

Theorem 6.3.1. *Let μ_0 be a prior on the latency coefficients α in a Bayesian congestion game G with positive, convex, non-decreasing, and continuously differentiable latency functions that are of the form $\ell_e(f_e) = \sum_{k=1}^{|\mathcal{K}|} \alpha_{e,k} \cdot \ell_k(f_e)$ where $\ell_k \in \mathcal{K}$. While using the signal-aware incentive policy $T^\star$ (as defined in Proposition 6.3.1, any signalling policy $\pi : A \to \Delta(S)$ has non-negative benefit to system cost, i.e.,*

$$\mathbf{B}(\pi; \mu_0, T^\star) \geq 0. \tag{6.9}$$

Proof of Theorem 6.3.1: Consider a realization of a congestion game G with latency coefficients α. Let $\mathcal{L}^\star(\alpha)$ denote the total latency in a Nash flow while using the incentive mechanism $T^\star$ as defined in Proposition 6.3.1. First, we characterize the Bayesian-Nash flow with incentives $T^\star$. If the signal $s \in S$ is sent to users, they update their belief via Bayesian inference to $\mu_s(\alpha) = \frac{\pi(s|\alpha)\cdot\mu_0(\alpha)}{\psi(s)}$. In a flow f, user $x \in [0, r_i]$ taking path $P_x \in \mathcal{P}_i$ experiences an expected cost of

$$\begin{aligned} J_x(P_x; f, \mu_s) &= \mathop{\mathbb{E}}_{\alpha\sim\mu_s}\left[\sum_{e\in P_x}\sum_{k=1}^{|\mathcal{K}|}\alpha_{e,k}\cdot\ell_k(f_e) + \tau_e^\star(s)\right] \\ &= \sum_{e\in P_x}\sum_{k=1}^{|\mathcal{K}|}\mathbb{E}[\alpha_{e,k}|s]\ell_k(f_e) + \tau_e^\star(s). \end{aligned}$$

Note that if f were not a Nash flow in the congestion game with coefficients $\mathbb{E}[\alpha|s]$, then by (6.3) there would exist a user x would be able to deviate their strategy $\sigma_x(s)$ and experience lower cost. Therefore, the only Bayes-Nash flows occur when $f(s)$ is a Nash flow with respect to $\mathbb{E}[\alpha|s]$ and tolls $\tau_e^\star(s)$ for all $s \in S$. From Proposition 6.3.1, this is the optimal flow in the network with coefficients $\mathbb{E}[\alpha|s]$.

Next, consider the prior distribution μ_0 on α, and let f be a flow in the network. The expected total latency

$$\begin{aligned} \mathcal{L}(f;\mu) &= \mathop{\mathbb{E}}_{\alpha\sim\mu}\left[\sum_{e\in E} f_e\cdot\sum_{k=1}^{|\mathcal{K}|}\alpha_{e,k}\ell_k(f_e)\right] \\ &= \sum_{e\in E} f_e\cdot\sum_{k=1}^{|\mathcal{K}|}\mathop{\mathbb{E}}_{\alpha\sim\mu}[\alpha_{e,k}]\,\ell_k(f_e) \\ &= \mathcal{L}(f; \mathop{\mathbb{E}}_{\alpha\sim\mu}[\alpha]), \end{aligned}$$

which follows from the linearity of expected value.

Combining the previous two observations, we obtain that the total latency in a Nash flow in the congestion game G with latency coefficients α when using $T^\star$, can be expressed

as

$$\mathcal{L}(\mathbf{f}^{\mathrm{BNe}};\mu_0,\pi,T^\star) = \sum_{s\in S}\psi(s)\mathcal{L}^\star(\mathbb{E}[\alpha|s]). \tag{6.10}$$

Next, we observe that $\mathcal{L}^\star(\alpha)$ is concave. $\mathcal{L}^\star(\alpha)$ can be expressed as the pointwise infimum over $f\in\mathcal{F}(G)$ for a given α,

$$\mathcal{L}^\star(\alpha) = \inf_{f\in\mathcal{F}(G)}\mathcal{L}(f;\alpha) = \inf_{f\in\mathcal{F}(G)}\sum_{e\in E}\sum_{k=1}^{|\mathcal{K}|}\alpha_{e,k}f_e\cdot\ell_k(f_e).$$

Observing that $\sum_{e\in E}\sum_{k=1}^{|\mathcal{K}|}\alpha_{e,k}f_e\ell_k(f_e)$ is affine in α (and thus concave in α) for each f, we can invoke that the pointwise infimum over a class of functions that are each concave is itself, concave [178]. Thus $\mathcal{L}^\star(\alpha)$ is concave though need not be affine.

Now, consider the total latency in a Bayes-Nash flow with signal policy π and incentive $T^\star$,

$$\begin{aligned}
\mathcal{L}(\mathbf{f}^{\mathrm{BNe}};\mu_0,\pi,T^\star) &= \sum_{s\in S}\psi(s)\cdot\mathcal{L}^\star(\mathbb{E}[\alpha|s]) && (6.11a)\\
&\le \mathcal{L}^\star\left(\sum_{s\in S}\psi(s)\cdot\mathbb{E}[\alpha|s]\right) && (6.11b)\\
&= \mathcal{L}^\star(\mathbb{E}[\alpha]) = \mathcal{L}(\mathbf{f}^{\mathrm{BNe}};\mu_0,\emptyset,T^\star), && (6.11c)
\end{aligned}$$

where the (6.11a) holds from (6.10) and (6.11b) holds from the concavity of $\mathcal{L}^\star$. From this, we can see that $\mathbf{B}(\pi;\mu_0,T^\star)\ge 0$, i.e., while utilizing incentive scheme $T^\star$, signalling cannot increase system cost. □

6.4 Polynomial Routing Games on Parallel Networks

In this section, we seek to further understand the connection between signalling and incentivizing by characterizing closed-form bounds on the benefit a signalling policy can provide with and without incentives. We do this in the context of Bayesian congestion games on parallel networks (i.e., one source-terminal pair with n directed edges directly connecting them) with polynomial latency functions (i.e., $\ell_k(\cdot)$ is a monomial). For the remainder of this work consider that the basis latency set $\mathcal{K}$ is determined by the set of monomial basis functions with degrees $\mathcal{D} = \{d_1, \dots, d_{|\mathcal{K}|}\}$ where $d_i \in \mathbb{Z}_{\geq 0}$ expresses the degree of basis latency function ℓ_i[4], e.g., $\mathcal{D} = \{0, 1\}$ represents affine congestion rates, $\mathcal{D} = \{0, 4\}$ can represent the well-known Bureau of Public Roads (BPR) latency functions, commonly used to model the congestion characteristics of physical roads [176, 177], and $\mathcal{D} = \{0, \dots, D\}$ can represent any positive, convex, increasing polynomial up to degree D [7]. Without loss of generality, we will index α by the polynomial degree $d \in \mathcal{D}$. We assume these latency functions are positive, increasing, convex polynomials; additionally, we note from Remark 2 in Appendix B that, without loss of generality, we can normalize to a unit demand, i.e., $r = 1$.

Though this class is restricted relative to the general class of congestion networks, our findings illustrate insights on the effects of signalling that can be observed only more dramatically more broadly and is a new step in generality from many similar works. Additionally, the proofs of Theorem 6.4.1 and Theorem 6.4.2 develop new tools leveraging the geometry and gradient of the system cost function, which can in future work be applied to more general classes of problems as well as more specific case studies.

We highlight two important possible realizations of the random variable α that will be used throughout: $\check{\alpha} \in \mathbb{R}_{\geq 0}^{|E|\cdot|\mathcal{D}|}$ such that $\check{\alpha}_{e,d} = \inf\{\text{supp}(\alpha_{e,d})\}$ (where $\text{supp}(\cdot)$ denotes the support of α) in which each parameter takes its lowest value, and $\hat{\alpha} \in \mathbb{R}_{\geq 0}^{|E|\cdot|\mathcal{D}|}$ such that $\hat{\alpha}_{e,d} = \sup\{\text{supp}(\alpha_{e,d})\}$ in which each parameter takes its largest value. Note that $\check{\alpha}$ and $\hat{\alpha}$ need not be in the support of α, but rather represent the corners of the smallest box that contains the support of α that are closest and furthest from the origin respectively.

[4]We assume that $0 \in \mathcal{D}$ is always satisfied.

Further, to avoid degenerate cases, we institute the following assumption on Bayesian congestion games.

Assumption 1. *In a Bayesian Congestion Game G with prior μ_0, $0, 1 \in \mathcal{D}$ is always satisfied, and, for each edge $e \in E$, $\check{\alpha}_{e,0}, \check{\alpha}_{e,1} > 0$.*

This assumption prevents cases where traffic can be routed with zero delay and has zero effect on congestion.

6.4.1 Signalling Alone

When a system designer seeks to improve system performance by solely using a public information-signalling system, Theorem 6.4.1 provides bounds on the benefit a signalling policy can provide.

Theorem 6.4.1. *Consider the class of parallel Bayesian congestion games with polynomial latency functions whose degrees come from the set $\mathcal{D}$. For any distribution over the latency coefficients μ_0 and any signalling policy π, the benefit in the expected total latency of a Bayes-Nash flow from signalling satisfies*

$$-\Theta\|\mathbb{E}[\alpha] - \check{\alpha}\|_2 \leq \mathbf{B}(\pi; \mu_0) \leq \Theta\|\mathbb{E}[\alpha] - \check{\alpha}\|_2, \tag{6.12}$$

where $\Theta := |\mathcal{D}| + \frac{\rho^+ - \rho_0^-}{2\rho_1^-}(|E| + |\mathcal{D}| - 1)$, $\rho_0^- = \min_{e\in E} \check{\alpha}_{e,0}$, $\rho_1^- = \min_{e\in E} \check{\alpha}_{e,1}$, $\rho^+ = \max_{e\in E} \sum_{d\in\mathcal{D}} (d+1)\hat{\alpha}_{e,d}$, $\mathbb{E}[\alpha] = \int_{z\in A} z\mu_0(z)dz$, and $\check{\alpha} \in \mathbb{R}_{\geq 0}^{|E|\cdot|\mathcal{D}|}$ such that $\check{\alpha}_{e,d} = \inf\{\mathrm{supp}(\alpha_{e,d})\}$ for each $e \in E$, $d \in \mathcal{D}$. Additionally, there exists a μ_0 such that for any truthful $\pi \neq \emptyset$ (i.e., $\pi : A \to S$ is deterministic),

$$\mathbf{B}(\pi; \mu_0) = \sqrt{|\mathcal{D}|} \cdot \|\mathbb{E}[\alpha] - \check{\alpha}\|_2. \tag{6.13}$$

Further, if $d \in \mathcal{D}$ where $d > 0$, then there exists a μ_0 such that for any truthful $\pi \neq \emptyset$,

$$\mathbf{B}(\pi; \mu_0) = -\|\mathbb{E}[\alpha] - \check{\alpha}\|_2. \tag{6.14}$$

The proof of Theorem 6.4.1 appears at the end of this section.

Theorem 6.4.1 reveals the capabilities a signalling policy has in improving system performance. It also highlights the reality that revealing information can make system performance worse. Though the bounds in (6.12) are not tight, the most interesting aspect of these bounds is what parameters they are conditioned on, providing some insights on how the different primitives of a network routing problem affect the efficacy of signalling. The bounds for the benefit of a signalling policy depend on the number of terms considered in each latency function $|\mathcal{D}|$, the size of the network $|E|$, as well as the distance between the average system state and the edge of its support $\|\mathbb{E}[\alpha] - \check{\alpha}\|_2$ and other terms that change with the support. One can think that the number of latency terms $|\mathcal{D}|$ characterizes the complexity of the model of network congestion while $\|\mathbb{E}[\alpha] - \check{\alpha}\|_2$ measures the amount of uncertainty about the system parameters. Additionally, (6.13) and (6.14) show that there exist situations where regardless of what signalling policy is chosen, revealing information can greatly benefit or hinder system performance.

For many of the proofs in this section, we will utilize the following Lemma, demonstrated in the proof of Theorem 6.3.1 and proven in Appendix A.

Lemma 4. *With a prior μ_0 and a signalling policy π, the Bayes-Nash flow $\mathbf{f}^{\mathrm{BNe}}$ can be characterized by $\{\overline{f}^{\mathrm{Ne}}(s)\}_{s \in S}$, where $\overline{f}^{\mathrm{Ne}}(s)$ is the Nash flow in the network G with coefficients $\overline{\alpha}_s = \mathbb{E}[\alpha|s]$. Additionally, for a given flow f, the expected total latency with distribution μ over the coefficients α is equal to the total latency in the network with the expected coefficients, i.e., $\mathcal{L}(f; \mu) = \mathcal{L}(f; \mathbb{E}_{\alpha \sim \mu}[\alpha])$. Together, these facts show that the expected total latency in a Bayes-Nash flow is equal to the weighted average of the total* *latency in the expected network after receiving each signal, i.e.,*

$$\mathcal{L}(\mathbf{f}^{\mathrm{BNe}}; \mu_0, \pi) = \sum_{s \in S} \psi(s) \mathcal{L}^{\mathrm{Ne}}(\mathbb{E}_{\alpha \sim \mu_s}[\alpha]), \tag{6.15}$$

where $\mathcal{L}^{\mathrm{Ne}}(\alpha)$ *denotes the total latency in a Nash flow in the deterministic congestion game* G *with latency coefficients* α.

To prove Theorem 6.4.1, in Lemma 5 we provide a fact about the function $\mathcal{L}^{\mathrm{Ne}}(\alpha)$ which bounds the difference in total latency between any two realizations of edge latency coefficients.

Lemma 5. *Consider the class of parallel congestion games with polynomial latency functions with degrees drawn from the set* $\mathcal{D}$ *with coefficients* $\alpha \in A$. *Let* $a, b \in \mathbb{R}_{\geq 0}^{|E| \cdot |\mathcal{D}|}$ *be two possible sets of coefficients for a congestion game with edge set* E, *then*

$$|\mathcal{D}| + \frac{\rho^+ - \rho_0^-}{2\rho_1^-} \left(|E| + |\mathcal{D}| - 1\right) \geq \frac{\mathcal{L}^{\mathrm{Ne}}(a) - \mathcal{L}^{\mathrm{Ne}}(b)}{||a - b||_2}, \tag{6.16}$$

where $\rho_0^- = \min_{e \in E} \check{\alpha}_{e,0}$, $\rho_1^- = \min_{e \in E} \check{\alpha}_{e,1}$, *and* $\rho^+ = \max_{e \in E} \sum_{d \in \mathcal{D}} (d+1) \hat{\alpha}_{e,d}$.

The proof appears in Appendix B.

Proof of Theorem 6.4.1: We start by proving the lower bound on $\mathbf{B}(\pi; \mu_0)$, which quantifies how much the use of a signalling policy can worsen the system performance. Consider the prior μ_0 and signalling policy $\pi : A \to \Delta(S)$. If μ_s is the posterior formed from receiving signal s, let $\overline{\alpha}_s = \mathbb{E}_{\alpha \sim \mu_s}[\alpha]$.

To prove the lower bound, first, define the set

$$\hat{A} = \prod_{e \in E, d \in \mathcal{D}} [\inf\{\mathrm{supp}(\alpha_{e,d}), \sup\{\mathrm{supp}(\alpha_{e,d})]$$

that is the smallest box in $\mathbb{R}_{\geq 0}^{|E| \cdot |\mathcal{D}|}$ that contains $A = \mathrm{supp}(\alpha)$. Note that $\check{\alpha}$ is the corner of this box that is closest to the origin. Let, $\hat{\mathcal{L}}^{\mathrm{Ne}}$ be the concave closure of the function $\mathcal{L}^{\mathrm{Ne}}$ over $\hat{A}$, i.e.,

$$\hat{\mathcal{L}}^{\mathrm{Ne}} = \sup\left\{z|(\alpha, z) \in \mathrm{Conv}_{\hat{A}}(\mathcal{L}^{\mathrm{Ne}})\right\},$$

where $\mathrm{Conv}_{\hat{A}}(\mathcal{L}^{\mathrm{Ne}})$ denotes the convex hull of the graph of $\mathcal{L}^{\mathrm{Ne}}$ over the domain $\hat{A}$. With this, we show

$$\begin{aligned}
\mathbf{B}(\pi; \mu_0) &= \mathcal{L}(\mathbf{f}^{\mathrm{BNe}}; \mu_0, \emptyset) - \mathcal{L}(\mathbf{f}^{\mathrm{BNe}}; \mu_0, \pi) && \text{(6.17a)}\\
&= \mathcal{L}^{\mathrm{Ne}}(\overline{\alpha}_0) - \sum_{s \in S}^{m} \psi(s) \cdot \mathcal{L}^{\mathrm{Ne}}(\overline{\alpha}_s) && \text{(6.17b)}\\
&\geq \mathcal{L}^{\mathrm{Ne}}(\breve{\alpha}) - \hat{\mathcal{L}}^{\mathrm{Ne}}(\overline{\alpha}_0) && \text{(6.17c)}\\
&= \hat{\mathcal{L}}^{\mathrm{Ne}}(\breve{\alpha}) - \hat{\mathcal{L}}^{\mathrm{Ne}}(\overline{\alpha}_0) && \text{(6.17d)}\\
&\geq -\Theta\|\mathbb{E}[\alpha] - \breve{\alpha}\|_2, && \text{(6.17e)}
\end{aligned}$$

where (6.17b) holds from Lemma 4, (6.17c) holds from $\mathcal{L}^{\mathrm{Ne}}$ monotonically increasing in α and $\breve{\alpha}_{e,d} \leq \alpha_{e,d}$ for all $e \in E$, $d \in \mathcal{D}$, and $\alpha \in A$, as well as the concave closure $\hat{\mathcal{L}}^{\mathrm{Ne}}$ being greater than any concave combination of points in A and $\sum_{s \in S} \psi(s)\overline{\alpha}_s = \overline{\alpha}_0$, (6.17d) holds from $\mathcal{L}^{\mathrm{Ne}}(\breve{\alpha}) = \hat{\mathcal{L}}^{\mathrm{Ne}}(\breve{\alpha})$ due to $\breve{\alpha}$ being a corner of $\hat{A}$, and (6.17e) holds from Lemma 5 and the definition of Θ from the theorem statement along with the observation that the maximum gradient in the concave closure $\hat{\mathcal{L}}^{\mathrm{Ne}}$ must also occur in the original function $\mathcal{L}^{\mathrm{Ne}}$ by the intermediate value theorem.

Now, we prove the upper bound using similar methods:

$$\begin{aligned}
\mathbf{B}(\pi;\mu_0) &= \mathcal{L}(\mathbf{f}^{\mathrm{BNe}};\mu_0,\emptyset) - \mathcal{L}(\mathbf{f}^{\mathrm{BNe}};\mu_0,\pi) && (6.18a)\\
&= \mathcal{L}^{\mathrm{Ne}}(\overline{\alpha}_0) - \sum_{s\in S}^{m} \psi(s)\cdot \mathcal{L}^{\mathrm{Ne}}(\overline{\alpha}_s) && (6.18b)\\
&\leq \mathcal{L}^{\mathrm{Ne}}(\overline{\alpha}_0) - \mathcal{L}^{\mathrm{Ne}}(\check{\alpha}) && (6.18c)\\
&\leq \Theta\|\mathbb{E}[\alpha] - \check{\alpha}\|_2. && (6.18d)
\end{aligned}$$

Together, these two bounds show the range of attainable performance improvements by utilizing a signalling policy π. The fact that $\mathbf{B}(\pi;\mu_0)$ can be negative shows that providing information to users need not always help. In fact, for any signalling policy π, there exist scenarios where adding information can be detrimental to system performance.

To see (6.13), consider a two link parallel network where $\ell_1(f_1) = 1$, and $\ell_2(f_2) = \sum_{d\in\mathcal{D}} \beta(f_2)^d$, i.e., $\alpha_{2,d} = \beta$ for all $d \in \mathcal{D}$. When $\beta \leq 1/|\mathcal{D}|$, $f_2 = 1$ and $\mathcal{L}^{\mathrm{Ne}}(\beta) = \beta\cdot|\mathcal{D}|$. When $\beta > 1$, $f_1 > 0$ and $\mathcal{L}^{\mathrm{Ne}}(\beta) = 1$. Consider a distribution μ_0 where $\beta = 0$ with probability $1-\epsilon$ and $\frac{1}{\epsilon\cdot|\mathcal{D}|}$ with probability ϵ. The expected value of β is thus $1/|\mathcal{D}|$ The expected total latency without signalling is $\mathcal{L}^{\mathrm{Ne}}(\beta = 1/|\mathcal{D}|) = 1$. Because any truthful signal π will reveal the two possible realizations of β, the expected total latency with π is

$$\mathcal{L}^{\mathrm{Ne}}(0)(1-\epsilon) + \mathcal{L}^{\mathrm{Ne}}\left(\frac{1}{\epsilon\cdot|\mathcal{D}|}\right)\epsilon \to \mathcal{L}^{\mathrm{Ne}}(0) = 0$$

as $\epsilon \to 0$. The benefit of π is thus

$$\mathcal{L}(\mathbf{f}^{\mathrm{BNe}};\mu_0,\emptyset) - \mathcal{L}(\mathbf{f}^{\mathrm{BNe}};\mu_0,\pi) = 1 = -\sqrt{|\mathcal{D}|}\cdot\|\mathbb{E}[\alpha] - \check{\alpha}\|_2,$$

where α is a vector with $|\mathcal{D}|$ entries of value $\frac{1}{|\mathcal{D}|}$.

To see (6.14), consider the two link parallel network where $\ell_1(f_1) = (f_1)^d + \beta$ and

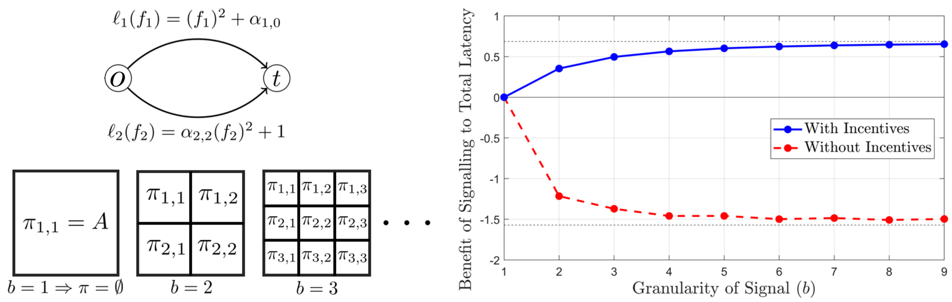

Figure 6.5: The benefit of revealing information with and without the concurrent use of incentives. π is the uniform-grid signal structure where the support A is partitioned into a grid with granularity b; as b increases, more information is revealed to the users. At left, the benefit of using the uniform-grid signalling policy π^b is shown for the setting described in Section 6.4.3 with and without the concurrent use of the incentive mechanism $T^\star$. When incentives are used, revealing information provides a positive benefit and improves performance, which is shown to be generally true in Theorem 6.3.1. With no incentives, the benefit becomes negative, and revealing information worsens system cost, which was shown to be possible in Example 4 and Theorem 6.4.1.

$\ell_2(f_2) = \beta(f_2)^d + 1$ and $f_1 + f_2 = 1$; in this congestion game, $\beta = \alpha_{1,0} = \alpha_{2,d}$ is a single parameter that represents two, correlated coefficients. It is difficult to characterize the Nash flow in closed form; however, we can utilize the following two facts (1)$\frac{\partial}{\partial\beta}\mathcal{L}^{\rm Ne}(\beta)|_{\beta=0} = 0$ and (2) $\frac{\partial}{\partial\beta}\mathcal{L}^{\rm Ne}(\beta)|_{\beta\to\infty} = 1$. Let μ_0 be a distribution on β. From the first fact, $\mathcal{L}(\mathbf{f}^{\rm BNe}; \mu_0, \emptyset) = \mathcal{L}^{\rm Ne}(\overline{\beta}) \to \mathcal{L}^{\rm Ne}(0)$ as $\overline{\beta} \to 0$. Now, consider the prior distribution $\mu_0(\alpha) = \{0, \text{ w.p. } 1-\epsilon, \ \overline{\beta}/\epsilon, \text{ w.p. } \epsilon\}$. Any signalling policy π will reveal which β as 0 or $\overline{\beta}/\epsilon$, as such,

$$\mathcal{L}(\mathbf{f}^{\rm BNe}; \mu_0, \pi) = \mathcal{L}^{\rm Ne}(0)(1-\epsilon) + \mathcal{L}^{\rm Ne}(\overline{\beta}/\epsilon)\epsilon.$$

From fact (2) above, as $\epsilon \to 0$, $\mathcal{L}(\mathbf{f}^{\rm BNe}; \mu_0, \pi) \to \overline{\beta}$. From these two facts, with sufficiently small $\overline{\beta}$,

$$\mathcal{L}(\mathbf{f}^{\rm BNe}; \mu_0, \emptyset) - \mathcal{L}(\mathbf{f}^{\rm BNe}; \mu_0, \pi) \to -\overline{\beta} = -\|\mathbb{E}[\beta] - \underline{\beta}\|_2,$$

where $\underline{\beta} = 0$ to match (6.14). □

6.4.2 Signals & Incentives

Theorem 6.4.1 showed that revealing information has the possibility of increasing or decreasing system cost. It is already known from Theorem 6.3.1 that concurrently utilizing appropriate monetary incentives removes the possibility of worsening performance; however, it is not yet clear how these incentives affect a signalling policy's ability to improve performance.

Theorem 6.4.2 provides bounds on the benefit a signalling policy can provide while also utilizing the signal-aware incentive mechanism $T^\star$. We see that by concurrently utilizing incentives and signalling, the system designer can guarantee the benefit of signalling is non-negative and still have room for significant improvement.

Theorem 6.4.2. *Consider the class of parallel Bayesian congestion games with polynomial latency functions whose degrees come from the set $\mathcal{D}$. For any distribution over the latency coefficients μ_0 and any signalling policy π, the decrease in the expected total latency of a Bayes-Nash flow from signalling satisfies*

$$0 \leq \mathbf{B}(\pi;\mu_0,T^\star) \leq \Xi \|\mathbb{E}[\alpha] - \check{\alpha}\|_2, \tag{6.19}$$

where $\Xi := |\mathcal{D}| + \frac{\rho^+ - \rho_0^-}{4\rho_1^-}\left(|E| + \sum_{d\in\mathcal{D}\setminus\{0\}}(d+1)^d\right)$, $\rho_0^- = \min_{e\in E}\check{\alpha}_{e,0}$, $\rho_1^- = \min_{e\in E}\check{\alpha}_{e,1}$, $\rho^+ = \max_{e\in E}\sum_{d\in\mathcal{D}}(d+1)\hat{\alpha}_{e,d}$, $\mathbb{E}[\alpha] = \int_{z\in A} z\mu_0(z)dz$, *and* $\check{\alpha} \in \mathbb{R}_{\geq 0}^{|E|\cdot|\mathcal{D}|}$ *such that* $\check{\alpha}_{e,d} = \inf\{\mathrm{supp}(\alpha_{e,d})\}$ *for each* $e \in E$, $d \in \mathcal{D}$.

Comparing the bounds on the benefit of a signalling policy with and without the use of incentives (i.e., (6.12) and (6.19)), we see that incentives can make the use of signals more robust (non-negative benefit) while allowing for similar opportunities to improve performance. We further support this conclusion in Section 6.4.3 by providing a numerical example and comparing the benefit of revealing information with and without incentives.

Before we prove Theorem 6.4.2, we state the following lemma that is similar to Lemma 5 but applies to $\mathcal{L}^\star$.

Lemma 6. *Consider the class of parallel congestion games with polynomial latency functions with degrees coming from the set $\mathcal{D}$ with coefficients $\alpha \in A$. Let $a, b \in \mathbb{R}_{\geq 0}^{|E|\cdot|\mathcal{D}|}$ be two possible sets of coefficients for a congestion game with edge set E, then*

$$|\mathcal{D}| + \frac{\rho^+ - \rho_0^-}{4\rho_1^-}\left(|E| + \sum_{d\in\mathcal{D}\setminus\{0\}} (d+1)^d\right) \geq \frac{\mathcal{L}^\star(a) - \mathcal{L}^\star(b)}{||a-b||_2}. \tag{6.20}$$

where $\rho_0^- = \min_{e\in E} \check{\alpha}_{e,0}$, $\rho_1^- = \min_{e\in E} \check{\alpha}_{e,1}$, $\rho^+ = \max_{e\in E} \sum_{d\in\mathcal{D}} (d+1)\hat{\alpha}_{e,d}$.

The proof of Lemma 6 is in Appendix B.

Proof of Theorem 6.4.2: Consider the prior μ_0 and signalling policy $\pi : A \to \Delta(S)$. If μ_s is the posterior formed from receiving signal s, then let $\overline{\alpha}_s = \mathbb{E}_{\alpha\sim\mu_s}[\alpha]$. When utilizing the incentive mechanism $T^\star$ which assigns incentives as stated in (6.8), then Proposition 6.3.1 states that the equilibrium flow that emerges when signal s is received will be $f^\star(s) \in \arg\min \mathcal{L}(f; \overline{\alpha}_s)$; as such $\mathcal{L}^\star(\overline{\alpha}_i)$ is the total latency that occurs.

The lower bound of (6.19) is immediate from Theorem 6.3.1. For the upper bound, we show that

$$\begin{aligned}
\mathbf{B}(\pi; \mu_0, T^\star) &= \mathcal{L}(\mathbf{f}^{\mathrm{BNe}}; \mu_0, \emptyset, T^\star) - \mathcal{L}(\mathbf{f}^{\mathrm{BNe}}; \mu_0, \pi, T^\star) \\
&= \mathcal{L}^\star(\overline{\alpha}_0) - \sum_{s\in S} \psi(s) \cdot \mathcal{L}^\star(\overline{\alpha}_s) && (6.21a) \\
&\leq \mathcal{L}^\star(\overline{\alpha}_0) - \mathcal{L}^\star(\check{\alpha}) && (6.21b) \\
&\leq \Xi \|\mathbb{E}[\alpha] - \check{\alpha}\|_2, && (6.21c)
\end{aligned}$$

where the (6.21a) holds from Lemma 4, (6.21b) from $\mathcal{L}^\star$ non-decreasing with α, and

(6.21c) from Lemma 6 and the definition of Ξ in the theorem statement.

The bound can be proven tight by considering an example in which $d_i = 0$ for each $d \in \mathcal{D}$ (i.e., all latency terms are constant. Consider an Bayesian congestion game in a two link parallel network in which the first edge has latency $\ell_1(f_1) = 1$ and the second has $\ell_2(f_2) = \sum_{d \in \mathcal{D}} \zeta$, where $\zeta = \alpha_{2,d} \geq 0$ for each $d \in \mathcal{D}$ is a single unknown latency parameter that represents $|\mathcal{D}|$, perfectly correlated coefficients. Let μ_0 be a prior distribution on ζ such that $\mu_0(\zeta = 0) = 1 - \epsilon$ and $\mu_0(\zeta = 1/(|\mathcal{D}| \cdot \epsilon)) = \epsilon$; as such $\mathbb{E}_{\zeta \sim \mu_0}[\zeta] = 1/|\mathcal{D}|$ and using Lemma 4 tells us the expected total latency without signalling is $\mathcal{L}(\mathbf{f}^{\mathrm{BNe}}; \mu_0, \emptyset, T^\star) = 1$. Now, consider the use of a signalling policy π that reveals the state to the users; again, using Lemma 4, we see the total latency with the signalling policy π is $\mathcal{L}(\mathbf{f}^{\mathrm{BNe}}; \mu_0, \pi, T^\star) = 0$. To see this matches our bound in (6.19), we note that the mean coefficient vector is $\mathbb{E}[\alpha] = [1, 0, \ldots, 0, 1/|\mathcal{D}|, \ldots, 1/|\mathcal{D}|]^T$ and the bottom of our support is $\check{\alpha} = [1, 0, \ldots, 0]^T$. Substituting this in, we get that the bound in (6.19) equates to 1. $\square$

6.4.3 Benefit of Truthful Signalling

To understand how the benefit of signalling changes as more truthful information is revealed, we offer the following numerical example. Consider a Bayesian congestion game with two edges, $\ell_1(f_1) = f_1^2 + \alpha_{1,0}$ and $\ell_2(f_2) = \alpha_{2,2} f_2^2 + 1$, where $\alpha_{1,0}$ and $\alpha_{2,2}$ are parameters unknown to the user. We consider that these two parameters are drawn from a truncated normal distribution, i.e., let

$$z \sim \mathcal{N}\left(\begin{bmatrix} 30 \\ 30 \end{bmatrix}, 180 \cdot \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix} \right),$$

and define the prior over $\alpha_{1,0}, \alpha_{2,2}$ as $\mu_0(\alpha) = \mathbb{P}[\alpha = z | z \in A]$ where $A = [0, 60]^2$ is their support.

Now, we analyze the benefit of the *uniform-grid signalling policy* with and without the concurrent use of the signal-aware incentive mechanism $T^\star$ as defined in Proposition 6.3.1. This signalling policy is truthful in that each user is informed accurately of what partition the realization of the latency coefficient parameter α is in and reveals more information as the number of partitions increases. Let b be an integer representing the granularity of the signalling mechanism, i.e., the number of times A is partitioned along each dimension as shown in Fig. 6.5, i.e.,

$$\pi_{i,j} = \left[\frac{60}{b}(i-1), \frac{60}{b}i\right] \times \left[\frac{60}{b}(j-1), \frac{60}{b}j\right],$$

essentially forming a uniform grid over A.

In Fig. 6.5, we plot the benefit of using the uniform-grid signalling policy with and without the concurrent use of the incentive mechanism $T^\star$, i.e., $\mathbf{B}(\pi; \mu_0)$ and $\mathbf{B}(\pi; \mu_0, T^\star)$. Observe that when no incentives are used, increasing the amount of information revealed to users (i.e., larger b) causes the benefit to become increasingly negative; meaning as more information is revealed, the signalling policy makes the system performance worse. Conversely, while using incentive mechanism $T^\star$, as more information is added, the benefit becomes increasingly positive and revealing information now improves performance.

6.5 Optimal Signal Design

In the preceding sections of this paper, we showed the range of possible benefit a signalling policy can provide with and without the concurrent use of monetary incentives. In this section, we address how one can compute an optimal signalling policy $\pi^\star$. In

Latency Functions		
State 1 w.p. 0.3	State 2 w.p. 0.4	State 3 w.p. 0.3
$\ell_1(f_1) = 25f_1 + 5$	$\ell_1(f_1) = 30f_1 + 25$	$\ell_1(f_1) = 30f_1 + 25$
$\ell_2(f_2) = 17f_2 + 10$	$\ell_2(f_2) = 35f_2 + 13$	$\ell_2(f_2) = 17f_2 + 10$
$\ell_3(f_3) = 30f_3 + 25$	$\ell_3(f_3) = 13f_3 + 15$	$\ell_3(f_3) = 30f_3 + 25$
$\ell_4(f_4) = 10f_4 + 25$	$\ell_4(f_4) = 10f_4 + 25$	$\ell_4(f_4) = 11f_4 + 35$

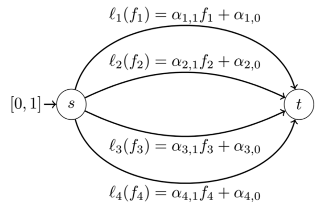

Info./Incentive Setting	Total Lat.
No signal / No tolls	23.99
True signal / No tolls	22.25
Opt. signal / No tolls	21.91*
No signal / w/ ag. tolls	23.41
Opt. signal / w/ ag. tolls	21.35*
No signal / w/ aw. tolls	23.41
Opt. signal / w/ aw. tolls	21.29*

Figure 6.6: Simulation results for the system cost (expected total latency) in a four-link parallel congestion game with affine latency functions. The comparison is made between seven information/incentive settings: no, true/full, and optimal signalling, as well as no signalling and optimal signalling with the use of concurrent signal-agnostic (ag. tolls) and signal-aware (aw. tolls) incentives, respectively. The optimal signals and associated total latency are found using (PT) and (P) with and without constraints (6.23); the asterisk denotes that the solution is approximate, found using the GloptiPoly solver. We find that the optimal signals provide notable improvements over signalling naively (true signal) and that both types of tolls further aid the benefit of signalling.

general, the optimal signal can be NP-hard to find [179]; however, in many problems, this is not the case [180]. We remain in the context of parallel, polynomial-latency congestion games. However, if the signalling platform is limited (e.g., physical signs or discrete UI options), the designer may be limited to a small/finite number of possible signals, i.e., $|S| = \Lambda \leq |A|$. In Section 6.5.1, we survey the result in [160], which shows a method of transcribing the optimal signalling policy problem as a generalized moment problem (GMP) which can be approximately solved with existing solvers [181]. Because we have observed in this work that the concurrent use of monetary incentives can aid in information signalling, we propose two extensions to solve for co-designed signal/incentive pairs with both signal-aware and signal-agnostic incentives. In Section 6.5.2, we show that the signal-agnostic co-design problem can be done by an expansion of the decision variables and the problem remains a GMP. In Section 6.5.3, we show that the signal-aware co-design allows for a simplification where the polynomial constraints can be removed, making the program geometric and solvable via convex programming techniques.

6.5.1 Computing Optimal Signals

We assume that α is realized from a prior distribution μ_0 with finite support $A = \{\alpha^1, \ldots, \alpha^m\}$. A signalling policy $\pi : A \to \Delta(S)$ can now be represented by an $m \times \Lambda$ column stochastic matrix, where $\Lambda \leq m$ defines the designer's constraint on their available signals[5], and $\pi(s,k) = \mathbb{P}[s|\alpha^k]$. A signal-dependent flow tuple $\mathbf{f}$ can be represented by a $\Lambda \times n$ matrix, where $\mathbf{f}(s,e)$ is the flow on edge e in the flow that emerges after receiving signal s. The expected system cost of a signalling policy π with flows $\mathbf{f}$ can thus be written as

$$\begin{aligned}\mathcal{L}(\mathbf{f};\mu_0,\pi) &= \sum_{k=1}^{m}\sum_{s=1}^{\Lambda}\mathcal{L}(\mathbf{f}(s,\text{-});\alpha^k)\cdot\mathbb{P}[s\cap\alpha^k]\\ &= \sum_{k=1}^{m}\sum_{s=1}^{\Lambda}\sum_{e\in E}\sum_{d\in\mathcal{D}}\alpha^k_{e,d}(\mathbf{f}(s,e))^{d+1}\cdot\pi(s,k)\mu_0(k). \end{aligned} \tag{6.22}$$

Note that (6.22) is polynomial in $\mathbf{f}$ and π.

In order to find the optimal signalling mechanism, we must introduce a constraint that $\mathbf{f}$ is a Bayes-Nash equilibrium; from Lemma 4, we can do this by requiring $\mathbf{f}(s,\text{-})$ to be a Nash flow with the expected latency coefficients given the signal, i.e.,

$$\sum_{k=1}^{m}\ell_{e,k}(\mathbf{f}(s,e))\mu_s(k) \leq \sum_{k=1}^{m}\ell_{e',k}(\mathbf{f}(s,e'))\mu_s(k),$$
$$\forall e\in E \text{ s.t. } \mathbf{f}(s,e)>0,\ e'\in E,\ s\in S,$$

[5]From [88], we need not consider signal sets with more than m signals.

where $\mu_s(k) = \frac{\pi(s,k)\mu_0(k)}{\psi(s)}$; this can be rewritten as

$$\mathbf{f}(s,e) \cdot \sum_{k=1}^{m} \big(\ell_{e,k}(\mathbf{f}(s,e)) - \ell_{e',k}(\mathbf{f}(s,e'))\big)\pi(s,k)\mu_0(k) \leq 0, \quad \forall e, e' \in E,\ s \in S. \tag{6.23}$$

Using (6.22) and (6.23), along with other constraints, we can write the following optimization problem, whose solution is the signalling mechanism that minimizes expected total latency in a Bayes-Nash flow:

$$\begin{aligned} \underset{\mathbf{f}\in\mathbb{R}_{\geq 0}^{\Lambda\times n},\ \pi\in\mathbb{R}_{\geq 0}^{\Lambda\times m}}{\text{minimize}} \quad & \mathcal{L}(\mathbf{f};\mu_0,\pi) \\ \text{subject to} \quad & (6.23), \\ & \mathbb{1}_\Lambda^T \pi = \mathbb{1}_m^T, \\ & \mathbf{f}\mathbb{1}_n = \mathbb{1}_\Lambda \end{aligned} \tag{P}$$

Note that (P) has a polynomial objective, polynomial inequality constraints, and linear equality constraints. Problems of this form can be cast as instances of the generalized problem of moments and solved approximately using a semidefinite programming approach [181], as discussed in [160]. In Section 6.5.4, we will use the solution to (P) to numerically investigate the benefit of optimal signals.

6.5.2 Optimal Signal-Agnostic Co-design

Example 5 showed that signals and tolls may be less effective when designed separately. We will show how the Co-design with signal-agnostic incentives can be done with little more complication than the signal design case. Let $\tau \in \mathbb{R}_{\geq 0}^n$ be a vector for the signal-agnostic incentive levied on each edge. The objective of the optimization problem will remain the same as in (6.22); however, the equilibrium constraint will be affected by τ. The new equilibrium constraints become

$$\mathbf{f}(s,e)\sum_{k=1}^{m}\big(\ell_{e,k}(\mathbf{f}(s,e))+\tau_e-\ell_{e',k}(\mathbf{f}(s,e'))-\tau_e'\big)\pi(s,k)\mu_0(k) \leq 0,\ \forall e,e'\in E,\ s\in S. \tag{6.24}$$

Using (6.22) and (6.24), we get the new program

$$\begin{aligned} \underset{\mathbf{f}\in\mathbb{R}_{\geq 0}^{\Lambda\times n},\ \pi\in\mathbb{R}_{\geq 0}^{\Lambda\times m},\ \tau\in\mathbb{R}_{\geq 0}^{n}}{\text{minimize}} \quad & \mathcal{L}(\mathbf{f};\mu_0,\pi) \\ \text{subject to} \quad & (6.24), \\ & \mathbb{1}_\Lambda^T\pi=\mathbb{1}_m^T, \\ & \mathbf{f}\mathbb{1}_n=\mathbb{1}_\Lambda \end{aligned} \tag{PT}$$

The program (PT) belongs to the same class of GMPs as (P), but with more constraints. In Section 6.5.4 and Fig. 6.6, we discuss how the co-designed mechanisms may improve performance.

6.5.3 Optimal Signal-Aware Co-design

In this section, we seek to solve for optimal signals/incentive pairs with signal-aware incentives. The incentive design portion of this task is handled by Proposition 6.3.1, which states that the incentive mechanism $T^\star$ is optimal for any π. We thus look for how to design a signalling policy while concurrently using these incentives.

Remark 1. *The optimal signal-incentive pair* $(\pi^\star, T^\star)$ *uses monetary incentives from Proposition 6.3.1 and signalling policy from the solution to* (P) *without constraints* (6.23).

Remark 1 follows from the fact that $T^\star$ causes the Bayes-Nash flow $\mathbf{f}(s,\cdot)$ after receiving a signal s to be one that minimizes the expected total latency given s. As such, removing (6.23) from (P) allows $\mathbf{f}(s,\cdot)$ to be any feasible flow; in the minimization problem $\mathbf{f}(s,\cdot)$ thus becomes one that minimizes the expected total latency, or be what emerges from using $T^\star$.

Additionally, after removing (6.23) from (P), the problem has only linear equality constraints and a posynomial objective. This problem is thus a geometric program and can be transformed into a convex optimization problem [178]. We will use the solution to this program to compare the effectiveness of the optimal signal-incentive pair with signal-aware incentives to other settings in Section 6.5.4.

6.5.4 Value of Optimal Signalling

To quantify the performance of optimal signalling mechanisms, we discuss a generated numerical example and draw several conclusions. The example is described in Fig. 6.6 and depicts a setting in which users must traverse a parallel network with four edges whose travel delays grow in an affine manner. Users are uncertain of these latency functions but know they come from three possible states (potentially caused by road accidents or weather-related hazards). In this problem, we compute the expected total latency in seven different settings: no signalling, truthful signalling which reveals the exact state, and the optimal signal, as well as no signalling and optimal signalling alongside the optimal signal-agnostic and signal-aware incentives, respectively. The optimal signals are found using the polynomial optimization solver GloptiPoly [182], which casts the problem as a generalized moment problem and finds an approximate solution via semi-definite programming (the asterisk in Fig. 6.6 is to denote the signalling mechanisms are found by this approximate solution method). We identify the following observation from the simulation:

1) Signalling can offer notable performance improvements. Simply revealing the truth offered a 7.25% reduction in system cost, and signalling optimally offered an 8.67% reduction.

2) Incentives can further aid in the capabilities of signalling. The optimal signal-incentive pairs – for both signal aware and agnostic incentives – offered the most significant performance improvements over signalling alone or incentivizing alone.

3) Signal-aware incentives give the best performance and make optimal mechanisms easiest to compute. This is apparent from the last row of the right table in Fig. 6.6 and Remark 1.

6.6 Conclusion

In this paper, we study the effectiveness of information signalling in the context Bayesian congestion games. Our main observations are that designing signalling mechanisms and monetary incentives concurrently can offer improvements that cannot be offered by either alone; one such improvement is that concurrently using appropriate monetary incentives and information signals can help avoid cases where revealing information worsens expected total travel latency. To further this understanding, we derive bounds on the possible benefit of signalling with and without the concurrent use of monetary incentives and provide methods to compute the optimal signalling policies.

Future work may investigate the capabilities of a system operator with less reliable mechanisms (e.g., uncertainty of their own about the system state and heterogeneity in users' beliefs and responses to incentives). Additionally, further studies may uncover if these conclusions exist in settings outside of Bayesian congestion games and apply the proposed techniques to more empirical problems.

Chapter 7

Value of Information in Incentive Design: A Case-Study in Simple Congestion Networks

It is well-known that system performance can experience significant degradation from the self-interested choices of human users. Accordingly, in this manuscript we study the question of how a system operator can exploit system-level knowledge to derive incentives to influence societal behavior and improve system performance. Throughout, we focus on a simple class of routing games where the system operator has uncertainty regarding the network characteristics (i.e., latency functions) and population characteristics (i.e., sensitivity to monetary taxes). Specifically, we address the question of what information can be most effectively exploited in the design of taxation mechanisms to improve system performance. Our main results characterize an optimal marginal-cost taxation mechanism and associated performance guarantee for varying levels of network and population information. The value of a piece of information cannot be known a priori, so we adopt a worst-case interpretation of the value a piece of information is guaranteed to provide.

Several interesting observations emerge about the relative value of information, including the fact that the value of population information saturates unless we also acquire more network knowledge.

7.1 Introduction

The self-interested decision making of system users can cause significant degradation in overall performance [33]. This emergent inefficiency caused by selfish behavior is commonly characterized by the ratio between the worst-case social welfare resulting from choices of self-interested users and the optimal social welfare; this quantity is often referred to as the *price of anarchy* [119] and is a well studied metric of system level inefficiency in the areas of resource allocation [183], distributed control [83], and transportation [184]. A common line of research studies how incentive mechanisms can be designed to influence users to make decisions more in line with the social optimal [121]. For the implementation of such incentives to be effective, a system designer must consider how the users will respond.

When designing an incentive scheme, a system designer is benefited by having more information about the problem, e.g., a more accurate model of the system infrastructure or of the human users' behavior. This paper seeks to understand how and what pieces of information aid in the incentive design task. Though increased understanding of the problem setting seems beneficial, in settings such as road traffic [185], power grids [186], supply chains [42], advertising [187], among others, there is an abundance of potential information sources. Though we can devise ways to learn different pieces of information [188], there is some cost that must be invested in acquiring it. The central questions that we focus on in this manuscript are as follows:

(i) What are the incentive mechanisms that optimize the efficiency of the emergent collective behavior for a given level of informational awareness?

(ii) What type of information, i.e., what specific information about the network or population, can be *best* exploited to improve the efficiency of the emergent collective behavior through an appropriately designed incentive mechanism?

We answer question (i) to understand how to effectively use the available information; however; the main message of this manuscript revolves around the answers to question (ii). In particular, is it better for a societal planner to invest in getting more detailed information regarding the infrastructure characteristics or population characteristics? How do the granularity and quality of the information impact the attainable performance guarantees?

We are particularly motivated by problems relating to congestion and traffic, where the system operator may wish to influence users away from more congestible roads. Recent work has studied this problem in practice [189], where the authors performed experiments in Bangalore to see how users respond to incentives persuading them to alter their commute from the common, direct route to a less direct route. The author finds that without additional information, incentives provide limited opportunities and identified value-of-time/price-sensitivity as an important factor. Inspired by this and related studies [190, 191], we introduce a formal model where each user (driver) has to choose between two decisions (direct commute or long commute) and are influenced by their perceived cost (travel time) as well as an imposed tax. To capture the system designer's uncertainty, we consider that each user has their own price sensitivity, affecting how they relate temporal and monetary costs. We model this as a *congestion game* with two links and a population of heterogeneous users, where the behavior that emerges from users self-interested decisions is a *Nash flow*. It is widely known that the system-level behavior can be suboptimal and the degree of suboptimality is typically characterized by the price of anarchy [35]. Typically, this form of analysis is relegated to studying worst-case scenarios; we seek to extend this by considering how available information may alter what performance guarantees are attainable.

Related Works — Research has sought to explore the use of tolling or taxation mechanisms to improve the system cost in congestion games [7, 71, 176]. These monetary incentives alter users' preferences in a manner that reduces the price of anarchy; however, the majority of this work does not consider the effects of user heterogeneity.

In the works that do study heterogeneous users in congestion games, there are a number of positive and negative results pertaining to the effectiveness of taxation mechanisms [192, 120, 5, 193]. On the positive side, there always exists a taxation mechanism that can completely mitigate any efficiency loss [122, 125, 123]. On the negative side, this taxation mechanism intimately depends on the detailed information pertaining to both the network (i.e., topology, edge latency functions, etc.) as well as the population (i.e., demands, sensitivities, etc.), which significantly limits its applicability. Accordingly, recent work in [136] focuses on robust taxation mechanisms that do not require such extensive knowledge. While the derived taxation mechanism does not necessarily guarantee optimality on a network by network basis, it does provide strictly better performance guarantees than the uninfluenced behavior in broad classes of networks. Hence, these results hint at an apparent trade-off between robustness and optimality.

Contributions — In this work, we seek to bridge the gap between optimal taxation mechanisms that require detailed information, and robust tolls that require less information but may fail to perfectly optimize routing. We consider a case study in 8 information domains and derive the tolling scheme that makes use of the available information *optimally* as well as the resulting price of anarchy bound.

Section 7.2.3 highlights these comparisons, and section Section 7.3 provides formal proofs. Though the system model we consider is simple relative to the general class of congestion games, the observations of this work (1) provide lower bounds on the possible inefficiency that can occur more generally, and (2) discover phenomona that, if can occur in simple settings, can occur more broadly and require consideration in future planning. Chief among these observations is that acquiring environmental knowledge (about the congestion rates of roads) proves more valuable than population knowledge (the exact price sensitivity of each user). Additional findings are discussed in the body of this manuscript.

7.2 Model and Performance Metrics

7.2.1 Congestion Routing Game

Consider a population of users $N = [0, 1]$, represented by a closed interval. To model situations with a very large number of users, a player has infinitesimal mass and is indexed by a real number in $x \in [0, 1] = N$. To model a simple road traffic scenario, the users must traverse a graph from an origin o to a destination d by taking one of two routes, represented by parallel edges e_1 and e_2; this is designed to represent two commute options: a congestible, direct route and an open but indirect route. Let $E = \{e_1, e_2\}$ denote the set of edges. The function $\mathbf{e} : N \to E$ (assumed to be Lebesgue integrable) captures the action of each user, i.e., each user $x \in N$ takes an action by selecting a route $\mathbf{e}(x) \in E$. A flow on edge e is the mass of users taking that route as their action, or $f_e(\mathbf{e}) = \int_{x \in N} \mathbb{1}[\mathbf{e}(x) = e] dx$ where $\mathbb{1}[\cdot]$ is the indicator function. For notational convenience we will omit the flow f reliance on $\mathbf{e}$ when clear from context. Let $f = (f_1, f_2) \in \Delta(E)$ denote a network flow, where $\Delta(E)$ denotes the standard probability simplex over the set E; that is, $\sum_{e \in E} f_e = 1$. To characterize transit delay, each edge

$e \in E$ in the network has a latency function of the form

$$\ell_e(f_e) = a_e f_e + b_e. \tag{7.1}$$

where $a_e \geq 0$ and $b_e \geq 0$ are coefficients used to model how transit delay on an edge grows with more traffic. The latency on an edge is thus a non-decreasing, non-negative function of the flow on that edge. Though this model does not capture all the relevant features of traffic, this setting does capture the decision making of a population of human users. Additionally, this simple model has been used to describe the driving patterns and congestion rates of commuters in real word traffic systems [189], proving useful in characterizing how incentives and users' decision making affect global performance.

For a flow f, the system cost is characterized by the *total latency* in the network, defined as

$$\mathcal{L}(f) = \sum_{e \in E} f_e \cdot \ell_e(f_e), \tag{7.2}$$

and we denote the flow that minimizes this total latency as $f^{\text{opt}} \in \arg\min_{f \in \Delta E} \mathcal{L}(f)$. We specify a particular network by the tuple $G = (E, \{\ell_e\}_{e \in E})$.

This work examines taxation mechanisms as tools to influence the self-interested, price-sensitive user population to reach more efficient equilibria. We model this routing problem as a congestion game where each edge $e \in E$ is assigned a flow dependent tolling function $\tau_e : [0,1] \to \mathbb{R}^+$. A user $x \in N$ has a price-sensitivity $s(x) > 0$; this price-sensitivity is subjective for each user and relates the user's cost from being tolled to their cost from experiencing delays and is the reciprocal of the user's value of time. Without loss of generality, we order players' indices by their individual price-sensitivity, i.e., $s(x) \geq s(y)$ if $x \geq y$. The function $s : N \to \mathbb{R}_{\geq 0}$ thus captures the distribution of price-sensitivity over the users in population N. In a flow f, the cost function for a user

x that is on an edge $\mathbf{e}(x) \in E$ can be expressed as

$$J_x(f) = \ell_{\mathbf{e}(x)}\left(f_{\mathbf{e}(x)}\right) + s(x)\tau_{\mathbf{e}(x)}\left(f_{\mathbf{e}(x)}\right). \tag{7.3}$$

Each user will choose to take the route that minimizes their own cost. When each user does so, the system reaches a *Nash equilibrium* $\mathbf{e}^{\mathrm{Ne}}$, satisfying

$$\mathbf{e}^{\mathrm{Ne}}(x) \in \arg\min_{e \in E}\{\ell_e\left(f_e(\mathbf{e})\right) + s(x)\tau_e\left(f_e(\mathbf{e})\right)\},\ \forall x \in N.$$

In a Nash equilibrium, we will call the resulting network flow a *Nash flow* $f^{\mathrm{Ne}} := f(\mathbf{e}^{\mathrm{Ne}})$ (again, we typically omit the reliance on $\mathbf{e}$ for brevity), also known as a Wardrop equilibrium [151]. A game is therefore characterized by a network G, price-sensitivity distribution $s : [0, 1] \to \mathbb{R}^+$, and a set of tolling functions $\{\tau_e\}_{e \in E}$, denoted by the tuple $(G, s, \{\tau_e\}_{e \in E})$. It is shown in [140] that a Nash flow will always exist in a congestion game of this form, and the total latency of a Nash flow is unique for each s.

7.2.2 Taxation Mechanisms & Performance Metrics

To understand the robustness of a tolling scheme, we consider the performance over a class of networks and users' sensitivities. For a network G, we identify the latency functions which constitute the network by $L(G)$; further, for a family of congestion games $\mathcal{G}$, let $L(\mathcal{G}) = \bigcup_{G \in \mathcal{G}} L(G)$ be the set of all latency functions that exist in the games in $\mathcal{G}$.

A *taxation mechanism* T maps latency functions ℓ_e to tolling functions τ_e. For a family of networks $\mathcal{G}$, this mapping is denoted $T : L(\mathcal{G}) \to \mathcal{T}$, where $\mathcal{T}$ is the set of all admissible tolling functions on $[0, 1]$. In this work, we consider a form of tolling function that is linear with the flow on that edge known as *scaled marginal-cost tolls*. We parameterize the tolls by

$$\tau_e(f_e) = kf_e \cdot \frac{d\ell_e}{df_e}(f_e) = ka_e f_e, \quad \forall\, e \in E, \tag{7.4}$$

where k is a parameter set by the system designer and a_e is the linear component of the edge latency function. Though broader forms of tolling mechanisms can be used to effectively influence users, scaled marginal-cost tolls offer several properties useful for analysis and implementation. If one considers a setting where the toll designer is under no constraint (outside of the implied information constraints), then unbounded incentives can be designed that guarantee optimal performance in nearly every setting by using unbounded step functions or unbounded incentives as described in [136, 5]. Because unbounded incentives are not reasonable in many settings, a toll designer has two options: they could choose to add a constraint bounding the magnitude of the incentives, or they could restrict their design to a class of tolling mechanisms that are intrinsically bounded. In this work, we focus on the latter by studying the design of optimal scaled marginal-cost tolls which are bounded whenever the latency is finite. Scaled marginal-cost tolls have been studied in congestion games with little available information on the network or users' price sensitivities [142]; further, in [136], it is shown that in some low information settings the optimal bounded tolls and associated performance guarantees can be found by solving for the optimal scaled marginal-cost toll; because of their desirable properties and connection to the literature, we consider scaled marginal-cost tolls throughout.

To formalize the notion of uncertainty in users' response, we consider families of sensitivity distributions that can occur when the system designer is only aware of the lower bound $S_{\rm L}$ and upper bound $S_{\rm U}$ on users' price sensitives. We define the set of possible sensitivity distributions as $\mathcal{S} = \{s : [0,1] \to [S_{\rm L}, S_{\rm U}]\}$. When the average price sensitivity $\overline{s}$ of the users is introduced to the system designer, the set of possible distributions becomes $\mathcal{S}(\overline{s}) = \{s \in \mathcal{S} | \int_0^1 s(x)dx = \overline{s}\}$; it is clear that $\mathcal{S}(\overline{s}) \subseteq \mathcal{S}$. To evaluate the performance of a tolling mechanism, let $\mathcal{L}^{\rm nf}(G, s, T)$ be the total latency on a network G, with price sensitivity distribution s, in the Nash flow $f^{\rm Ne}$ when tolls are assigned according to taxation mechanism[1] T, and let $\mathcal{L}^{\rm opt}(G)$ be the minimum total latency which occurs under the optimal flow $f^{\rm opt}$. The *price of anarchy* compares the Nash flow on a network with the optimal flow; this characterizes the inefficiency of the network and can be defined as

$$\mathrm{PoA}(G, s, T) = \frac{\mathcal{L}^{\rm nf}(G, s, T)}{\mathcal{L}^{\rm opt}(G)} \geq 1. \tag{7.5}$$

We extend this definition to include families of networks and sensitivity distributions, i.e.,

$$\mathrm{PoA}(\mathcal{G}, \mathcal{S}, T) = \sup_{G \in \mathcal{G}} \sup_{s \in \mathcal{S}} \left\{ \frac{\mathcal{L}^{\rm nf}(G, s, T)}{\mathcal{L}^{\rm opt}(G)} \right\}, \tag{7.6}$$

such that the price of anarchy is now the worst-case inefficiency over possible networks and populations. Note that the same taxation mechanism T is applied to any realized instance.

7.2.3 Optimal Tolling & Our Contributions

The system designer's goal when designing a taxation mechanism is to minimize worst-case inefficiency given uncertainties over the network and/or user sensitivities. Thus, we define an optimal tolling mechanism as

$$T^* \in \operatorname*{arg\,inf}_{T : L(\mathcal{G}) \to \mathcal{T}} \mathrm{PoA}(\mathcal{G}, \mathcal{S}, T),$$

[1] The taxation mechanism is a mapping from latency functions to tolling functions. A game with taxation mechanism T is therefore denoted $(G, s, \{T(\ell_e) | \ell_e \in G\})$. For brevity, we simply denote this as (G, s, T).

	$\mathcal{S}_{>0}$ *sensitivity-agnostic*	$\mathcal{S}$ *bound-aware* $(S_{\rm L}, S_{\rm U})$	$\mathcal{S}(\bar{s})$ *mean-aware* $(S_{\rm L}, S_{\rm U}, \bar{s})$	$s(x)$ *distribution-aware*
$\mathcal{G}$ *network-agnostic*	$\text{PoA}^*(\mathcal{G}, \mathcal{S}_{>0}) = 1.\overline{33}$ ([16], Prop. 3)	$\text{PoA}^*(\mathcal{G}, \mathcal{S}) \approx 1.176$ (A, Thm. 1)	$\text{PoA}^*(\mathcal{G}, \mathcal{S}(\bar{s})) \leq 1.1401$ (B, Thm. 2)	$\text{PoA}^*(\mathcal{G}, s) \leq 1.1401$ (C, Thm. 3)
G *network-aware*	$\text{PoA}^*(G, \mathcal{S}_{>0}) \leq 1.\overline{33}$ (D, Prop. 3)	$\text{PoA}^*(G, \mathcal{S}) \leq 1.09$ (E, Thm. 4)	$\text{PoA}^*(G, \mathcal{S}(\bar{s})) \leq 1.0494$ (F, Thm. 5)	$\text{PoA}^*(G, s) = 1$ (G, Thm. 6), [25]

Figure 7.1: Price of anarchy bounds under optimal taxation mechanisms with varying amounts of partial information ($S_{\rm U}/S_{\rm L} = 10$).

such that it is the taxation mechanism which minimizes the price of anarchy expressed in (7.6) for a given family of networks $\mathcal{G}$ and sensitivity distributions $\mathcal{S}$. To understand and compare the value of different pieces of information, we seek to quantify the performance guarantees under the optimal tolling mechanism in different information settings. Therefore, we define the price of anarchy bound under an optimal tolling mechanism as

$$\text{PoA}^*(\mathcal{G}, \mathcal{S}) \triangleq \inf_{T:L(\mathcal{G})\to\mathcal{T}} \text{PoA}(\mathcal{G}, \mathcal{S}, T), \tag{7.7}$$

which will serve as the measure of how useful information is to the system designer.

In this paper, we demonstrate the value of different pieces of information to a system designer by comparing the price of anarchy bounds of the optimal incentive in different information settings as shown in Fig. 7.1. We consider these questions in the class of two link parallel networks as these networks often display worst-case inefficiency over larger classes of networks and allow us to analyze the benefit of these partially informed tolls [141]. Many of the results generalize to parallel and more general networks; for those that do not, these results provide lower bounds on the price of anarchy. In Section 7.4, we discuss this context in more detail. For uniformity of presentation, all results are expressed for two link networks. Additionally, the purpose of this work is to identify information factors that affect the incentive design task. If interesting observations can occur in simple problems, then certainly they can occur more generally.

Fig. 7.1 depicts a snap-shot of the theoretical results for $S_{\mathrm{U}}/S_{\mathrm{L}} = 10$. In the top left, the toll designer possesses no information about the network or users' price sensitivities, making the zero toll ($\tau_e = 0 \forall e \in E$) optimal and recovering the price of anarchy bound for this class of networks of $4/3 = 1.\overline{33}$ [35]; we show this formally in Proposition 7.3.3. As the toll designer acquires more information, their performance improvements are captured by moving down and to the left. The information available to the system designer is encoded in the arguments of the price of anarchy expression defined in (7.7). With regard to network information, we consider two possible cases of information available to the system designer:

- ***network-agnostic*** $\mathrm{PoA}^*(\mathcal{G}, \ \cdot \ )$ - the system designer is unaware of the specific problem instance and only knows the class of possible networks and must choose a taxation mechanism that is applied to each[2].
- ***network-aware*** $\mathrm{PoA}^*(G, \cdot)$ - the system designer is aware of the specific network and may design tolls for that specific instance[3].

When the system designer is aware of the exact network characteristics, they will be able to design tolls more effectively. From this fact, we expect a network-aware toll to perform no worse than a network-agnostic one. The benefit of this information for different settings can be seen by comparing the two rows of Fig. 7.1.

Similarly, we consider several settings for the system designer's knowledge of the user-sensitivity distribution:

[2]Though it was assumed prior that the demand in the network is always of unit size, when the system designer is network-agnostic it is without loss of generality that they are also unaware of the demand in the network. We thus consider demand as an implied piece of network information.

[3]Though the network structure is consistent throughout each routing problem, we use the nomenclature of network-aware to match the literature, where network-agnostic tolls must be assigned with only local edge latency characteristics while network-aware tolls are designed with information of each edge's latency function.

- ***sensitivity-agnostic*** $\mathrm{PoA}^*(\ \cdot\ , \mathcal{S}_{>0})$ - the system designer knows nothing about the users' sensitivities except they are bounded away from zero.

- ***bound-aware*** $\mathrm{PoA}^*(\ \cdot\ , \mathcal{S})$ - the system designer knows the lower-bound S_{L} and upper-bound S_{U}, on users' possible sensitivities.

- ***mean-aware*** $\mathrm{PoA}^*(\ \cdot\ , \mathcal{S}(\overline{s}))$ - the system designer knows the lower-bound S_{L} and upper-bound S_{U} as well as the mean $\overline{s}$ of users' sensitivities.

- ***distribution-aware*** $\mathrm{PoA}^*(\cdot, s)$ - the system designer knows the exact distribution on user sensitivities.

The sensitivity distribution serves as a model for the population's behavior. Refining the set of possible distributions reduces the designer's uncertainty and allows them to design more effective tolls. The benefit of increasing the information available to the system designer can be seen by comparing the columns of Fig. 7.1.

The main focus of this work is to investigate which pieces of information give the greatest gains in the performance of tolls with respect to the price of anarchy ratio. Though it is clear that additional information will help tolls provide better guarantees, it is not obvious what will provide a better gain in performance when introduced to an uninformed system designer: network-awareness or population-awareness. Further, the value of a piece of information is highly contextual, and it is impossible to know a priori what value new information provides, thus we adopt a worst-case approach for comparing performance. Comparing the worst-case performance bounds of each setting (such as those demonstrated in Fig. 7.1) shows:

1. Comparing elements (B) and (C) shows that the full distribution of users' price sensitivities need not be any more helpful than the mean alone, thus the value of population information saturates.

2. Comparing elements (D) and (G) shows that, in the absence of any population information, network-awareness may be of no help, however, in the presence of full population information, network-awareness allows for tolls that can always incentivize optimal self-routing.

3. Comparing elements (B) and (E) shows that the guaranteed value of information about network characteristics is more valuable than the guaranteed value of the mean of users' price sensitivities.

To further illustrate the results and highlight that the relationships between information settings hold more generally, we provide a plot of each price of anarchy bound for varying levels of population heterogeneity. Fig. 7.2 shows the best attainable price of anarchy bounds under scaled marginal cost tolling in the previously described information settings for each $S_{\mathrm{L}}/S_{\mathrm{U}} \in [0,1]$. As the $S_{\mathrm{L}}/S_{\mathrm{U}}$ approaches 1, there is less discrepancy between the different users' price sensitivities and all tolls can optimize performance; as $S_{\mathrm{L}}/S_{\mathrm{U}}$ approaches zero, the differences between users' responses can be arbitrarily large and no toll is effective. For values of $S_{\mathrm{L}}/S_{\mathrm{U}}$ in between, we see that the previously described relationships hold, and network information proves more valuable than population information. These findings illuminate several important considerations for incentive designers. In problems closely related to our described setting (such as the tolling experiment in [189]), our results inform what information the toll designer should invest in acquiring. Further work is needed to understand if identical conclusions are true in other settings, however identifying them in our setting highlights that these comparisons need to be considered more broadly.

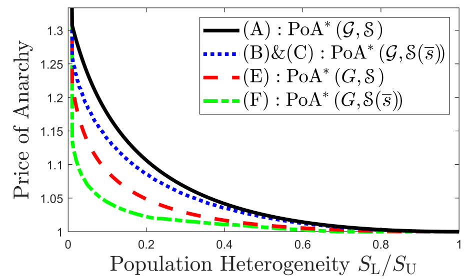

Figure 7.2: Worst-case price of anarchy in each information setting over levels of user heterogeneity $S_\mathrm{L}/S_\mathrm{U} \in [0,1]$. Each plot represents the best achievable price of anarchy bound using scaled marginal cost tolls for one of the information settings: (A) network agnostic, mean agnostic, (B) & (C) network agnostic, mean (or full distribution) aware, (E) network aware, mean agnostic, and (F) network aware, mean aware. Each toll guarantees superior performance than the untolled price of anarchy of 4/3. The values of each line at $S_\mathrm{L}/S_\mathrm{U} = 0.1$ are those presented in Fig. 7.1. By varying the value of $S_\mathrm{L}/S_\mathrm{U}$, we can see that the relationship between settings holds when looking at worst-case performance guarantees.

7.3 Main Results

In each setting considered, we provide a result reporting the price of anarchy bound under an optimal toll, as well as a subsidiary result reporting the the optimal scaled marginal-cost toll when applicable. The main conclusions can be observed numerically in Fig. 7.1, while the analytic expressions are given in the following subsections. In this work, we limit the search for optimal tolls to a search over scaled marginal-cost tolling mechanisms. Taxation mechanisms of this form can be parameterized by a single scaling factor k, and will be denoted $T(k)$ which assigns to an edge e a toll $\tau_e(f_e) = k a_e f_e$. As discussed previously, these tolls posses desirable properties in that they are naturally bounded when the latency is finite, and they can be reasonably implemented in network-aware and network-agnostic settings; further, in [136, Lemma 2.2] the authors show that the optimal *bounded* toll can be found by searching for the optimal scaled marginal-cost toll in the network-agnostic case. In the network-aware case, it is an open question as to what form an optimal, bounded taxation mechanism will take.

7.3.1 Network-agnostic, bound-aware

The first scenario we consider is when the system designer is agnostic of the exact network characteristics and knows only the lower and upper bound on user sensitivities; we provide a bound on the price of anarchy under an optimal scaled marginal-cost taxation mechanism in this setting.

Theorem 7.3.1. *When only S_{L} and S_{U} are known, the price of anarchy under an optimal scaled marginal-cost tolling mechanism is*

$$\mathrm{PoA}^*(\mathcal{G}, \mathcal{S}) = \frac{\left(q - 1 + \sqrt{q^2 + 14q + 1}\right)^2}{8q\left(-q - 1 + \sqrt{q^2 + 14q + 1}\right)}, \tag{7.8}$$

where $q := S_{\mathrm{L}}/S_{\mathrm{U}}$.

This result gives a performance guarantee for the setting where the system designer has minimal information about the population and network characteristics. This result is a generalization of [142, Theorem 1], where we now consider networks that need not have traffic on every edge in a Nash flow. As a means to find this upper bound, we first derive the optimal scaled marginal-cost toll. In the proof of Theorem 7.3.1, at the end of this subsection, we will use this toll to show the associated price of anarchy bound.

Proposition 7.3.1. *When only S_{L} and S_{U} are known, the optimal network-agnostic marginal-cost toll scaling factor is*

$$k_{\mathrm{agn}} = \frac{-S_{\mathrm{L}} - S_{\mathrm{U}} + \sqrt{S_{\mathrm{L}}^2 + 14 S_{\mathrm{L}} S_{\mathrm{U}} + S_{\mathrm{U}}^2}}{2 S_{\mathrm{L}} S_{\mathrm{U}}}. \tag{7.9}$$

Proof. We start by finding the scaling factor k_{agn} for the optimal scaled marginal-cost toll. In this information setting, the optimal scaling factor was found in [194] to be the

Table 7.1: Table of Commonly Used Notation & Abbreviations

E	Edge set indexed by $e \in E$
f_e	Flow on edge e
$\ell_e(\cdot)$	Latency on edge e as a function of edge flow
a_e	Linear coefficient of latency function on edge e
b_e	Constant coefficient of latency function on edge e
$\mathcal{L}(\cdot)$	Total latency as a function of network flow
$\mathcal{L}^{\mathrm{nf}}(\cdot)$	Total latency in a Nash flow
$\mathcal{L}^{\mathrm{opt}}(\cdot)$	Total latency in the optimal flow
G	Congestion routing game problem instance
$\mathcal{G}$	Family of Congestion routing game problems
$\tau_e(\cdot)$	Toll on edge e as a function of edge flow
$s(\cdot)$	User price sensitivity as a function of user index
$\mathcal{S}$	Set of population sensitivity distributions
$S_{\mathrm{L}}, S_{\mathrm{U}}$	Lower and upper bound on users' price sensitivity
$\overline{s}$	Average user price sensitivity
$\tau_e(\cdot)$	Toll applied to edge e as a function of edge flow
$T(\cdot)$	Incentive mechanisms that maps edges to tolls
k_{agn}	scalar of network/sensitivity agnostic toll
$k_{(\overline{s})}$	scalar of network agnostic/mean aware toll
$k_{(G)}$	scalar of network aware/sensitivity agnostic toll
$k_{(\overline{s},G)}$	scalar of network/mean aware toll
k^{gm}	geometric mean scaling factor $1/\sqrt{S_{\mathrm{L}}S_{\mathrm{U}}}$
$\mathrm{PoA}(\cdot)$	Price of anarchy as a function of a class of congestion games, price sensitivity set, and incentive mechanisms
$\mathrm{PoA}^{\star}(\cdot)$	Price of anarchy using optimal toll as a function of a class of congestion games and price sensitivity set
Common Abbreviations	
opt	optimal over respective domain
Nf	Nash flow (often w.r.t. toll, population, and game)

solution to the equation

$$\frac{4}{4\left(1+k_{\mathrm{agn}}S_{\mathrm{L}}\right)-\left(1+k_{\mathrm{agn}}S_{\mathrm{L}}\right)^{2}}=\frac{\left(1+k_{\mathrm{agn}}S_{\mathrm{U}}\right)^{2}}{4k_{\mathrm{agn}}S_{\mathrm{U}}}. \tag{7.10}$$

It is shown in [194] that when $S_{\mathrm{L}} < S_{\mathrm{U}}$, (7.10) always has exactly one solution on the interval $(1/S_{\mathrm{U}}, 1/S_{\mathrm{L}})$, and that solution is the desired optimal scale factor. (7.9) is a solution to (7.10), so we show here that (7.9) describes this desired solution by showing that k_{agn} is in the interval $(1/S_{\mathrm{U}}, 1/S_{\mathrm{L}})$. Define the term $p = S_{\mathrm{U}}/S_{\mathrm{L}}$. Because $p > 1$, we have

$$1+14p+p^{2} > 1+14p+p^{2}+8(1-p)=(p+3)^{2}.$$

Thus, (7.9) can be lower bounded by

$$k_{\mathrm{agn}} > \frac{-1-p+\sqrt{(p+3)^{2}}}{2S_{\mathrm{U}}}=\frac{1}{S_{\mathrm{U}}}. \tag{7.11}$$

Likewise, define $q = S_{\mathrm{L}}/S_{\mathrm{U}}$ (so that $q < 1$), we have

$$1+14q+q^{2} < 1+14q+q^{2}+8(1-q)=(q+3)^{2},$$

yielding a lower bound on k_{agn} of

$$k_{\mathrm{agn}} < \frac{-1-q+\sqrt{(q+3)^{2}}}{2S_{\mathrm{U}}}=\frac{1}{S_{\mathrm{L}}}. \tag{7.12}$$

Thus, the scaling factor k_{agn} defined in (7.9) exists in the interval $(1/S_{\mathrm{U}}, 1/S_{\mathrm{L}})$. It can be shown by substitution that (7.9) satisfies (7.10). □

Proof of Theorem 7.3.1: Using the scaling factor from Proposition 7.3.1, (7.8) can be found by substituting (7.9) into (7.10). □

7.3.2 Network-agnostic, mean-aware

We consider the mean sensitivity as additional information to just the bounds $S_{\rm L}, S_{\rm U}$. If the system designer is aware that the mean user sensitivity is $\bar{s}$, the set of possible sensitivity distributions is reduced to the set $\mathcal{S}(\overline{s}) \subset \mathcal{S}$. Using this information, the toll designer is able to refine the optimal mechanism and improve the performance guarantees. The price of anarchy of an optimal toll in this setting is shown in Fig. 7.1, where value (B) is noticeably lower than (A) demonstrating the value of information to the system designer.

In deriving this bound, we perform a series of reductions in the set of feasible instances to one which realizes the worst-case price of anarchy. In Lemma 7, we show that any Nash flow can emerge by a population with a bimodal sensitivity distribution, thus reducing our search for worst-case instances to those with bimodal sensitivity populations. In Lemma 8 we further identify two distributions, one of which will realize worst-case inefficiency. In Lemma 9 we reduce the search over networks to those with only one linear and one constant edge, and in Lemma 10 we identify two such networks that constitute worst-case instances. We will first prove these lemmas that will aid in proving several later theorems.

We say users x, y have the same *type* if $s(x) = s(y)$. Further let a *bimodal* distribution be one in which there exist exactly two user types; the set of such distributions is denoted $\mathcal{S}^{\rm bi}(\overline{s}) \subset \mathcal{S}(\overline{s})$. We denote a bimodal distribution with types S_1 and S_2 by (S_1, S_2). Note that for a given $\overline{s}$, S_1 and S_2, the mass of users with each sensitivity is well defined. Additionally, we adopt the convention used elsewhere that the network links are indexed such that $b_1 \leq b_2$.

Lemma 7. *A Nash flow f for a sensitivity distribution $s \in \mathcal{S}(\overline{s})$, under a linear tax T, is likewise a Nash flow for some distribution $s' \in \mathcal{S}^{\mathrm{bi}}(\overline{s})$ in which one type of user is indifferent between the two edges and all users on each edge are of a single type. This implies the price of anarchy over sensitivities in $\mathcal{S}(\overline{s})$ is equal to the price of anarchy over bimodal distributions in $\mathcal{S}^{\mathrm{bi}}(\overline{s})$, i.e.,*

$$\mathrm{PoA}(\mathcal{G}, \mathcal{S}(\overline{s}), T) = \mathrm{PoA}(\mathcal{G}, \mathcal{S}^{\mathrm{bi}}(\overline{s}), T). \tag{7.13}$$

Proof. Let $s_1 \in \mathcal{S}(\overline{s})$ be some distribution of users' sensitivities, and let S_{ind} be the sensitivity that has equal cost between the two links in the Nash flow f^{Ne}, i.e., solution to

$$(1 + S_{\mathrm{ind}}k)a_1 f_1^{\mathrm{Ne}} + b_1 = (1 + S_{\mathrm{ind}}k)a_2 f_2^{\mathrm{Ne}} + b_2. \tag{7.14}$$

Note that in the case where $S_{\mathrm{ind}} > S_{\mathrm{U}}$ or $S_{\mathrm{ind}} < S_{\mathrm{L}}$, any distribution $s \in \mathcal{S}$ will have the same Nash flow with all users choosing the same edge. First, consider the case where $S_{\mathrm{ind}} < \mu(s_1)$, where $\mu(\cdot)$ is the mean of the distribution. From Claim 1.1.2 in [142], if a user has a sensitivity $S < S_{\mathrm{ind}}$, then they strictly prefer the first link; if they have a sensitivity $S > S_{\mathrm{ind}}$ then they strictly prefer the second.

Now, let s_2 be a new distribution where each user who had chosen edge 1 now has sensitivity S_{ind}. The Nash flows from s_1 and s_2 are the same, as the same number of users have a sensitivity $S \leq S_{\mathrm{ind}}$ and thus the same users choose the first edge. It is clear that $\mu(s_2) > \mu(s_1)$ as no user has a lower sensitivity and some have higher.

Now, consider a third distribution s_3, where users who chose edge 2 now have some sensitivity $S' \in (S_{\mathrm{ind}}, S_{\mathrm{U}}]$; these users will now strictly prefer the second edge of the network but the Nash flow will remain unchanged. If we pick $S' = S_{\mathrm{U}}$, the mean has surely increased again; if we pick $S' = S_{\text{ind}}$, because we are in the case $S_{\text{ind}} < \overline{s}$, the mean is lower than $\mu(s_1)$. Because $\mu(s_3)$ is continuous with S', we can select S' so that $\mu(s_3) = \mu(s_1)$. The case of $S_{\text{ind}} > \mu(s_1)$ is similar.

The distribution $s_3 = (S_{\text{ind}}, S')$ induces the same Nash flow as s_1 and now, one set of users is indifferent and users of the same type exist on the same edge only. □

Having shown in Lemma 7 that any feasible Nash flow can be realized by a population with a bimodal sensitivity distribution, we note that the worst-case price of anarchy can be realized by a bimodal distribution. Our search further reduces as we characterize two specific distributions that give worst-case inefficiency.

Lemma 8. *For a given network $G \in \mathcal{G}$ and scaled marginal-cost tax T with toll scaling factor k, two distributions $s_\ell^{(\overline{s},G,k)}$ and* supp*, that maximize and minimize (respectively) the flow on the first edge of the network, realize the price of anarchy over those in $\mathcal{S}^{\text{bi}}(\overline{s})$,*

$$\text{PoA}(G, \mathcal{S}^{\text{bi}}(\overline{s}), T) = \text{PoA}\left(G, \{s_\ell^{(\overline{s},G,k)}, \text{supp}\}, T\right) \tag{7.15}$$

Proof. The proof follows from the fact that total latency is quadratic in the flow, thus the largest price of anarchy will come from the flow that is furthest from optimal. From Lemma 7, we see that any flow induced by a distribution $s \in \mathcal{S}(\overline{s})$ can be realized by a bimodal distribution that has one set of users observing equal cost between the links and each edge containing only one sensitivity type. We therefore define $s_\ell^{(\overline{s},G,k)}$ as the distribution that maximizes f_1^{Ne} and supp as the distribution which maximizes f_1^{Ne}. □

Next, we focus on which networks exhibit worst-case inefficiency and reduce our search to the set of instances with have one linear latency function and one constant; the set of such networks is defined as

$$\mathcal{G}^{\mathrm{lc}} = \{G | \ell_1(f_1) = a_1 f_1,\ \ell_2(f_2) = b_2,\ a_1, b_2 \geq 0\}.$$

Lemma 9. *For any $G \in \mathcal{G}$, there exists a $\hat{G} \in \mathcal{G}^{\mathrm{lc}} \subset \mathcal{G}$ that, under the same scaled marginal-cost tolling mechanism $T(k)$, has a higher price of anarchy, implying*

$$\mathrm{PoA}(\mathcal{G}, \mathcal{S}(\overline{s}), T(k)) = \mathrm{PoA}(\mathcal{G}^{\mathrm{lc}}, \mathcal{S}(\overline{s}), T(k)). \tag{7.16}$$

The proof of Lemma 9 appears in the appendix.

Finally, we identify two specific networks that demonstrate worst-case inefficiency. For a given set of distributions $\mathcal{S}(\overline{s})$ and toll scaling factor k, we define two networks:

(1) $G_\beta \in \mathcal{G}^{\mathrm{lc}}$ with latency functions $\ell_1(f_1) = f_1$ and $\ell_2(f_2) = \beta$ and satisfies $s_\ell^{(\overline{s}, G_\beta, k)} = (S_{\mathrm{L}}, S_{\mathrm{U}})$, and

(2) $G_\alpha \in \mathcal{G}^{\mathrm{lc}}$ with latency functions $\ell_1(f_1) = f_1$ and $\ell_2(f_2) = \alpha$ and satisfies $s_u^{(\overline{s}, G_\alpha, k)} = (S_{\mathrm{L}}, S_{\mathrm{U}})$. Due to the discussion in the proof of Lemma 9, any network in $\mathcal{G}^{\mathrm{lc}}$ with cost functions satisfying $b_2/a_1 = \beta$ will have the same price of anarchy as G_β, and the same is true for for G_α.

Lemma 10. *For linear constant networks, under sensitivity distributions in $\mathcal{S}(\overline{s})$ with toll scaling factor k, the network G_α or G_β will realize the upper bound on the price of anarchy, i.e.,*

$$\mathrm{PoA}(\mathcal{G}^{\mathrm{lc}}, \mathcal{S}(\overline{s}), T(k)) = \mathrm{PoA}(\{G_\alpha, G_\beta\}, \mathcal{S}(\overline{s}), T(k)).$$

Proof. It can be seen by differentiation of (B.34), the price of anarchy increases with the value of the indifferent sensitivity when $f_1^{\mathrm{Ne}} < f_1^{\mathrm{opt}}$ and decreases when $f_1^{\mathrm{Ne}} < f_1^{\mathrm{opt}}$. Recall that $s_\ell^{(\bar{s},G,k)}$ has $f_{1\ell} > f_1^{\text{opt}}$ and indifferent sensitivity $S_{\ell 1}$; similarly, supp has $f_{1u} < f_1^{\text{opt}}$ and indifferent sensitivity S_{u2}. It is therefore true that having $S_{\ell 1} = S_{\text{L}}$ or $S_{u2} = S_{\text{U}}$ is a necessary condition for the network which maximizes the price of anarchy.

Further, in bimodal distributions (S_1, S_2) where users are homogeneous on either link, $f_1^{\text{Ne}} = (S_2 - \bar{s})/(S_2 - S_1)$. For $s_\ell^{(\bar{s},G,k)}$ when users with sensitivity $S_{\ell 1} = S_{\text{L}}$ are indifferent, the largest flow that can occur on f_1 occurs when $s_\ell^{(\bar{s},G,k)} = (S_{\text{L}}, S_{\text{U}})$. Similarly, for supp, when users with sensitivity $S_{u2} = S_{\text{U}}$ are indifferent, the least flow in f_1 has supp $= (S_{\text{L}}, S_{\text{U}})$. One of these two conditions must be met by a network $G \in \mathcal{G}^{\text{lc}}$ to maximize the price of anarchy. Those networks are the defined G_α and G_β. □

For ease of notation, we will define a change of variable $z(x) = \frac{1}{1+s(x)k}$ and $z_L = \frac{1}{1+S_{\text{L}}k}$ and $z_U = \frac{1}{1+S_{\text{U}}k}$.

Proposition 7.3.2. *When S_{L}, S_{U} and the mean sensitivity $\bar{s}$ are known, the optimal network-agnostic marginal-cost toll scaling factor $k_{(\bar{s})}$ will be the solution on $(1/S_{\text{U}}, 1/S_{\text{L}})$ to*

$$\text{PoA}(G_\beta, (S_{\text{L}}, S_{\text{U}}), T(k)) = \text{PoA}(G_\alpha, (S_{\text{L}}, S_{\text{U}}), T(k)), \tag{7.17}$$

where

$$\beta = R/z_L, \quad \alpha = \begin{cases} \frac{1}{2(z_U - z_U^2)}, & \frac{1}{2(1-z_U)} \geq R \\ R/z_U, & \text{otherwise.} \end{cases} \tag{7.18}$$

where $R := (S_{\text{U}} - \bar{s})/(S_{\text{U}} - S_{\text{L}})$.

The proof of Proposition 7.3.2 appears in the appendix.

Finally, the price of anarchy under the optimal tolling mechanism can be expressed as in the following theorem.

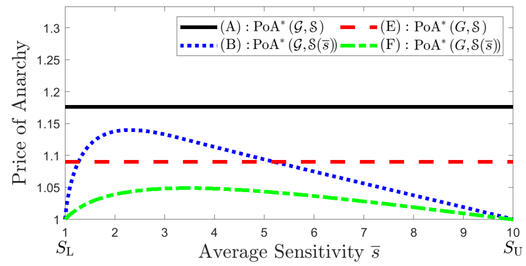

Figure 7.3: Price of anarchy in each information setting over various mean sensitivities. Each plot represents a bound for one of the introduced tolling mechanisms: (A) network agnostic, mean agnostic toll, (B) network agnostic, mean aware toll, (E) network aware, mean agnostic toll, and (F) network aware, mean aware toll. Each toll gives superior performance guarantees than the untolled price of anarchy of 4/3. Price sensitivity bounds $S_\mathrm{L} = 1$ and $S_\mathrm{U} = 10$ are shown; changing these values has a minimal effect on the relation between the lines.

Theorem 7.3.2. *When S_L, S_U and the mean sensitivity $\overline{s}$ are known, the price of anarchy under an optimal, scaled marginal-cost toll is given by*

$$\mathrm{PoA}^*\left(\mathcal{G}, \mathcal{S}(\overline{s})\right) = \frac{R^2 - \beta R + \beta}{\beta - \beta^2/4}, \tag{7.19}$$

where $R := (S_\mathrm{U} - \overline{s})/(S_\mathrm{U} - S_\mathrm{L})$, and $\beta = (1 + S_\mathrm{L} k_{(\overline{s})})R$, with $k_{(\overline{s})}$ being the solution to (7.17).

In Fig. 7.3, we show the price of anarchy bound of the network-agnostic, mean-aware tolls alongside the price of anarchy bound in several other settings. As mentioned before, the value of knowing certain pieces of information (in this case the mean) is highly contextual: when the mean $\overline{s}$ is close to one of the lower or upper bound S_L and S_U, the mean is very informative as users' sensitivities must be more concentrated around the average. However, in worst-case, the mean sensitivity does not offer as much value to the toll designer as the knowledge of the edge latency functions as in the network-aware, mean-agnostic case.

Proof of Theorem 7.3.2: From Lemma 10, a network G_β realizes the price of anarchy when the toll scaling factor is chosen optimally as in Proposition 7.3.2. The price of anarchy for this network is found by substituting β from (B.39) into the latency function ratio in (B.37). □

7.3.3 Network-agnostic, distribution-aware

When the system designer is informed of the average user sensitivity, they are able to improve the price of anarchy ratio by utilizing the new information. It would seem that having precise knowledge would allow further reductions in the price of anarchy; however, in Theorem 7.3.3 it is shown that full information on the user sensitivities does not improve the price of anarchy.

Theorem 7.3.3. *The worst-case performance guarantee for the network-agnostic taxation mechanism with knowledge of the full user sensitivity distribution is no better than that of the network-agnostic, mean-aware:*

$$\max_{s \in \mathcal{S}(\overline{s})} \mathrm{PoA}^*(\mathcal{G}, s) = \mathrm{PoA}^*(\mathcal{G}, \mathcal{S}(\overline{s})). \tag{7.20}$$

When the system designer is uncertain of the network characteristics, the full sensitivity distribution information is no more valuable than the average of the users sensitivities, in worst-case; this is highlighted in Fig. 7.1 where the price of anarchy in box (B) and (C) are equal.

Proof. The proof follows similarly from Section 7.3.2, utilizing Lemma 7, Lemma 8, Lemma 9, and Lemma 10. Observe the two worst-case problem instances: G_α with bimodal distribution $(S_{\mathrm{L}}, S_{\mathrm{U}})$ with mean $\overline{s}$, and G_β with bimodal distribution $(S_{\mathrm{L}}, S_{\mathrm{U}})$ with mean $\overline{s}$. Because the user sensitivity distribution is the same in both instances, if this distribution was known apriori, the networks G_α and G_β would still constitute worst-case instances and the optimal tolling mechanism must be selected as in Proposition 7.3.2

and give the same performance guarantee as if only the mean was known. □

7.3.4 Network-aware, sensitivity-agnostic

In the previous sections, it is shown that additional information about the population of users may help improve the performance guarantees of an optimal taxation mechanism. In the following sections we will also see how full knowledge of the network characteristics can improve the efficacy of tolls. Specifically, in this section we consider the case where the system designer has full information on the network characteristics but knows nothing about the users' sensitivities except that they are bounded away from zero; in this setting, the additional information will not help.

Proposition 7.3.3. *When the the exact network G is known, but users' sensitivities $s \in \mathcal{S}_{>0}$ are only known to be bounded away from zero, no taxation mechanism can improve the price of anarchy, i.e.,*

$$\sup_{G \in \mathcal{G}} \mathrm{PoA}^*(G, \mathcal{S}_{>0}) = \mathrm{PoA}^*(\mathcal{G}, \mathcal{S}_{>0}) = 4/3. \tag{7.21}$$

Proposition 7.3.3 shows shows that even if the exact network characteristics are known, some information about the population's response is required to improve worst-case performance guarantees. This is consistent with box (D) of Fig. 7.1; Further, this implies the top left box as well: when no information on network or population is present, no toll can lower the price of anarchy below 4/3.

Proof. If each user has the same sensitivity S; further, consider the classic Pigou network with two parallel edges, one with latency $\ell_1(f_1) = f_1$ and $\ell_e(f_2) = 1$. In the absence of any tolling, the Nash flow is $f^{\mathrm{Ne}} = (1, 0)$ and the optimal is $f^{\mathrm{opt}} = (1/2, 1/2)$, giving a price of anarchy of 4/3. When all users travel on the same link (either the first or the second) in the Nash flow, the price of anarchy is 4/3, the same as the untolled case. Any tolling mechanism that incentives users to utilize the second link (i.e., $\tau_1 > \tau_2$) can be made arbitrarily ineffective by letting $S \to \infty$, causing the Nash flow to be $f^{\mathrm{Ne}} = (0,1)$ and the price of anarchy to be 4/3. □

7.3.5 Network-aware, mean-agnostic

As seen in the previous section, network information will not help a system designer that has no knowledge of the population. When the system designer at least has bounds on the possible sensitivities of users, the optimal toll will be able to improve performance.

Theorem 7.3.4. *When only S_{L} and S_{U} are known, the price of anarchy under an optimal, network-aware, scaled marginal-cost toll is tightly upper bounded by*

$$\mathrm{PoA}^*(G, \mathcal{S}) \leq \frac{4}{3}\left(1 - \frac{\sqrt{q}}{\left(1+\sqrt{q}\right)^2}\right) \tag{7.22}$$

where $q := S_{\mathrm{L}}/S_{\mathrm{U}}$.

By comparing box (B), (C) and (E) in Fig. 7.1, one can observe that network information is significantly more valuable than additional population information beyond the lower and upper bound.

To prove this bound, we assume the toll designer can determine the Nash flow of a possible homogeneous low-sensitivity population associated with each toll scaling factor; this assumption is reasonable as the Nash flow of a homogeneous population can be found by solving a convex optimization problem [35]. Let $f_i^{\mathrm{Ne}}(G, S, k)$ be the mass of traffic on edge i in network G in a Nash flow of a population of users with homogeneous price-sensitivity S and tolling factor k.

Proposition 7.3.4. *For any network $G \in \mathcal{G}$ and any $S_{\mathrm{U}} \geq S_{\mathrm{L}} > 0$, let $k^{\mathrm{gm}} = (S_{\mathrm{L}}S_{\mathrm{U}})^{-1/2}$.*

The following is an optimal network-aware marginal-cost toll scaling factor:

$$k_{(G)} = \begin{cases} 0 & \text{if } f_2^{\mathrm{Ne}}(G, S_{\mathrm{L}}, k^{\mathrm{gm}}) = 0, \\ k^{\mathrm{gm}} & \text{otherwise.} \end{cases} \tag{7.23}$$

Proof. Consider the following cases, differentiated by the structure of Nash flows resulting from $k = k^{\mathrm{gm}} := (S_{\mathrm{L}} S_{\mathrm{U}})^{-1/2}$:

1. $f_2^{\mathrm{Ne}}(G, S_{\mathrm{L}}, k^{\mathrm{gm}}) > 0$, and

2. $f_2^{\mathrm{Ne}}(G, S_{\mathrm{L}}, k^{\mathrm{gm}}) = 0$.

It is shown in [142] that in Case (1), it must be true that $\mathcal{L}^{\mathrm{nf}}\left(G, S_{\mathrm{L}}, k\right) = \mathcal{L}^{\mathrm{nf}}\left(G, S_{\mathrm{U}}, k\right)$ and that this choice of k is uniquely optimal, resulting in the price of anarchy given in (7.22).

Consider Case (2). Here, the extreme low-sensitivity population with $s = S_{\mathrm{L}}$ strictly prefers link 1 when $k = k^{\mathrm{gm}}$, effectively stripping the designer of their influence over the price of anarchy. It can easily be shown (using, e.g., tools from [142]) that

$$k \leq k^{\mathrm{gm}} \implies \mathcal{L}^{\mathrm{nf}}\left(G, S_{\mathrm{L}}, k\right) = \mathcal{L}^{\mathrm{nf}}\left(G, \emptyset\right), \tag{7.24}$$

but that

$$k^{\dagger} > k^{\mathrm{gm}} \implies \mathcal{L}^{\mathrm{nf}}\left(G, S_{\mathrm{U}}, k^{\dagger}\right) > \mathcal{L}^{\mathrm{nf}}\left(G, \emptyset\right). \tag{7.25}$$

That is, in this regime, the designer cannot change the behavior of $s = S_{\mathrm{L}}$ without increasing tolls, but cannot increase tolls because this would cause the high-sensitivity population with $s = S_{\mathrm{U}}$ to route more inefficiently. That is, $k = 0$ is an optimal tolling coefficient in this case.[4] □

[4] In this case, the set of price-of-anarchy-minimizing tolling coefficients is not a singleton in general: any coefficient satisfying $\mathcal{L}^{\mathrm{nf}}(G, S_{\mathrm{L}}, k) \geq \mathcal{L}^{\mathrm{nf}}(G, S_{\mathrm{U}}, k)$ is optimal. Implication (7.24) means that this set always contains $k = 0$.

Proof of Theorem 7.3.4 It follows easily from the results in [142] that in Case (2) when $k \leq k^{\mathrm{gm}}$, it is true for any s that $\mathcal{L}^{\mathrm{nf}}(G, s, k) \leq \mathcal{L}^{\mathrm{nf}}(G, S_{\mathrm{U}}, k^{\mathrm{gm}})$; the price of anarchy bound for this scenario is thus precisely that in [142], where now we include games which need not have flow on every edge in an untolled Nash flow. □

7.3.6 Network-aware, mean-aware

Next, we consider when the system designer is network-aware and mean-aware to illustrate the gain in performance when the system designer has knowledge of the network and partial information of the population.

Theorem 7.3.5. *When S_{L}, S_{U} and the mean sensitivity $\overline{s}$ are known, under an optimal, network-aware, scaled marginal-cost tolling mechanism, the price of anarchy is tightly upper bounded by*

$$\mathrm{PoA}^*\left(G, \mathcal{S}(\overline{s})\right) \leq \frac{R^2 - \beta R + \beta}{\beta - \beta^2/4}, \tag{7.26}$$

where $R = (S_{\mathrm{U}} - \overline{s})/(S_{\mathrm{U}} - S_{\mathrm{L}})$ and β is the unique solution on the interval $[0, 2]$ to

$$\beta = R\left(1 + \sqrt{\frac{1 + R - \beta}{\overline{s}/S_{\mathrm{L}} + R - \beta}}\right). \tag{7.27}$$

In order to prove this, we start by making several of the same reductions as in Section 7.3.2. In this setting, the optimal network-aware toll is found and denoted by the scaling factor $k_{(\overline{s},G)}$.

Proposition 7.3.5. *For a network $G \in \mathcal{G}$ with price sensitivity distributions $s \in \mathcal{S}(\overline{s})$ with extreme sensitivity distributions $s_\ell^{(\overline{s},G,k)} = (S_{\ell 1}, S_{\ell 2})$ and* supp $= (S_{u1}, S_{u2})$*, the opti-*

mal toll scaling factor for a linear toll will take the form,

$$k_{(\bar{s},G)} = \frac{1}{\sqrt{S_{\ell 1} S_{u2}}}. \tag{7.28}$$

Proof of Proposition 7.3.5: From Lemma 7, under the same tolling mechanism, the set of Nash flows caused by $\mathcal{S}(\bar{s})$ is equal to those caused by distributions with bounds $[S_{\ell 1}, S_{u2}]$ and no mean constraint. The optimal scaling factor will therefore minimize the price of anarchy over this set of distributions. From [142], the optimal scaling factor for a linear toll will take this form. □

Lemma 9 shows that a transformation from a network $G \in \mathcal{G}$ to a network $\hat{G} \in \mathcal{G}^{\mathrm{lc}}$ will increase the price of anarchy; we also note that this transformation had no dependence on the toll scaling factor and we can thus choose a k that is optimal for the resulting network.

Corollary 4. *When making a reduction from $G \in \mathcal{G}$ to $\hat{G} \in \mathcal{G}^{\mathrm{lc}}$, the price of anarchy increases regardless of the toll scaling factor k, including when $k = k_{(\bar{s},G)}$ for each network before and after the reduction.*

Proof: In the proof of Lemma 9, the relation between k and the price of anarchy was not used; instead, it was shown that the price of anarchy increases as the network is transformed from any two link network, to one that was in $\mathcal{G}^{\mathrm{lc}}$. Consider having network G with the non-optimal toll scaling factor $k_{(\bar{s},\hat{G})}$. When the reduction from G to $\hat{G}$ is done, by Lemma 9 we have

$$\begin{aligned}\mathrm{PoA}(G, \mathcal{S}(\bar{s}), T(k_{(\bar{s},G)}) &\leq \mathrm{PoA}(G, \mathcal{S}(\bar{s}), T(k_{(\bar{s},\hat{G})}) \\ &\leq \mathrm{PoA}(\hat{G}, \mathcal{S}(\bar{s}), T(k_{(\bar{s},\hat{G})}). \qquad \square\end{aligned}$$

Proof of Theorem 7.3.5: It is shown in Lemma 10 that a set of two networks realizes the price of anarchy. The price of anarchy for the network G_β is found by (B.37). Now, let G', defined by β', be a network that has the same price of anarchy when the flow R is on the first link, One solution is clearly $\beta = \beta'$, the other is $\beta' = \frac{(4-\beta)R}{R^2-\beta R+\beta}$. Using the cost function of network G, we have $\beta = (1 + S_{\mathrm{L}} k_{(\bar{s},G)})R$. Thus, if β' satisfied supp $= (S_{\mathrm{L}}, S_{\mathrm{U}})$ for the same mean sensitivity, then

$$\beta' = \frac{(4-\beta)R}{R^2 - \beta R + \beta} = (1 + S_{\mathrm{U}} k_{(\bar{s},G')})R. \tag{7.29}$$

However, it can be shown that the right hand side of (B.38) is strictly less than the left hand side. This imposes that the flow $f_1 = R$ cannot be a Nash flow in G' under distributions in $\mathcal{S}(\bar{s})$ and therefore not achieve the same price of anarchy as G. This implies that the price of anarchy for G_β is greater than that of G_α when both are tolled optimally with respect to Proposition 7.3.5. As this network is optimally tolled, from Proposition 7.3.5, it will be the case that

$$S_{u2} = \frac{\bar{s} - S_{\mathrm{L}}}{1 + R - \beta} + S_{\mathrm{L}}. \tag{7.30}$$

Now, in sensitivity distribution $s_\ell^{(\bar{s},G,k)} = (S_{\mathrm{L}}, S_{\mathrm{U}})$, users with sensitivity S_{L} are indifferent with optimally scaled toll $k_{(\bar{s},G)}$. Using β from (B.38) and substituting the optimal scaling factor with extreme sensitivity from (7.30) leads to the characterization of β in the theorem statement, and the price of anarchy is found by substituting this into (B.37).

□

7.3.7 Network-aware, distribution-aware

In [125], the authors show that in the fully informed setting, there exist fixed tolls that will incentivize the users to self route optimally. We will extend this work to our framework to show that, when the system designer known the full user sensitivity distribution and the network characteristics, they can always design a toll that gives price of anarchy of one.

Let $F(S)$ be a cumulative distribution function of the users' sensitivities in population s. Further, let $F^{-1}(f)$ be a preimage of $[S_{\mathrm{L}}, S_{\mathrm{U}}]$ under $F(S)$ where if the preimage is non-singleton, the minimum sensitivity is used, i.e., $F^{-1}(f)$ is the sensitivity at which f mass of users have a lower sensitivity.

Theorem 7.3.6. *For the network G with population s, the linear toll*

$$T(af+b) = \frac{1}{F^{-1}(f_1^{\mathrm{opt}})} af, \tag{7.31}$$

where $f_1^{\mathrm{opt}} = \frac{2a_2+b_2-b_1}{2(a_1+a_2)}$ will have price of anarchy one, i.e., $\mathrm{PoA}^*(G,s) = 1$ *for any G and s.*

Box (G) in Fig. 7.1 shows that, when sufficient information is available, the inefficiency can be entirely eliminated, regardless of the problem instance.

Proof. First, note that for a network G, the optimal flow will be $f_1^{\mathrm{opt}} = \frac{2a_2+b_2-b_1}{2(a_1+a_2)}$. If S is the sensitivity at which a user is indifferent between the two paths in the optimal flow, then any user with lower sensitivity will use the first edge. If the indifferent sensitivity is $S^\star = F^{-1}(f_1^{\mathrm{opt}})$ then picking k such that

$$(1+S^\star k)a_1 f_1^{\mathrm{opt}} + b_1 = (1+S^\star k)a_2(1-f_1^{\mathrm{opt}}) + b_2 \tag{7.32}$$

is satisfied, the equilibrium flow will be f^{opt}. Substituting $S^\star$ and solving for k gives $k = \frac{1}{F^{-1}(f_1^{\mathrm{opt}})}$. □

7.4 Empirical Study

The theoretical claims of this work are presented in the setting of two-link, affine-latency congestion games with scaled marginal cost tolls. As mentioned previously, some of these results generalize beyond this simple class of networks; we presented each result in this reduced setting to improve uniformity in presentation. Specifically, by matching the upper bound in [35], Proposition 7.3.3 holds for all networks, even those with multiple source-destination pairs. Theorem 7.3.6 easily generalizes to multi-link parallel networks following the same steps as presented in this work. Theorem 7.3.1 and Theorem 7.3.4 can be shown to hold after a slight transformation for multi-link parallel networks by following steps from [136]. This leaves Theorem 7.3.2, Theorem 7.3.3, and Theorem 7.3.5 as results where it is not known if our findings generalize to other classes of networks. In this section, we will motivate why we believe these results offer insights more broadly and the relationships likely hold more generally.

To understand how the results of this work extend to more general networks, we focus on understanding how the derived, optimal, scaled marginal-cost tolls perform in networks of greater than two links. We do this empirically by randomly generating a large number of five-link parallel networks with different population sensitivity distributions and recording the price of anarchy. Fig. 7.4 shows the empirical price of anarchy values and demonstrates that every found parallel network has a strictly better price of anarchy than the 2-link bound.

We focus our simulation on the mean aware setting to understand how these results generalize. To garner these empirical results, we set $\overline{s} = 1$ and vary the lower and upper bound over several values. For many values of $\alpha \in (0,1)$, such that $S_{\mathrm{L}} = \overline{s} - \alpha$ and $S_{\mathrm{U}} = \overline{s} + \alpha$, we randomly generate 500 networks with five parallel edges, where the coefficients a_e and b_e in the latency function $\ell_e(f_e) = a_e f_e + b_e$ are independently

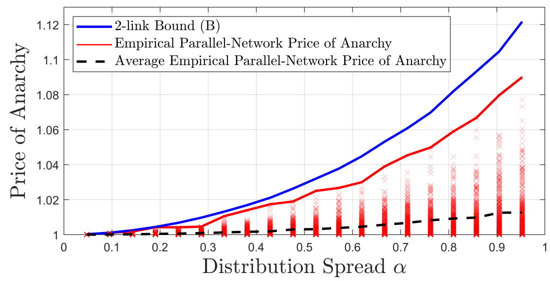

Figure 7.4: Empirical price of anarchy in 5-link parallel congestion networks compared to two-link bound. Specifically, the study is in the mean-aware, network-agnostic information setting, where in each game the optimal scaled marginal cost toll $k_{(\bar{s})}$ from Prop. 7.3.1 is used. With average sensitivity $\bar{s} = 1$, the population sensitivity distributions were chosen with lower bound $S_{\mathrm{L}} = \bar{s} - \alpha$ and upper bound $S_{\mathrm{U}} = \bar{s} + \alpha$. For varying values of $\alpha \in (0, 1)$, 500 parallel networks with 5 links and 100 population sensitivity distributions (including demonstrably worst-case distributions) were randomly generated and the price of anarchy while using $k_{(\bar{s})}$ was recorded for each realization.

drawn uniformly at random from $[0, 1]$. For each realized network, a set of sensitivity distribution were generated; the generated distributions included ones randomly created with 2, 3, 4, and 5 sensitivity values with positive weight. Fig. 7.4 shows a scatter plot of each recorded price of anarchy along with the maximum over these empirical samples.

It is not surprising that the 2-link bound appears to hold over the class of parallel networks; in [141], it is shown that the price of anarchy in non-atomic congestion games is independent of the network structure, and worst-case examples are realized by two-link networks. Though it is not obvious this relationship holds with the introduction of tolling, in [142], the authors show that scaled-marginal cost tolls similarly experience worst-case performance over parallel networks in two-link networks.

7.5 Conclusion

This work studies the value of different types of information to a toll designer. When comparing the performance guarantees awarded to toll designers with differing available information, we observe that, though possessing additional information can give a system designer greater capabilities, it is not trivial which specific pieces of information are most helpful in influencing user behavior. The results of this work offer comparisons between the value of different types of information, including when additional information is helpful and when it is not.

Chapter 8

Conclusions

This thesis presented several works on the role of information in the control of multi-agent systems. Tools from game theory are used throughout to model the interactions of multiple decision-makers in distributed autonomous and social environments. The broadest conclusions from this thesis are that 1) significant opportunities for improving the performance and capabilities of multi-agent systems are available by introducing new channels of information communication, and 2) information needs to be used and shared carefully lest system performance be inadvertently worsened.

Appendix A

Appendix Title A

A.1 Omitted proofs of Chapter 2

Proof of Proposition 2.2.1: To show existence, we can simply observe that $a^{\text{opt}} \in \arg\max_{a \in \mathcal{A}} W(a)$ is a k-strong Nash equilibrium for any $k \in [n]$. Because $W(a^{\text{opt}}) \geq W(a')$ for all $a' \in \mathcal{A}$, the global optimal satisfies $W(a^{\text{opt}}) \geq W(a'_\Gamma, a^{\text{opt}}_{-\Gamma}),\ \forall a'_\Gamma \in \mathcal{A}_\Gamma,\ \Gamma \in \mathcal{C}_{[k]}$. □

Proof of Theorem 2.3.1: The proof can be outlined by four parts: first, the problem of finding $\text{SPoA}_k(\mathcal{G}_n, w)$ is transformed and relaxed, second, the parameterization used in the proof of Proposition 2.3.2 is used to turn the relaxed problem into a linear program, next, an example is constructed to show the linear program provides a tight bound, finally, we take the dual of said linear program.

1) Relaxing the problem: Quantifying $\text{SPoA}_k(\mathcal{G}_n, w)$ can be expressed as taking the minimum k-strong price of anarchy over all games in $\mathcal{G}_n$, i.e.,

$$\min_{G \in \mathcal{G}_n} \quad \frac{\min_{a^{k\text{SNE}} \in k\text{SNE}(G)} W(a^{k\text{SNE}})}{\max_{a^{\text{opt}} \in \mathcal{A}} W(a^{\text{opt}})} \tag{D1}$$

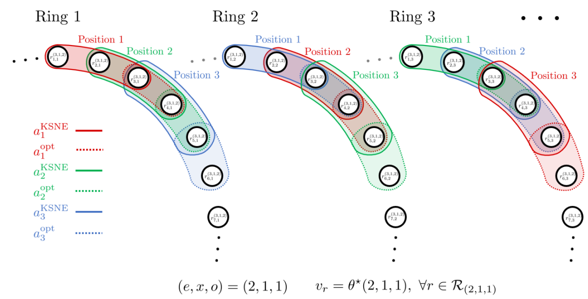

Figure A.1: Game construction for worst-case k-strong price of anarchy. Three of the n players' action sets are shown (color coded in red, green, and blue respectively) on three of $n!$ rings for the label $(e, x, o) = (2, 1, 1)$. A ring has n positions, one for each player. For a label (e, x, o) we generate $n!$ rings for all the orderings of players over positions. This is repeated for each label. Players still only have two actions, but each action covers resources from each ring. The value of a resource is equal to the value of $\theta^\star$, a solution to (D), for the label with which it is associated.

$$\begin{aligned}
\max_{\theta \in \mathbb{R}^{|\mathcal{I}|}_{\geq 0}} \quad & \sum_{e,x,o} w(o+x)\theta(e,x,o) \\
\text{s.t.} \sum_{e,x,o} & \left(\binom{n}{\zeta} w(e+x) - \sum_{\substack{0 \leq \alpha \leq e \\ 0 \leq \beta \leq o \\ \alpha+\beta \leq \zeta}} \binom{e}{\alpha}\binom{o}{\beta}\binom{n-e-o}{\zeta-\alpha-\beta} w(e+x+\beta-\alpha) \right) \theta(e,x,o) \geq 0 \\
& \qquad \forall \zeta \in \{1, \dots, k\} \\
& \sum_{e,x,o} w(e+x)\theta(e,x,o) = 1
\end{aligned} \tag{D}$$

To make this problem more approachable, we introduce several transformations and relaxations. First, rather than searching over the entire set of game $\mathcal{G}_n$, we search over the set of games $\hat{\mathcal{G}}_n$, in which each agent has exactly two actions. This reduction of the search-space can be down without loss of generality, i.e., $\text{SPoA}_k(\mathcal{G}_n, w) = \text{SPoA}_k(\hat{\mathcal{G}}_n, w)$. Trivially, $\hat{\mathcal{G}}_n \subset \mathcal{G}_n$. Further consider any game $G \in \mathcal{G}_n$; if for every player each of their actions is removed except their action in the optimal allocation a_i^{opt} and their action in their worst k-strong Nash equilibrium $a_i^{k\text{SNE}}$, the new problem will maintain the same k-strong price of anarchy, but will now exist in $\hat{\mathcal{G}}_n$. With this reduction, we will denote each player's action set as $\hat{\mathcal{A}}_i = \{a_i^{\text{opt}}, a_i^{k\text{SNE}}\}$. Second, we normalize each resource value v_r such that the equilibrium welfare is one. This too can be done without loss of generality by scaling each resource identically thus not altering the SPoA ratio. Third, we invert the objective and consider the maximization of $W(a^{\text{opt}})/W(a^{k\text{SNE}})$. Finally, we sum over each of the k-coalition equilibrium constraints. For each $\zeta \in [k]$, rather than satisfying each inequality in (2.3), sum over every combination of the ζ out of n players, denoted $\mathcal{C}_\zeta$. Applying these reductions to (D1) gives,

$$\begin{aligned} \max_{G \in \hat{\mathcal{G}}_n} \quad & W(a^{\text{opt}}) \\ \text{s.t.} \quad & \binom{n}{\zeta} W(a^{k\text{SNE}}) \geq \sum_{\Gamma \in \mathcal{C}_\zeta} W(a_\Gamma^{\text{opt}}, a_{-\Gamma}^{k\text{SNE}}), \ \forall \zeta \in [k], \\ & W(a^{k\text{SNE}}) = 1 \end{aligned} \tag{D2}$$

(D2) provides a lower bound on $\text{SPoA}_k(\mathcal{G}_n, w)$ as the feasible set was expanded. Later, we will show that the bound is tight by constructing an example that realizes it.

2) Parameterization: We use the parameterization introduced in the proof of Proposition 2.3.2 with respect to the joint actions $a = a^{k\text{SNE}}$ and $a' = a^{\text{opt}}$. By considering any $\theta \in \mathbb{R}_{\geq 0}^{|\mathcal{I}|}$, we can parameterize any game $G \in \hat{\mathcal{G}}_n$; to find the worst-case price of anarchy, we search over all such parameters, i.e., look over the entire class of games. The linear program (D) is the result of the search for the vector θ that results in the highest price of anarchy.

3) Constructing an example: Consider the following resource allocation problem: for each label $(e, x, o) \in \mathcal{I}$ and permutation of the n player $\sigma \in \Sigma_n$, define a ring of n resources. Total, there are $nn!|\mathcal{I}|$ resources. Let $r_{i,j}^{(e,x,o)}$ denote the resource with label (e, x, o) at position i in the jth ring. Consider, for instance, the $n!$ rings associated with the label $(e, x, o) = (2, 1, 1)$ as depicted in Fig. A.1. We will construct the actions $a_i^{k\text{SNE}}$ and a_i^{opt} so that for each resource in these rings, $e + x = 3$ agents have it in only their equilibrium action, and $x + o = 2$ agents have it only in their optimal action. In the first ring (with the monotonic permutation $\sigma = (1, 2, 3, \ldots, n)$), agent i has actions $a_i^{k\text{SNE}} = \{r_{i,1}^{(2,1,1)}, r_{i+1\%n,1}^{(2,1,1)}, r_{i+2\%n,1}^{(2,1,1)}\}$ and $a_i^{\text{opt}} = \{r_{i+2,1}^{(2,1,1)}, r_{i+3\%n,1}^{(2,1,1)}\}$, where % denotes the modulo operator so the selected resources wrap around the ring. This pattern continues for each ring $j \in [n!]$ with a different permutation of players $\sigma \in \Sigma_n$. At a ring with label (e, x, o) and permutation σ, player i has the actions $a_i^{k\text{SNE}} = \{r_{\sigma(i),j}^{(e,x,o)}, \ldots, r_{\sigma(i)+e+x-1\%n,j}^{(e,x,o)}\}$ and $a_i^{\text{opt}} = \{r_{\sigma(i)+e\%n,j}^{(e,x,o)}, \ldots, r_{\sigma(i)+e+x+o-1\%n,j}^{(e,x,o)}\}$. Finally, each resource of type (e, x, o) has a value $\theta(e, x, o)$ where θ is a fixed parameter. The function which encodes the welfare from player overlap is w.

In the joint action $a^{k\text{SNE}}$, each resource is covered by exactly $e + x$ agents and the system welfare can be written

$$W(a^{k\text{SNE}}) = \sum_{e,x,o} nn!\theta(e, x, o)w(e + x). \tag{A.1}$$

Similarly, joint action a^{opt} satisfies

$$W(a^{\text{opt}}) = \sum_{e,x,o} nn!\theta(e,x,o)w(o+x). \tag{A.2}$$

Now, consider a coalition $\Gamma \in \mathcal{C}_{[k]}$ and denote by ζ its cardinality. The system welfare of this group deviating their action to a^{opt}_{Γ} is

$$\begin{aligned} W(a^{\text{opt}}_{\Gamma}, a^{k\text{SNE}}_{-\Gamma}) &= \sum_{e,x,o}\sum_{j=1}^{n!}\sum_{i=1}^{n} \theta(e,x,o)w(|a^{\text{opt}}_{\Gamma}, a^{k\text{SNE}}_{-\Gamma}|_r) \\ &= \sum_{e,x,o}\theta(e,x,o) \sum_{\substack{0\le\alpha\le e\\0\le\beta\le o\\\alpha+\beta\le\zeta}} nn!\binom{e}{\alpha}\binom{o}{\beta}\binom{n-e-o}{\zeta-\alpha-\beta}w(e+x+\beta-\alpha) \end{aligned} \tag{A.3}$$

where we let r be the shorthand for $r^{(e,x,o)}_{i,j}$. The second equality holds by defining α and β as the number of players in Γ who invested in resource r exclusively in their action $a^{k\text{SNE}}$ or a^{opt} respectively. By counting arguments, there are exactly $\binom{e}{\alpha}\binom{o}{\beta}\binom{n-e-o}{\zeta-\alpha-\beta}$ positions for the players in Γ which yield the profile (α,β) for a resource at some fixed position in the ring, there are $\zeta!$ ways to order the players in Γ, $(n-\zeta)!$ ways to order the players not in Γ, and n resource in each ring.

Verifying $a^{k\text{SNE}}$ is a k-strong Nash equilibrium boils down to showing (A.1) is greater than or equal to (A.3). We can see that this holds whenever θ is a feasible point in (D). Accordingly, the k-strong price of anarchy satisfies

$$\frac{1}{Q^\star} \le \text{SPoA}_k(\mathcal{G}_n, w) \le \frac{1}{\sum_{e,x,o}\theta(e,x,o)w(o+x)}, \tag{A.4}$$

where the first inequality holds from the reductions made in part 1, and the second holds as the k-strong price of anarchy is upper bounded by any particular problem; comparing (A.1) and (A.2) gives the final expression. Letting θ take on the solution to (D) shows the bound is tight.

4) Taking the Dual: Before considering the dual program to (D), we first show that the primal is feasible. It is easy to verify the feasible set is non-empty by considering the point $\theta(1,0,0) = 1/w(1)$ and zero otherwise. Now, we must show that the feasible set is compact and thus the value of (D) is bounded. From the equality constraint, we can obtain

$$1 \geq \min_{y>0} w(y) \sum_{\substack{e,x,o \\ e+x>0}} \theta(e,x,o).$$

Because we assume $w(y) > 0$ for all $y > 0$, we show that each value of $\theta(e,x,o)$ such that $e + x > 0$ is bounded. For the remaining values of $\theta(0,0,o)$, consider the equilibrium constraint[1] when $\zeta = 1$. By rearranging terms, and observing the bounded terms from the previous argument, we observe $L \geq w(1)\sum_{o\in[n]} o\theta(0,0,o)$, where L is a bounded value. Because $w(1) > 0$, the remaining decision variables are also bounded, and thus the feasible set is finite.

Now, we find the dual program to (D). Because (D) is a linear program, we can rewrite it in the more concise form

$$\begin{aligned} \max_{\theta\in\mathbb{R}^{|\mathcal{I}|}} \quad & b^\top\theta \\ \text{s.t.} \quad & c_\zeta^\top\theta \geq 0,\ \forall\zeta\in[k] \qquad (\nu_\zeta), \\ & d^\top\theta - 1 = 0 \qquad (\rho), \\ & \theta \geq 0 \qquad (\phi) \end{aligned}$$

where $\nu \geq 0$, ρ, and $\phi \geq 0$ are the associated dual variables. The Lagrangian function is defined as $\mathcal{L}(\theta,\nu,\rho,\phi) = b^\top\theta + (\sum_{\zeta\in[k]} \nu_\zeta c_\zeta^\top\theta) - \rho(d^\top\theta - 1) + \phi^\top\theta$. Let $g(\nu,\rho,\phi) = \sup_{\theta\in\mathbb{R}^{|\mathcal{I}|}} \mathcal{L}(\theta,\nu,\rho,\phi)$ serve as an upper bound to (D). The dual program is derived by min-

[1] The $\zeta = 1$ constraint is present in (D) for all $k \geq 1$.

imizing $g(\nu, \rho, \phi)$; note that this value is only unbounded above unless $b^\top + \sum_{\zeta\in[k]} \nu_\zeta c_\zeta^\top - \rho d^\top + \phi^\top = 0$. Substituting this into the objective, and removing the free variable ϕ so that the equality constraint becomes an inequality, the dual problem becomes

$$\begin{aligned} \min_{\rho, \{\nu_\zeta \in \mathbb{R}_{\geq 0}\}_{\zeta\in[k]}} \quad & \rho \\ \text{s.t.} \quad & b^\top - \rho d^\top + \sum_{\zeta\in[k]} \nu_\zeta c_\zeta \leq 0, \end{aligned} \tag{P1}$$

From strong duality, (P1) provides the same value as (D). Expanding the hidden terms shows that (P1) is equivalent to (P[k]). □

Proof of Proposition 2.5.2: The proof is straighforward, and simply requires generalizing the constraint set of (P[k]). Consider taking the same steps as the proof of Theorem 2.3.1, but with the equilibrium constraint defined by the utility rule u. This will result in the same linear program as in (P[k]), but now with the constraint set

$$0 \geq w(o+x) - \rho w(e+x) + \sum_{\zeta\in[k]} \nu_\zeta \left(\binom{n}{\zeta} u(e+x) - \sum_{\substack{0\leq\alpha\leq e \\ 0\leq\beta\leq o \\ \alpha+\beta\leq\zeta}} \binom{e}{\alpha}\binom{o}{\beta}\binom{n-e-o}{\zeta-\alpha-\beta} u(e+x+\beta-\alpha) \right) \quad \forall (e,x,o) \in \mathcal{I}. \tag{A.5}$$

At this point, the new linear program will provide tight bounds on a specified utility rule u.

Finally, we substitute the new decision variable $u_\zeta \in \mathbb{R}^n_{\geq 0}$ into each occurrence of $\nu_\zeta u$. This enlarges the feasible set, which now subsumes all the feasible points that would evaluate a utility rule u by satisfying $u = u_\zeta$ for all $\zeta \in [k]$. As we do not enforce this constraint, the value of the final program (Q[k]) provides a lower bound on the original program, or its reciprocal provides an upper bound on the k-strong price of anarchy under the optimal utility design. □

A.2 Omitted proofs of Chapter 4

A.2.1 Proof of Proposition 4.3.1

We note that this proof follows similarly to that of [72, 105] but now with the presence of stubborn agents. Here we go through the construction of the linear program and important steps of the proof, but direct the reader to [72, 105] for a more detailed explanation. We begin by identifying the problem of characterizing the price of anarchy in a class of games $\mathcal{G}^m_{\mathcal{W},\mathcal{F}}$ for a class of games with a single basis value function w (i.e., $\mathcal{W} = \{\alpha w \mid \alpha > 0\}$) while using a local utility rule $f = \mathcal{F}(w)$. Each resource welfare function can therefore be written as $w_r(x) = v_r w(x)$, where v_r is the 'value' of that specific resource. We will discuss at the end how the solution to a single basis function can extend to the original statement. In looking for price of anarchy bounds we note that a class of resource covering problems $\mathcal{G}$ with utility rule f has the same price of anarchy as the class of problems $\mathcal{G}^*$ where each agent has exactly two actions $\mathcal{A}_i = \{a_i^{\mathrm{Ne}}, a_i^{\mathrm{opt}}\}$,[2] thus we will search for price of anarchy bounds in these two-action games and note they hold more generally. The price of anarchy over $\mathcal{G}^m_{\mathcal{W},\mathcal{F}}$ while utilizing utility rule f, $\mathrm{PoA}(\mathcal{G}^m_{\mathcal{W},\mathcal{F}})$,

[2]This can be seen by transforming each game $G \in \mathcal{G}$ into one with two actions by removing all actions but the worst equilibrium a^{Ne} and the optimal allocation a^{opt}. Because a^{Ne} remains a Nash equilibrium, the price of anarchy is unchanged.

can be written as

$$\min_{G \in \mathcal{G}^m_{\mathcal{W},\mathcal{F}}} \quad \frac{W(a^{\mathrm{Ne}})}{W(a^{\mathrm{opt}})} \tag{A.6}$$
$$\text{s.t. } U_i(a^{\mathrm{Ne}}; d) \geq U_i(a_i^{\mathrm{opt}}, a_{-i}^{\mathrm{Ne}}; d) \quad \forall i \in N.$$

Where $G = (N, \mathcal{A}, \{U_i\}_{i \in N}, W)$ encodes all of the information about a problem instance. This program is not efficient to solve in general, however, we will make use of a parameterization that will greatly ease the computation of the price of anarchy. First, we will modify (A.6) by normalizing $W(a^{\mathrm{Ne}}) = 1$, which can be done by homogeneously scaling each resource value and will not alter the problems price of anarchy. Next, we relax the equilibrium constraint from holding for every agent $i \in N$ to only hold as a summation over all agents, i.e., $\sum_{i \in N} U_i(a^{\mathrm{Ne}}; d) - U_i(a_i^{\mathrm{opt}}, a_{-i}^{\mathrm{Ne}}; d) \geq 0$. Note, that this relaxation will cause the new program to provide a lower bound for the original, however we will show that this bound is tight. Finally, we take the reciprocal of the objective and turn the minimization problem into a maximization problem. The new program, which the solution of will be a lower bound for the original, can be written

$$\max_{G \in \mathcal{G}^m_{\mathcal{W},\mathcal{F}}} \quad W(a^{\mathrm{opt}}) \tag{A.7}$$
$$\text{s.t. } \sum_{i \in N} U_i(a^{\mathrm{Ne}}; d) - U_i(a_i^{\mathrm{opt}}, a_{-i}^{\mathrm{Ne}}; d) \geq 0,$$
$$W(a^{\mathrm{Ne}}) = 1.$$

Now, we make use of a parameterization that was also described in the proof of Theorem 4.4.1. Each resource is given a label (x_r, y_r, z_r, d_r) defined by $x_r = |a^{\mathrm{Ne}} \setminus a^{\mathrm{opt}}|_r$, $z_r = |a^{\mathrm{opt}} \setminus a^{\mathrm{Ne}}|_r$, $y_r = |a^{\mathrm{opt}} \cap a^{\mathrm{Ne}}|_r$, and d_r is the number of stubborn agents. The label denotes the number of agents that utilize a resource in only their Nash action x_r, only their optimal action z_r, or both y_r. The set of all such labels is $\mathcal{I}_n = \{(x,y,z) \in \mathbb{N}^3_{\geq 0} \mid 1 \leq x+y+z \leq n\}$. For each label we define a parameter $\theta(x,y,z,d) = \sum_{r \in \mathcal{R}(x,y,z,d)} v_r$, where $\mathcal{R}(x,y,z,d)$ is the set of resources with label (x,y,z,d). We can express several quantities using this paramerization as follows:

$$\begin{aligned}
\sum_{i \in N} U_i(a^{\mathrm{Ne}}; d) &= \sum_{r \in \mathcal{R}} v_r[(x_r + y_r) f(x_r + y_r + d_r)] \\
&= \sum_{x,y,z,d} (x+y) f(x+y+d) \theta(x,y,z,d), \\
\sum_{i \in N} U_i(a_i^{\mathrm{opt}}, a_{-i}^{\mathrm{Ne}}; d) &= \sum_{x,y,z,d} [y f(x+y+d) + z f(x+y+d+1)] \theta(x,y,z,d), \\
W(a^{\mathrm{Ne}}) &= \sum_{x,y,z,d} w(x+y) \theta(x,y,z,d), \\
W(a^{\mathrm{opt}}) &= \sum_{x,y,z,d} w(z+y) \theta(x,y,z,d).
\end{aligned}$$

Note that we write the sum over all labels in $\mathcal{I}_n$ as $\sum_{x,y,z,d}$ for brevity. Rewriting (A.7) using this parameterization gives

$$\begin{aligned}
p^\star = \max_{\theta \in \mathbb{R}^{|\mathcal{I}_n|}} \quad & \sum_{x,y,z,d} w(z+y) \theta(x,y,z,d) \\
\text{s.t.} \quad & \sum_{x,y,z,d} [x f(x+y+d) - z f(x+y+d+1)] \theta(x,y,z,d) \geq 0, \\
& \sum_{x,y,z,d} w(x+y) \theta(x,y,z,d) = 1, \\
& \theta \geq 0.
\end{aligned} \tag{A.8}$$

As discussed when introducing (A.7), $p^\star$ offers a lower bound on the price of anarchy. We further show that using the solution to (A.8), $\theta^\star$, one can construct a game whose price of anarchy is a tight upper bound.

For each label (x,y,z,d) such that $\theta^\star(x,y,z,d) > 0$, introduce n resources each with value $\theta^\star(x,y,z,d)/n$. As in Fig. A.2, define each players action set to cover $x+y$ of these resources in their equilibrium action $a_i^{\rm Ne}$ and $z+y$ of these resources in their optimal action $a_i^{\rm opt}$ where y of the resources are in both actions. By considering the n resources in a ring, and offsetting each agents action sets by one resource, each agent can experience this set of resources symmetrically. Finally, let d stubborn agents be placed on each of these resources. If this is repeated for each label, then one can observe that player i will have utility

$$\begin{aligned} U_i(a^{\rm Ne};d) &= \sum_{x,y,z,d} (x+y) f(x+y+d)\theta(x,y,z,d) \\ &\geq \sum_{x,y,z,d} [y f(x+y+d) + z f(x+y+d+1)]\theta(x,y,z,d) \\ &= U_i(a_i^{\rm opt}, a_{-i}^{\rm Ne}; d), \end{aligned}$$

where the inequality holds from the constraint in (A.8) and $\theta^\star$ being a feasible solution; thus $a^{\rm Ne}$ is a Nash equilibria, and the price of anarchy of this game is at most $\frac{W(a^{\rm Ne})}{W(a^{\rm opt})} = \frac{\sum_{x,y,z,d} w(x+y)\theta^\star(x,y,z,d)}{\sum_{x,y,z,d} w(z+y)\theta^\star(x,y,z,d)} = \frac{1}{\sum_{x,y,z,d} w(z+y)\theta^\star(x,y,z,d)} = \frac{1}{p^\star}$, where the second equality holds from the constraint in (A.8). The constructed game therefore offers an upper bound on the price of anarchy of $1/p^\star$, the solution to (A.8), offers a matching lower bound, proving the bound is tight.

Notice that (A.8) is a linear program with decision variable θ. Next we find the dual

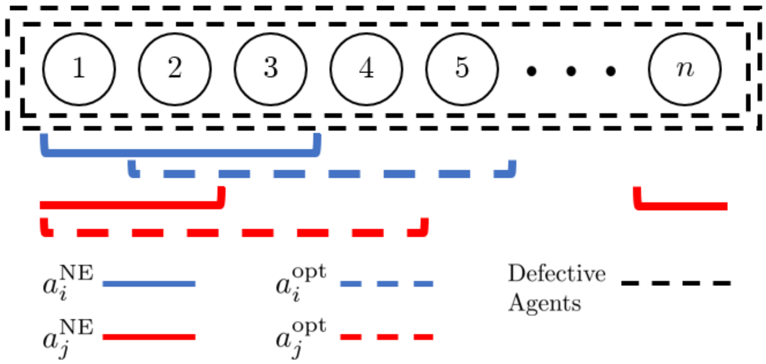

Figure A.2: Game construction for resource allocation problems utilizing the solution to (A.6). For each tuple (x, y, z, d), n resources are created with value $\theta^\star(x, y, z, d)/n$. For a resource with label (x, y, z, d), design the action set of agent i to utilize the first $x + y$ of these resources in their first action, a_i^{Ne}, and the $x + 1$ to $x + y + z$ resources in their other action a_i^{opt}. For the proceeding agent, follow the same process but increasing the index of the starting resource by 1. If the agent were to use the non-existent $n + 1$ or greater resource, start the assignment from the beginning, essentially forming a ring. Once each action set is assigned for all n agents, each resource will be used by $x + y$ agents in the action a^{Ne} and $y + z$ agents in a^{opt}, matching the label it was assigned. We can observe that a^{Ne} is a Nash equilibrium from the constraints in (A.6). For a similar, but more detailed explanation, see [72].

of (A.8) to be

$$\begin{aligned} d^\star = \min_{\lambda \geq 0, \mu \in \mathbb{R}_{\geq 0}} \quad & \mu \\ \text{s.t. } & w(z+y) - \mu w(x+y) \\ & + \lambda[xf(x+y+d) - zf(x+y+d+1)] \leq 0 \\ & \forall (x,y,z) \in \mathcal{I}_n, \ d \in \{1, \ldots, m\}. \end{aligned} \tag{A.9}$$

Because, (A.8) is a linear program, and thus convex, by the principle of strong duality, $d^\star = p^\star = \mathrm{PoA}(\mathcal{G}^m_{\mathcal{W},\mathcal{F}})^{-1}$. Finally, to optimize the price of anarchy over local utility rules, we need only minimize (A.9) over $f : [n + m] \to \mathbb{R}_{\geq 0}$, which can be treated as a vector in $\mathbb{R}^{n+m}$. Allowing f to be a decision variable in (A.9) would cause each constraint to be bilinear in $f(i)$ and λ; however, every occurrence of λ is multiplied by an $f(i)$ for some

i and vice versa, therefore the two decision variables can be combined into one giving a program of the form (4.9).

Finally, we note that (i) an optimal utility rule can be composed as the optimal utility rule for each basis function, i.e., for a resource with value $w_r = \sum_{b=1}^{B} \alpha_b w_b$ for some $\{\alpha_b\}_{b=1}^{B}$, then $f_r^{\text{opt}} = \sum_{b=1}^{B} \alpha_b f_b^{\text{opt}}$ where f_b^{opt} is the optimal utility rule for the basis function w_b described prior, and (ii) the worst case price of anarchy over the set of games with resource value functions in $\mathcal{W} = \{\sum_{b=1}^{B} \alpha_b w_b | \alpha_b \geq 0 \ \forall b \in [B]\}$, is equal to the maximum of the sets of games with just one of these basis functions. These two observations have been shown in [7, 105] and follow identically here. This gives the final form of the optimal local utility design and the associated performance guarantee. □

A.2.2 Proof of Theorem 4.4.1

In this appendix, we give the full proof of Theorem 4.4.1 as well as several supporting lemmas. As in the proof of Proposition 4.3.1, we restrict our search to games where each player has two actions $\mathcal{A} = \{a_i^{\text{Ne}}, a_i^{\text{opt}}\}$ and note that the price of anarchy over this class is the same as the original with larger agent action sets. The price of anarchy bounds in (4.13) and (4.14) are tight along the Pareto-optimal frontier. To prove that each is an upper bound, we will make use of several examples; three structures of parameterized problem instances are shown in Fig. A.3. To show that these are lower bounds, we will make use of smoothness inequalities introduced in [70]. If, given a utility rule f, each Nash equilubria $a^{\text{Ne}} \in \text{NE}(G_f)$ satisfies

$$W(a^{\text{Ne}}) \geq \lambda \cdot W(a^{\text{opt}}) + \mu \cdot W(a^{\text{Ne}}), \tag{A.10}$$

for some $\lambda, \mu \in \mathbb{R}$, then the price of anarchy will satisfy $\text{PoA}(\mathcal{G}_f) \geq \frac{\lambda}{1-\mu}$. We will provide lower bounds by finding values of λ and μ for different settings (e.g., with and without

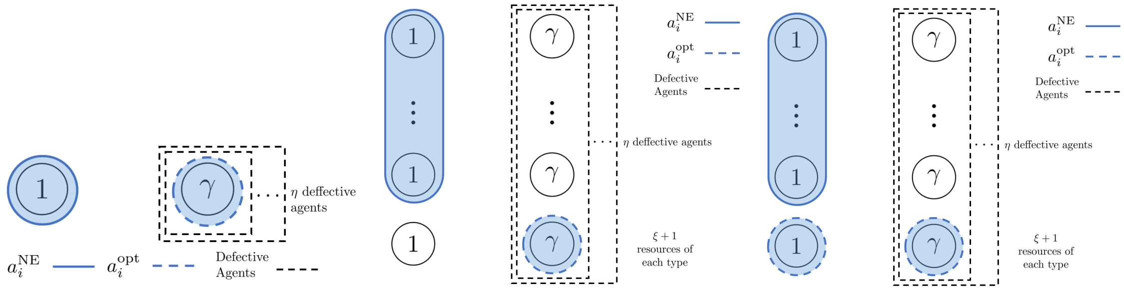

Figure A.3: (Left) Example A: G^A
A problem instance with one agent having two choices: a resource with value one and a resource with value γ covered by η defective agents. When $\gamma \leq 1/f(\eta+1)$, the agent may pick the resource of value one in equilibrium leading to $\mathrm{PoA}(G^A_f) = \frac{1}{\gamma} \geq f(\eta+1)$.
(Center) Example B: G^B
A problem with $\xi+1$ agents each with two choices: selecting ξ resources of value 1, a_i^{Ne}, or one resource of value γ with η defective agents, a_i^{opt}. The agents' equilibrium and optimal actions are distinct from one another, implying in the allocation a^{Ne}, xi agents cover each resource of value 1, and in a^{opt}, each resource of value γ is covered by one agent. When $\gamma \leq \frac{\xi f(\xi)}{f(\eta+1)}$, a^{Ne} is an equilibrium allocation with $\mathrm{PoA}(G^B_f) = \frac{1}{\gamma}$.
(Right) Example C: G^C
A problem with $\xi+1$ agents each with two choices: selecting ξ resources of value 1, a_i^{Ne}, or the remaining resource of value 1 and one resource of value γ with η defective agents, a_i^{opt}. The agents' equilibrium and optimal actions are distinct from one another, implying in the allocation a^{Ne}, ξ agents cover each resource of value 1, and in a^{opt}, every resource is covered by one agent. When $\gamma \leq \frac{\xi f(\xi) - f(\xi+1)}{f(\eta+1)}$, a^{Ne} is an equilibrium allocation with $\mathrm{PoA}(G^C_f) = \frac{1}{1+\gamma}$.

stubborn agents); often, to do so, we will utilize the fact that the welfare of a Nash equilibria can be lower bounded by

$$W(a^{\mathrm{Ne}}) \geq \sum_{i \in N} u_i(a_i^{\mathrm{opt}}, a_{-i}^{\mathrm{Ne}}) - \sum_{i \in N} u_i(a^{\mathrm{Ne}}) + W(a^{\mathrm{Ne}}), \tag{A.11}$$

which holds from the definition of a Nash equilibrium (4.1) where $u_i(a_i^{\mathrm{opt}}, a_{-i}^{\mathrm{Ne}}) \leq u_i(a^{\mathrm{Ne}})$ for all $i \in N$, implying $\sum_{i \in N} u_i(a_i^{\mathrm{opt}}, a_{-i}^{\mathrm{Ne}}) - \sum_{i \in N} u_i(a^{\mathrm{Ne}}) \leq 0$. Additionally, using the prameterization discussed in Section 4.3, where, in an allocation $(a, \overline{a})$, each resource $r \in \mathcal{R}$ is given a label (x_r, y_r, z_r, d_r) defined by $x_r = |a^{\mathrm{Ne}} \backslash a^{\mathrm{opt}}|_r$, $z_r = |a^{\mathrm{opt}} \backslash a^{\mathrm{Ne}}|_r$, $y_r = |a^{\mathrm{opt}} \cap a^{\mathrm{Ne}}|_r$, and d_r is the number of stubborn agents, where for two joint actions $a, a' \in \mathcal{A}$,

$|a\backslash a'|_r$ is the number of agents that utilize resource r in action a but not a' and $|a\cap a'|_r$ is the number of agents that utilize resource r in both a and a'. This parameterization allows us to write $W(a^{\mathrm{Ne}}) = \sum_{r\in\mathcal{R}} v_r \mathbb{1}_{[]}[x_r + y_r]$ and $W(a^{\mathrm{opt}}) = \sum_{r\in\mathcal{R}} v_r \mathbb{1}_{[]}[y_r + z_r]$; additionally, (A.11) can be rewritten as

$$W(a^{\mathrm{Ne}}) \geq \sum_{r\in\mathcal{R}} v_r[z_r f(x_r + y_r + d_r + 1) \\ - x_r f(x_r + y_r + d_r) + \mathbb{1}_{[]}[x_r + y_r]], \tag{A.12}$$

where the welfare function $w(x) = \mathbb{1}_{[]}[x]$ is the indicator function that the argument is greater than zero in covering games. Manipulating the right hand side of (A.12) into the form of (A.10) will be the primary method of lower bounding the price of anarchy of a utility rule f in a class of games.

From [36], the optimal utility rule in covering games with no stubborn agents and arbitrarily many regular agents $\mathcal{G}^0$ is

$$f^0(j) := (j-1)!\frac{e - \sum_{i=0}^{j-1}\frac{1}{i!}}{e-1} \;\forall j \geq 1, \tag{A.13}$$

and $f^0(0) = 0$. This can also be seen by taking n to infinity in (4.12). The performance guarantee of f^0 is $\mathrm{PoA}(\mathcal{G}^0_{f^0}) = 1 - \frac{1}{e}$, which can be seen from the following lemma.

Lemma 11 (Gairing 2009 [36])**.** *In the class of problems $\mathcal{G}^0$, with utility rule*

$$f^0(j) = (j-1)!\frac{e - \sum_{i=0}^{j-1}\frac{1}{i!}}{e-1} \;\forall j \geq 1, \tag{A.14}$$

(A.10) *is satisfied with $\lambda = 1$ and $\mu = -1/(e-1)$.*

This utility rule is useful in constructing the optimal utility rules in the setting with stubborn agents. Additionally, the following claim is useful in proving several lower-bounds.

Lemma 12. *The local utility rule f^0 defined in* (A.13) *satisfies*

$$jf^0(j) - f^0(j+1) = \frac{1}{e-1} \quad \forall\, j \in \mathbb{N}. \tag{A.15}$$

Proof. The claim can be proven directly by substitution:

$$\begin{aligned} jf^0(j) - f^0(j+1) &= j\left((j-1)!\frac{e - \sum_{i=0}^{j-1} \frac{1}{i!}}{e-1}\right) - (j)!\frac{e - \sum_{i=0}^{j} \frac{1}{i!}}{e-1} \\ &= \frac{j!}{e-1}\left(\sum_{i=0}^{j} \frac{1}{i!} - \sum_{i=0}^{j-1} \frac{1}{i!}\right) = \frac{1}{e-1}. \end{aligned}$$

□

The following several lemmas will define and quantify the smoothness coefficients of some useful local utility rules.

Lemma 13. *In the class of problems $\mathcal{G}^m$, with utility rule*

$$\overline{f}^m(j) = \begin{cases} \frac{m}{j} f^0(m), & \text{if } j \in \{1, \dots, m\} \\ f^0(j), & \text{otherwise,} \end{cases} \tag{A.16}$$

(A.10) *is satisfied with $\lambda = f^0(m+1)$ and $\mu = 1 - mf^0(m)$.*

Proof. Let $\mathcal{R}_a \subset \mathcal{R}$ be the set of all resources where $x_r + y_r + d_r \geq m+1$ and let $\mathcal{R}_b \subset \mathcal{R}$ be the set of all resources where $x_r + y_r + d_r \leq m$, forming a partition of $\mathcal{R}$.

For the resources in $\mathcal{R}_a$,

$$\sum_{r\in\mathcal{R}_a} v_r[z_r f^0(x_r+y_r+d_r+1) - x_r f^0(x_r+y_r+d_r)+1]$$

$$\geq \sum_{r\in\mathcal{R}_a} v_r[(z_r+y_r)f^0(x_r+y_r+d_r+1) - (x_r+y_r+d_r)f^0(x_r+y_r+d_r)+1] \tag{A.17a}$$

$$\geq \sum_{r\in\mathcal{R}_a} v_r[f^0(x_r+y_r+d_r+1)\mathbb{1}_{[]}[z_r+y_r] - (x_r+y_r+d_r)f^0(x_r+y_r+d_r)+1] \tag{A.17b}$$

$$\geq \sum_{r\in\mathcal{R}_a} v_r[f^0(x_r+y_r+d_r+1)\mathbb{1}_{[]}[z_r+y_r] - f^0(x_r+y_r+d_r+1) - \frac{1}{e-1}+1] \tag{A.17c}$$

$$\geq \sum_{r\in\mathcal{R}_a} v_r[f^0(m+1)\mathbb{1}_{[]}[z_r+y_r] - f^0(m+1) - \frac{1}{e-1}+1] \tag{A.17d}$$

$$= \sum_{r\in\mathcal{R}_a} v_r[f^0(m+1)\mathbb{1}_{[]}[z_r+y_r] + (1-mf^0(m))\mathbb{1}_{[]}[x_r+y_r]], \tag{A.17e}$$

where (A.17a) and (A.17d) hold from f^0 decreasing, (A.17b) holds from f^0 positive, and (A.17c) and (A.17e) hold from Lemma 12.

For the resources in $\mathcal{R}_b$,

$$\sum_{r\in\mathcal{R}_b} v_r[z_r\overline{f}^m(x_r+y_r+d_r+1) - x_r\overline{f}^m(x_r+y_r+d_r) + \mathbb{1}_{[]}[x_r+y_r]]$$

$$= \sum_{r\in\mathcal{R}_b} v_r[\frac{z_r}{x_r+y_r+d_r+1}(mf^0(m)) - \frac{x_r}{x_r+y_r+d_r}mf^0(m) + \mathbb{1}_{[]}[x_r+y_r]] \tag{A.18a}$$

$$\geq \sum_{r\in\mathcal{R}_b} v_r[(z_r+y_r)f^0(m+1) - \frac{x_r}{x_r+y_r+d_r}(f^0(m+1)+\frac{1}{e-1}) + \mathbb{1}_{[]}[x_r+y_r]] \tag{A.18b}$$

$$\geq \sum_{r\in\mathcal{R}_b} v_r[f^0(m+1)\mathbb{1}_{[]}[z_r+y_r] + (1-f^0(m+1)-\frac{1}{e-1})\mathbb{1}_{[]}[x_r+y_r]] \tag{A.18c}$$

$$= \sum_{r\in\mathcal{R}_b} v_r[f^0(m+1)\mathbb{1}_{[]}[z_r+y_r] + (1-mf^0(m))\mathbb{1}_{[]}[x_r+y_r]]. \tag{A.18d}$$

where (A.18b) holds from Lemma 12 and (A.18c) holds from $x_r/(x_r+y_r+d_r) \leq 1$, providing the same lower bound for the price of anarchy.

It follows that $\lambda = f^0(m+1)$ and $\mu = 1 - mf^0(m)$ satisfy (A.10). ☐

Lemma 14. *In the class of problems $\mathcal{G}^0$, with utility rule*

$$\overline{f}^m(j) = \begin{cases} \frac{m}{j}f^0(m), & \text{if } j \in \{1,\ldots,m\} \\ f^0(j), & \text{otherwise,} \end{cases} \tag{A.19}$$

(A.10) *is satisfied with $\lambda = mf^0(m)$ and $\mu = \frac{e-2}{e-1} - mf^0(m)$.*

Proof. Let $\mathcal{R}_c \subset \mathcal{R}$ denote the set of resources where $x_r > 0$ or $y_r > 0$, and let $\mathcal{R}_d \subset \mathcal{R}$ be the set of resources where $x_r = y_r = 0$. First recall the bound from (A.17e) and

(A.18d) that together give

$$\sum_{r\in\mathcal{R}} v_r[z_r\overline{f}^m(x_r+y_r+1) - x_r\overline{f}^m(x_r+y_r) + \mathbb{1}_{[\cdot]}[x_r+y_r]] \tag{A.20}$$

$$\geq \sum_{r\in\mathcal{R}} v_r[f^0(m+1)\mathbb{1}_{[\cdot]}[z_r+y_r] + (1-mf^0(m))\mathbb{1}_{[\cdot]}[x_r+y_r]], \tag{A.21}$$

in the special case where $d_r = 0$, as is the case for games the class $\mathcal{G}^0$. For the set $\mathcal{R}_c$,

$$\sum_{r\in\mathcal{R}_c} v_r[f^0(m+1)\mathbb{1}_{[\cdot]}[z_r+y_r] + (1-mf^0(m))\mathbb{1}_{[\cdot]}[x_r+y_r]]$$

$$= \sum_{r\in\mathcal{R}_c} v_r[\left(mf^0(m)\mathbb{1}_{[\cdot]}[z_r+y_r] + \frac{1}{e-1}\right) + (1-mf^0(m))\mathbb{1}_{[\cdot]}[x_r+y_r]] \tag{A.22a}$$

$$\geq \sum_{r\in\mathcal{R}_c} v_r[(mf^0(m))\mathbb{1}_{[\cdot]}[z_r+y_r] + \frac{1}{e-1} + (1-mf^0(m))\mathbb{1}_{[\cdot]}[x_r+y_r]] \tag{A.22b}$$

$$= \sum_{r\in\mathcal{R}_c} v_r[(mf^0(m))\mathbb{1}_{[\cdot]}[z_r+y_r] + \left(\frac{e-2}{e-1} - mf^0(m)\right)\mathbb{1}_{[\cdot]}[x_r+y_r]], \tag{A.22c}$$

where (A.22a) holds from Lemma 12, (A.22b) holds from $\mathbb{1}_{[\cdot]}[x] \leq x$ for all non-negative integer x, and (A.22c) holds from definition of $\mathcal{R}_c$ that $\mathbb{1}_{[\cdot]}[x_r+y_r] = 1$. For the remaining resources in $\mathcal{R}_d$,

$$\sum_{r\in\mathcal{R}_d} v_r[z_r\overline{f}^m(x_r+y_r+1) - x_r\overline{f}^m(x_r+y_r) + \mathbb{1}_{[\cdot]}[x_r+y_r]]$$

$$= \sum_{r\in\mathcal{R}_d} v_r[z_r f^0(1)] \geq \sum_{r\in\mathcal{R}_d} v_r[z_r m f^0(m)] \tag{A.23a}$$

$$\geq \sum_{r\in\mathcal{R}_d} v_r[(mf^0(m))\mathbb{1}_{[\cdot]}[z_r+y_r] + \left(\frac{e-2}{e-1} - mf^0(m)\right)\mathbb{1}_{[\cdot]}[x_r+y_r]], \tag{A.23b}$$

where (A.23a) holds from the definition of $\overline{f}^m$ and $mf^0(m) < f^0(1) = 1$, and (A.23b) holds from f^0 positive and $x_r = y_r = 0$. From (A.22c) and (A.23b), $\lambda = mf^0(m)$ and $\mu = \frac{e-2}{e-1} - mf^0(m)$ satisfy (A.10). □

Lemma 15. *In the class of problems $\mathcal{G}^m$, with utility rule*

$$f^0(j) = (j-1)! \frac{e - \sum_{i=0}^{j-1} \frac{1}{i!}}{e-1} \ \forall j \geq 1, \tag{A.24}$$

(A.10) *is satisfied with $\lambda = f^0(m+1)$ and $\mu = 0$.*

Proof. As in Lemma 13, let $\mathcal{R}_a \subset \mathcal{R}$ be the set of all resources where $x_r + y_r + d_r \geq m+1$ and let $\mathcal{R}_b \subset \mathcal{R}$ be the set of all resources where $x_r + y_r + d_r \leq m$, forming a partition of $\mathcal{R}$. For the resources in the set $\mathcal{R}_a$, follow the steps of (A.17a)-(A.17e) and note that $jf^0(j) \leq 1$ for all j, therefore (A.17e) is further lower-bounded by

$$\sum_{r \in \mathcal{R}_a} v_r[\mathbb{1}_{[]}[z_r + y_r] f^0(m+1)]. \tag{A.25}$$

For the resources in $\mathcal{R}_b$,

$$\begin{aligned}
&\sum_{r \in \mathcal{R}_b} v_r[z_r f^0(x_r + y_r + d_r + 1) - x_r f^0(x_r + y_r + d_r) + \mathbb{1}_{[]}[x_r + y_r]] \\
&\quad = \sum_{r \in \mathcal{R}_b} v_r[(z_r + y_r) f^0(x_r + y_r + d_r + 1) \\
&\qquad\qquad - (x_r + y_r) f^0(x_r + y_r + d_r) + \mathbb{1}_{[]}[x_r + y_r]] && \text{(A.26a)} \\
&\quad \geq \sum_{r \in \mathcal{R}_b} v_r[(z_r + y_r) f^0(x_r + y_r + d_r + 1)] && \text{(A.26b)} \\
&\quad \geq \sum_{r \in \mathcal{R}_b} v_r[\mathbb{1}_{[]}[z_r + y_r] f^0(m+1)], && \text{(A.26c)}
\end{aligned}$$

where (A.26a) holds from f^0 decreasing, (A.26b) holds from $jf^0(j) \leq 1$ for all $j \geq 0$, and (A.26c) holds from $\mathbb{1}_{[]}[x] \leq x$ for all non-negative integer x. From (A.25) and (A.26c), $\lambda = f^0(m+1)$ and $\mu = 0$ satisfy (A.10). $\square$

Proof of Theorem 4.4.1: To prove that the curve defined by (4.13) and (4.14) represent a Pareto-optimal frontier of the multi-criterion problem of minimizing $\text{PoA}(\mathcal{G}_f^m)$ and

PoA($\mathcal{G}^0_f$), we first give a parameterized utility rule that draws the curve then show a tight lower and upper bound on it's price of anarchy, and finally show this utility rule is indeed Pareto-optimal. Let $f^t(j) = t\overline{f}^m(j) + (1-t)f^0(j)$ for some $t \in [0,1]$, be a local utility rule parameterized by $t \in [0,1]$. Through some rearanging, this is equivalent to (4.15). We will show the price of anarchy guarantees of this utility rule draw the Pareto-optimal frontier.

Part 1: Upper Bound We will give problem instances that upper bound the price of anarchy over the set $\mathcal{G}^m$ and $\mathcal{G}^0$ for the utility rule f^t. For the nominal price of anarchy, let $G^C \in \mathcal{G}^0$ be a covering game as described in Fig. A.3 (right) with $\eta = 0$, $\gamma = \frac{\xi f(\xi) - f(\xi+1)}{f(1)}$. By selecting $\xi \geq m+1$ agents in the game (where m is the number of defective agents for which f^t is designed), from Lemma 12

$$\gamma = \frac{1}{(e-1)f^t(1)} = \frac{1}{(e-1)(tmf^0(m) + (1-t))}.$$

Defining $\Gamma_m = mf^0(m) - 1 = m!\frac{e - \sum_{i=0}^{m-1}\frac{1}{i!}}{e-1} - 1$, the price of anarchy of the described game is

$$\text{PoA}(G^C_{f^t}) = \frac{1}{1+\gamma} = \frac{(e-1)(1+t\Gamma_m)}{1+(e-1)(1+t\Gamma_m)}.$$

Because $G^C \in \mathcal{G}^0$, $\text{PoA}(\mathcal{G}^0_{f^t}) \leq \text{PoA}(G^C_{f^t})$. For the price of anarchy in the perturbed agent setting, let $G^A \in \mathcal{G}^m$ be a covering game as described in Fig. A.3 (left) with $\eta = m$ and $\gamma = f^t(1)/f^t(m+1)$. From the definition of f^t and Lemma 12, the price of anarchy of this game, with utility rule f^t is

$$\text{PoA}(G^A_{f^t}) = \frac{f^0(m+1)}{1+t(mf^0(m)-1)} = \frac{\Gamma_m + \frac{e}{e-1}}{1+t\Gamma_m}.$$

Because $G^A \in \mathcal{G}^m$, $\text{PoA}(\mathcal{G}^m_{f^t}) \leq \text{PoA}(G^A_{f^t})$. This provides our upper bounds for the price of anarchy over $\mathcal{G}^0$ and $\mathcal{G}^m$ while using the utility rule f^t.

Part 2: Lower Bound To lower bound the price of anarchy, we again look for coefficients λ, μ that satisfy (A.10). From the definition of f^t, (A.12) can be rewritten

$$\begin{aligned} W(a^{\text{Ne}}) \geq \sum_{r \in \mathcal{R}} & tv_r[z_r \overline{f}^m(x_r + y_r + d_r + 1) \\ & \quad - x_r \overline{f}^m(x_r + y_r + d_r) + \mathbb{1}_{[]}[x_r + y_r]] \\ & + (1-t)v_r[z_r f^0(x_r + y_r + d_r + 1) \\ & \quad - x_r f^0(x_r + y_r + d_r) + \mathbb{1}_{[]}[x_r + y_r]], \end{aligned} \tag{A.27}$$

where $\overline{f}^m$ is as defined in (A.16). For any game in $\mathcal{G}^m$, from Lemma 13 and Lemma 15, (A.27) can be lower bounded by

$$f^0(m+1) \cdot W(a^{\text{opt}}) + t(1 - mf^0(m)) \cdot W(a^{\text{Ne}}),$$

producing for the lower bound on the price of anarchy of

$$\text{PoA}(\mathcal{G}^m_{f^t}) \geq \frac{f^0(m+1)}{t(1 - mf^0(m))} = \frac{\Gamma_m + \frac{e}{e-1}}{1 + t\Gamma_m}. \tag{A.28}$$

For the price of anarchy over the nominal setting $\mathcal{G}^0$ with utility law f^t, (4.14) needs to be lower bounded for the case where $d_r = 0$ for all $r \in \mathcal{R}$. From Lemma 14 and Lemma 11, this lower bound is

$$\begin{aligned} & (tmf^0(m) + (1-t))W(a^{\text{opt}}) \\ & \qquad + \left(t\left(\frac{e-2}{e-1} - mf^0(m) \right) + (1-t)\frac{-1}{e-1} \right) W(a^{\text{Ne}}). \end{aligned}$$

This gives a lower bound on the nominal price of anarchy while using f^t of

$$\begin{aligned}\mathrm{PoA}(\mathcal{G}^0_{f^t}) &\geq \frac{tmf^0(m) + (1-t)}{1 - \left(t\left(\frac{e-2}{e-1} - mf^0(m)\right) + (1-t)\frac{-1}{e-1}\right)} \\ &= \frac{(e-1)(1+t\Gamma_m)}{1+(e-1)(1+t\Gamma_m)}.\end{aligned}$$

Part 3: Pareto-Optimality Consider a local utility rule f with nominal price of anarchy guarantee

$$\mathrm{PoA}(\mathcal{G}^0_f) > x \tag{A.29}$$

for some $x \in [0,1]$. Consider a game $G^C \in \mathcal{G}^0$ following Fig. A.3 (right) where $\eta = 0$ and $\xi = m+1$. If $\gamma = ((m+1)f(m) - f(m+2))/f(1)$, then

$$\mathrm{PoA}(G^C_f) = \frac{1}{1 + \frac{1}{(e-1)f(1)}},$$

from the assumption that $f(j) = f^0(j)\ \forall j \geq m+1$ and Lemma 12. To satisfy the price of anarchy guarantee in (A.29),

$$f(1) > \frac{x}{(e-1)(1-x)}. \tag{A.30}$$

Now, consider the game $G^A \in \mathcal{G}^m$ described by Fig. A.3 (left) where $\eta = m$ and $\gamma = f(1)/f(m+1) = f(1)/f^0(m+1)$. The price of anarchy of this game is $\mathrm{PoA}(G^A_f) = 1/\gamma$. From (A.30),

$$\mathrm{PoA}(G^A_f) < \frac{(e-1)f^0(m+1)(1-x)}{x}. \tag{A.31}$$

In (A.31), choose $x = \frac{(e-1)(1+t\Gamma_m)}{1+(e-1)(1+t\Gamma_m)}$ for some $t \in [0,1]$ and

$$\mathrm{PoA}(\mathcal{G}^m_f) \leq \mathrm{PoA}(G^A_f) < \frac{\Gamma_m + \frac{e}{e-1}}{1+t\Gamma_m} \tag{A.32}$$

from the fact $\Gamma_m = f^0(m+1)+\frac{1}{e-1}$. The monotonicity of each price of anarchy expression shows the logic is reversible, matching the theorem. A similar argument could be followed for other values of the utility rule. $\square$

Appendix B

Appendix Title B

B.1 Omitted proofs of Chapter 5

We prove Lemma 2 using the definition of the Nash flow, and by showing this transformation does not affect user preferences.

Proof of Lemma 2: Let f' be a Nash flow for a game $G \in \mathcal{G}$ under influencing mechanism T. User $x \in N_i$ observes cost

$$J_x(P_x, f') = \sum_{e \in P_x} \ell_e(f'_e) + \tau_e(f'_e), \tag{B.1}$$

and by the definition of Nash flow, will have preferences satisfying

$$J_x(P_x, f') \leq J_x(P', f'), \quad \forall P' \in \mathcal{P}_i. \tag{B.2}$$

In the same flow f', but now under influencing mechanism $\hat{T}$, user x observes cost

$$\hat{J}_x(P_x, f') = \sum_{e \in P_x} \ell_e(f'_e) + \lambda\tau_e(f'_e) + (\lambda - 1)\ell_e(f'_e), \tag{B.3}$$

$$= \sum_{e \in P_x} \lambda(\tau_e(f'_e) + \ell_e(f'_e)) \tag{B.4}$$

$$= \lambda J_x(P_x, f'). \tag{B.5}$$

Observe that through the same process, it can be shown that $\hat{J}_x(P, f') = \lambda J_x(P, f')$ for every $P \in \mathcal{P}_i$. From (B.2),

$$(1/\lambda)\hat{J}_x(P_x, f') \leq (1/\lambda)\hat{J}_x(P_x, f'), \quad \forall P_x \in \mathcal{P}_i \tag{B.6}$$

$$\hat{J}_x(P_x, f') \leq \hat{J}_x(P_x, f'), \quad \forall P_x \in \mathcal{P}_i. \tag{B.7}$$

(B.7) holds for all $x \in N$, satisfying that f' is a Nash equilibrium in G under $\hat{T}$. It is therefore the case that any equilibrium in any game $G \in \mathcal{G}$ under T is also an equilibrium under $\hat{T}$, thus

$$\mathcal{L}^{\text{nf}}(G, T) = \mathcal{L}^{\text{nf}}(G, \hat{T}), \tag{B.8}$$

and, because this holds for every game $G \in \mathcal{G}$, it certainly holds for the supremum over the set which is the same as (5.16) by definition. □

Proof of Proposition 5.2.1: We first look at the optimal bounded toll and its associated price of anarchy bound. Trivially, when $\beta > 1$ the optimal toll is the marginal cost toll that gives price of anarchy of one. For a bounding factor $\beta \in [0, 1)$, a feasible bounded toll must satisfy

$$\tau_e^+(f_e) \in [0, \beta \cdot \ell_e] = [0, \beta a_e f_e + \beta b_e]. \tag{B.9}$$

Because the tolls are network-agnostic, and must satisfy an additivity property discussed in [136] as well as in the proof of Theorem 5.3.2, we can therefore reduce the search for an optimal bounded toll to $\tau_e^+(f_e) = k_1 a_e f_e + k_2 b_2$ where $k_1, k_2 \in [0, \beta]$. We first show that the optimal toll will have $k_2 = 0$.

Let T^+ be a tolling mechanism that assigns bounded tolls with some $k_1, k_2 \in [0, \beta]$. A player $x \in N_i$ utilizing path P_x in a flow f observes cost

$$J_x(P_x, f) = \sum_{e \in P_x} (1 + k_1) a_e f_e + (1 + k_2) b_e. \tag{B.10}$$

Now, consider an incentive mechanism $\hat{T}$ where edges are assigned tolls $\tau_e(f_e) = (\frac{1+k_1}{1+k_2} - 1) a_e f_e$. Under this new incentive, the same player as before now observes cost

$$\hat{J}_x(P_x, f) = \sum_{e \in P_x} \frac{1 + k_1}{1 + k_2} a_e f_e + b_e. \tag{B.11}$$

Because the player's cost in (B.10) and (B.11) are proportional, the players preserve the same preferences and the Nash flows remains unaltered. Because $(\frac{1+k_1}{1+k_2} - 1) \leq k_1 \leq \beta$ the new incentive is bounded by β. Note that any toll that improves the price of anarchy satisfies $k_1 > k_2$; this can be seen by considering the worst-case example depicted in Fig. 5.2 with $p = 1$. Because $0 < k < \beta$, we need only consider tolls of the form $\tau_e(f_e) = k a_e f_e$ when in search of the optimal bounded toll. When $k < 0$ the price of anarchy is at least $4/3$ and is indeed not optimal[1].

For a tolling mechanism $T^+(af + b) = kaf$ with $k \in [0, \beta) \subseteq [0, 1)$, a player's cost takes the form

$$J_x(P_x, f) = \sum_{e \in P_x} (1 + k) a_e + b_e. \tag{B.12}$$

[1] Consider the classic Pigou network, as in Fig. 5.2 with $p = 1$. It is well known this network gives the worst case price of anarchy of $4/3$ with Nash flow of $f_1 = 1$. Consider using a taxation mechanism $T(af + b) = kaf$ for some $k < 0$ and observe that the Nash flow is unchanged, thus not reducing the price of anarchy for the class of affine congestion games.

When player cost functions take this form, the game is similar to that of an altruistic game (introduced in [192]) and has price of anarchy of

$$\mathrm{PoA}(\mathcal{G}^{\mathrm{aff}}, T^{+}) = \frac{4}{3 + 2k - k^2}. \tag{B.13}$$

The price of anarchy is decreasing with $k \in [0, 1)$ and thus the optimal toll occurs when k is maximized at $k = \beta$.

For the optimal subsidy, we now note that incentives must be bounded by $\tau_e(f_e) \in [-\beta\ell_e(f_e), 0]$. From Lemma 2, we can map any such subsidy to an equivalent toll, now constrained to the region $\hat{\tau}_e(f_e) \in [0, \hat{\beta}\ell_e(f_e)]$ where $\hat{\beta} = (\frac{1}{1-\beta} - 1)$. It was shown prior that the optimal tolling mechanism in this region is $\hat{T}(af + b) = \hat{\beta}af$. Finally, we can again use Lemma 2 to map back to the optimal bounded subsidy,

$$T^{\mathrm{opt}-}(af + b) = (\lambda - 1)(af + b) + \lambda\hat{T}(af + b), \tag{B.14}$$

with $\lambda = 1 - \beta$. The result is an optimal subsidy of the form $T^{\mathrm{opt}-}(af + b) = -\beta b$ for $\beta \in [0, 1/2)$. The price of anarchy bound comes from considering the equivalent toll. □

Proof of Lemma 3: First, we assume without loss of generality, that $S_{\mathrm{L}} = 1$. To see this, we make an equivalent problem where this is true and show the same price of anarchy bound holds. Let T be any incentive mechanism and $\mathcal{S}$ be a family of sensitivity distributions with lower bound S_{L} and upper bound S_{U}. In any game $G \in \mathcal{G}$, a player $x \in N_i$ observes costs as expressed in (5.3). Observe that if we normalize every sensitivity distribution $s \in \mathcal{S}$ by multiplying by $1/S_{\mathrm{L}}$ and correspondingly scale the incentive by S_{L} the player cost remains unchanged. It is therefore the case that any equilibrium is preserved and unchanged, enforcing that

$$\text{PoA}(\mathcal{G}, \mathcal{S}, T) = \text{PoA}\left(\mathcal{G}, \mathcal{S}/S_{\text{L}}, S_{\text{L}} \cdot T\right). \tag{B.15}$$

Accordingly, we will consider that $S_{\text{L}} = 1$ throughout.

Let f be a flow in $G \in \mathcal{G}$ induced by sensitivity distribution $s \in \mathcal{S}$, and let T be an incentive mechanism that assigns tolls τ_e^+. From Lemma 2 a nominally equivalent incentive mechanism can be found by using the transformation $\hat{T}(\ell_e; \lambda) = (\lambda - 1)\ell_e + \lambda T(\ell_e)$, where choosing λ sufficiently close to zero causes $\hat{T}$ to be a subsidy mechanism. We will show that for any $\lambda \in (0, 1)$, the incentive mechanism $\hat{T}$ performs worse than T at the introduction of player heterogeneity.

Let $\hat{s}$ be a new sensitivity distribution such that

$$\hat{s}_x = g(s_x, \lambda) = \frac{s_x}{\lambda + s_x - s_x\lambda}, \tag{B.16}$$

for all $x \in N$. Now, consider an agent's cost in flow f with sensitivity $\hat{s}$ under incentive mechanism $\hat{T}$. An agent $x \in N_i$ utilizing path P_x in f experiences cost,

$$\begin{aligned}
\hat{J}_x(P_x, f) &= \sum_{e \in P_x} \ell_e(f_e) + \hat{s}_x \hat{T}(\ell_e(f_e); \lambda) \\
&= \sum_{e \in P_x} \ell_e(f_e) + \hat{s}_x[(\lambda - 1)\ell_e + \lambda\tau_e^+(f_e)] \\
&= \frac{\lambda}{\lambda + s_x - s_x\lambda} \sum_{e \in P_x} (\ell_e(f_e) + s_x\tau_e(f_e)),
\end{aligned}$$

which is proportional to $J_x(P_x, f)$. By observing proportional costs, players preserve the same preferences over paths, preserving the same Nash flows.

Finally, we show that $\hat{s}$ is a feasible sensitivity distribution in $\mathcal{S}$. From the original bounds S_{L} and S_{U}, any generated distribution $\hat{s}$ exists between $g(S_{\mathrm{L}}, \lambda)$ and $g(S_{\mathrm{U}}, \lambda)$. From before, $S_{\mathrm{L}} = 1$, thus from (B.16), $g(S_{\mathrm{L}} = 1, \lambda) = 1 = S_{\mathrm{L}}$, for any $\lambda \in (0, 1)$. Now, observe that any generated distribution satisfies

$$g(S_{\mathrm{U}}, \lambda) = \frac{S_{\mathrm{U}}}{\lambda + S_{\mathrm{U}} - S_{\mathrm{U}}\lambda} \leq S_{\mathrm{U}}, \tag{B.17}$$

for any $\lambda \in (0, 1)$. Thus any generated distribution $\hat{s}$ is sufficiently bounded by S_{L} and S_{U} and is a feasible distribution in $\mathcal{S}$. By choosing f to be a Nash flow, we can see that any Nash flow that can be induced by some $s \in \mathcal{S}$ while using T can similarly be induced by $\hat{s} \in \mathcal{S}$ while using $\hat{T}$. It is therefore the case that the price of anarchy with user heterogeneity is non-decreasing as λ decreases, showing the monotonicity. Further, if $S_{\mathrm{L}} < S_{\mathrm{U}}$, then $S_{\mathrm{L}} \leq g(S_{\mathrm{L}}, \lambda) \leq g(S_{\mathrm{U}}, \lambda) < S_{\mathrm{U}}$, and if $\mathcal{G}$ is responsive to user heterogeneity, the price of anarchy is strictly increasing with λ. □

Proof of Theorem 5.3.1: Lemma 3 states that though two incentive mechanisms have the same price of anarchy when users are homogeneous (from Lemma 2), they need not perform the same when users are heterogeneous. Further, by increasing λ, one can lower the heterogeneous price of anarchy without altering the performance in the homogeneous setting. The proof of Theorem 5.3.1 is a simple extension of Lemma 3. Increasing λ reduces the effect of player heterogeneity on the price of anarchy, and by letting $\lambda \to \infty$ we can construct an incentive that recovers (5.26).

In Theorem 1 of [136], the authors propose a realization of this result when using marginal cost taxes. In the class of congestion games where marginal cost taxes are optimal in the homogeneous setting, they show that the taxation mechanism

$$T^u(\ell_e; k)[f_e] = k \left(\ell_e(f_e) + f_e \cdot \frac{d}{df_e} \ell_e(f_e) \right)$$

has a price of anarchy of 1 as k approaches infinity, i.e.,

$$\lim_{k\to\infty} \mathrm{PoA}(\mathcal{G}, \mathcal{S}, T^u(k)) = 1. \tag{B.18}$$

This same result can be recovered using Theorem 5.3.1. The marginal cost taxation mechanism defined in (5.5) has the same performance as

$$T_\lambda(\ell_e)[f_e] = (\lambda - 1)\ell_e(f_e) + \lambda T^{\mathrm{mc}}(\ell_e)[f_e].$$

By taking the limit as λ approaches infinity, this incentive becomes

$$\begin{aligned} T_\lambda(\ell_e)[f_e] &= \lambda\left(\ell_e(f_e) + f_e \cdot \frac{d}{df_e}\ell_e(f_e)\right) \\ &= T^u(\ell_e; \lambda)[f_e]. \end{aligned}$$

Not only does this give us the same toll, but by Lemma 3, we know that

$$\lim_{\lambda\to\infty} \mathrm{PoA}(\mathcal{G}, \mathcal{S}, T_\lambda) = \mathrm{PoA}(\mathcal{G}, T^{\mathrm{mc}}) = 1, \tag{B.19}$$

giving the final statement in Theorem 5.3.1. □

Proof of Theorem 5.3.2: First, consider a game $G \in \overline{\mathcal{G}}$ that has a unique equilibrium and optimal flow respectively, to obtain a heterogeneous price of anarchy of one, the equilibrium must be the same for any sensitivity distribution $s \in \mathcal{S}$. If a taxation mechanism is agnostic of the users' sensitivities, the only way this can be accomplished is by letting the magnitude of the subsidies become large compared to the latency function; for a player $x \in N$ this causes $J_x(P_x, f) \approx \sum_{e\in P_x} s_x T(\ell_e)[f_e]$. With this subsidy, the users price sensitivity does not affect their preference over paths.

Each of the following three conditions is necessary for a sufficiently large subsidy to incentivize optimal routing (we justify each but note the proof that any one is necessary is trivial).

1. *Additivity.* A network-agnostic incentive mechanism must satisfy $T(\alpha\ell_1 + \beta\ell_2) = \alpha T(\ell_1) + \beta T(\ell_2)$. A proof of this appears in [7]; intuitively, a single latency function can be represented as multiple in series and the total incentive must be the same in both cases to guarantee the same total cost.

2. *Incentives are Unbounded.* $|T(\ell)[f]| > M \ \forall\ M \in [0, \infty)\ \forall \ell \in L(\mathcal{G}),\ f > 0$. Any bounded incentive may allow different sensitivity distributions to induce different equilibrium flows. When the optimal flow is unique, a bounded incentive is incapable of enforcing each equilibrium flow be optimal.

3. *Related by Marginal Cost.* For any two edges ℓ_i, ℓ_j with respective flow f_i, f_j, if $\ell_i^{\mathrm{mc}}(f_i) \leq \ell_j^{\mathrm{mc}}(f_j)$ then $T(\ell_i)[f_i] \leq T(\ell_j)[f_j]$, where $\ell_i^{\mathrm{mc}}(f_i) = \ell_i(f_i) + f_i \frac{d}{df_i}\ell_i(f_i)$ is the marginal cost on edge i. Recall that T is defined as a cost and therefore negative for subsidies, thus this condition states that users must receive less subsidy on edges with higher marginal cost. It is shown in [132] that when users observe the marginal cost, the equilibrium flow is optimal.

Note that condition 2 implies player costs are negative: $\ell(f) + T(\ell)[f] < 0\ \forall \ell \in L(\mathcal{G}),\ f > 0$. Similarly, condition 3 implies that incentives are non-decreasing: if $f_1 > f_2$ then $T(\ell)[f_1] \geq T(\ell)[f_2]\ \forall \ell \in L(\mathcal{G})$.

We now show that no network-agnostic subsidy mechanism can satisfy each of these three conditions. Assume T is an optimal subsidy mechanism. By the symmetry of condition 3, we see that if $\ell_i^{\mathrm{mc}}(f_i) = \ell_j^{\mathrm{mc}}(f_j)$, then $T(\ell_i)[f_i] = T(\ell_j)[f_j]$. Consider a unit mass of traffic traversing a two link parallel network with edges possessing latency functions ℓ_1 and ℓ_2 that are strictly increasing. Let f_1 be the solution to $\ell_1^{\mathrm{mc}}(f_1) =$

$\ell_2^{\mathrm{mc}}(1-f_1)$, and by condition 3, $T(\ell_1)[f_1] = T(\ell_2)[1-f_1]$. Now, consider a similar network, but ℓ_2 is replaced by a scaled latency function $\frac{1}{2}\ell_2$. Now, define f_1' as the solution to $\ell_1^{\mathrm{mc}}(f_1') = \frac{1}{2}\ell_2^{\mathrm{mc}}(1-f_1')$; from ℓ_1, ℓ_2 strictly increasing, $f_1' < f_1$. Implied by condition 3, $T(\ell_1)[f_1'] < T(\ell_1)[f_1]$ and $T(\ell_2)[1-f_1] < T(\ell_2)[1-f_1']$. From conditions 1 and 2, $T(\ell_2)[1-f_1'] < \frac{1}{2}T(\ell_2)[1-f_1'] = T(\frac{1}{2}\ell_2)[1-f_1']$. Put together this gives,

$$\begin{aligned} T(\ell_1)[f_1'] < T(\ell_1)[f_1] &= T(\ell_2)[1-f_1] \\ &< T(\ell_2)[1-f_1'] \\ &< T(\frac{1}{2}\ell_2)[1-f_1'], \end{aligned}$$

implying $T(\ell_1)[f_1'] \neq T(\frac{1}{2}\ell_2)[1-f_1']$, contradicting condition 3. □

Proof of Proposition 5.3.1: The first part of the proposition comes from [142]. We thus find the nominally equivalent subsidy mechanism and find the associated price of anarchy bound.

For notational convenience, let $k = 1/\sqrt{S_{\mathrm{L}}S_{\mathrm{U}}}$; the robust marginal cost toll is thus $T^{\mathrm{smc}}(af+b) = kaf$. From Lemma 2, we can derive a nominally equivalent subsidy by $T^{\mathrm{nes}}(af+b) = (\lambda-1)(af+b) + \lambda(kaf)$, for any $\lambda > 0$. By letting $\lambda = 1/(1+k)$, we get the nominally equivalent subsidy to be $T^{\mathrm{nes}}(af+b) = -kb/(1+k) = -\frac{1}{1+\sqrt{S_{\mathrm{L}}S_{\mathrm{U}}}}b$.

To determine the price of anarchy of T^{nes} with player heterogeneity, we use the result of Theorem 5.3.3 to determine the equivalent level of heterogeneity on the nominally equivalent toll, T^{smc}. Let $s \in \mathcal{S}$ be a feasible sensitivity distribution, bounded by S_{L} and S_{U}. As it is defined above, we seek to find the preimage of $[S_{\mathrm{L}}, S_{\mathrm{U}}]$ under the function $g(S, 1/(1+k))$. Without loss of generality, we normalize $[S_{\mathrm{L}}, S_{\mathrm{U}}]$, to $[q, 1]$ and look for its preimage. Because g is continuous on $S \in [0,1]$, we look at the endpoints of the region. We first note that $g(1,\lambda) = 1$ for any $\lambda > 0$. Next, we determine $\hat{q}$ such that $g(\hat{q},\lambda) = q$

as

$$\hat{q} = \frac{\lambda q}{1 - q + \lambda q},$$

and by setting $\lambda = 1/(1+k) = \sqrt{S_\mathrm{L} S_\mathrm{U}}/(1 + \sqrt{S_\mathrm{L} S_\mathrm{U}})$ recover the equivalent amount of heterogeneity, $\hat{q}$, on T^{smc} as the original subsidy T^{nes} with heterogeneity q. By replacing q with $\hat{q}$ in (5.35) we obtain the price of anarchy for T^{nes} with heterogeneity. □

Proof of Proposition 5.4.2: The proof follows similar steps to that of Proposition 5.4.2, which appears in [136]. Let $G \in \mathcal{G}^{\mathrm{pa}}$ be a game instance and user be distributed with sensitivity $s \in \mathcal{S}$. Because each network in $\mathcal{G}$ is parallel, each path constitutes a single edge. Under an affine subsidy mechanism $T^{-}(k_1, k_2)$, player $x \in N$ utilizing edge e observes cost

$$J_x(e, f) = (1 - k_1 s_x) a_e f_e + (1 - k_2 s_x) b_e,$$

where $k_1, k_2 > 0$. Note that scaling users cost functions does not alter their preference over their paths, thus without loss of generality we can write player costs as

$$J_x(e, f) = \frac{(1 - k_1 s_x)}{(1 - k_2 s_x)} a_e f_e + b_e. \tag{B.20}$$

We define a new incentive mechanism $T'(af + b) = k' a f$. Now, let s' be a new sensitivity distribution such that players observe the same cost under T' as they did in (B.20) with sensitivity distribution s, i.e.,

$$\frac{(1 - k_1 s_x)}{(1 - k_2 s_x)} = (1 + k' s'_x). \tag{B.21}$$

The new distribution can be realized by the transformation

$$s'_x = \frac{s_x (k_2 - k_1)}{k'(1 - k_2 s_x)}. \tag{B.22}$$

The taxation mechanism T' constitutes a scaled marginal cost toll, for which, the following result exists:

Theorem B.1.1. Brown & Marden [142]: *For any network $G \in \mathcal{G}$ with flow on all edges in an un-tolled Nash flow, and any $s \in \mathcal{S}$, any scaled marginal cost taxation mechanism reduces the total latency of any Nash flow when compared to the total latency of any Nash flow associated with the un-tolled case, i.e., for any $k > 0$*

$$\mathcal{L}^{\mathrm{nf}}(G, s, T^A(k, 0)) < \mathcal{L}^{\mathrm{nf}}(G, s, \emptyset). \tag{B.23}$$

Furthermore, the unique optimal scaled marginal-cost tolling mechanism uses the scale factor

$$k^* = \frac{1}{\sqrt{S_{\mathrm{L}} S_{\mathrm{U}}}} = \arg\min_{k \geq 0} \{\mathrm{PoA}(\mathcal{G}, \mathcal{S}, T^A(k, 0))\}. \tag{B.24}$$

Finally, the price of anarchy resulting from the optimal scaled marginal-cost taxation mechanism is

$$\mathrm{PoA}(\mathcal{G}, \mathcal{S}, T^A(k^*, 0)) = \frac{4}{3}\left(1 - \frac{\sqrt{S_{\mathrm{L}}/S_{\mathrm{U}}}}{\left(1 + \sqrt{S_{\mathrm{L}}/S_{\mathrm{U}}}\right)^2}\right) \leq \frac{4}{3}. \tag{B.25}$$

Because of this, we set $k' = \frac{1}{\sqrt{S'_{\mathrm{L}} S'_{\mathrm{U}}}}$ to be the optimal scaled marginal cost taxation mechanism over the new family of sensitivity distributions $\mathcal{S}'$, generated by transforming each sensitivity distribution in $\mathcal{S}$ as in (B.22).

Now, in the original subsidy mechanism $T(k_1, k_2)$, let

$$k_1 = k' = \frac{1}{\sqrt{S'_{\mathrm{L}} S'_{\mathrm{U}}}}. \tag{B.26}$$

Combining (B.22) and (B.26) gives an expression for the accompanying choice of k_2 to satisfy (B.21),

$$k_2 = \frac{k_1^2 S_L S_U - 1}{2k_1 S_L S_U - S_L - S_U}. \tag{B.27}$$

Observe that (B.25) is decreasing with $S_L/S_U < 1$. For the similar taxation mechanism T',

$$\frac{S'_L}{S'_U} = \frac{S_L(1 - k_2 S_U)}{S_U(1 - k_2 S_L)}, \tag{B.28}$$

for $\mathcal{S}'$ found by (B.22). Notice (B.28) is decreasing with $k_2 < S_U$[2]. Therefore, from (B.27), decreasing k_1 decreases the price of anarchy; thus picking $k_1^* = 0$ is optimal. Substituting into (B.27) gives $k_2^* = 1/(S_L + S_U)$. For $\beta < 1/(S_L + S_U)$ it is easy to show by a similar transformation that $k_1^* = 0$ and $k_2^* = \beta$.

Finally, for subsidies of the form $T^A(0, -k)$ with $k < 1$, we show the price of anarchy bound. Let $s \in \mathcal{S}$ be the users' sensitivity distribution. Let $T^+ = (\lambda - 1)\ell(x) + \lambda T^A(0, -k)$ and let s' be a new sensitivity distribution. Letting $\lambda = 1/(1 - k)$ and $s'_x = \frac{s_x(1-k)}{1-ks_x}$,

$$\begin{aligned}
&\ell_e(f_e) + s'_x T^+(\ell_e)[f_e] = \\
&a_e f_e + b_e + \frac{s_x(1-k)}{1-ks_x}\left[\Big(\frac{1}{1-k} - 1\Big)(a_e f_e + b_e) - \frac{k}{1-k} b_e\right] \\
&= \frac{1}{1-ks_x}(a_e f_e + b_e - ks_x b_e) \\
&\propto \ell_e(f_e) + s_x T^A(0, -k)[f_e].
\end{aligned}$$

From player costs being proportional, we can analyze T^+ to get the price of anarchy bound. The new incentive manifests as $T^+(af + b) = \frac{k}{1-k} af$. Because $k \le \frac{1}{S_L + S_U}$ for an optimal bounded subsidy (and by our assumption that $S_L S_U = 1$ for ease of notation),

$$k \le \frac{1}{S_L + S_U - 1} = \frac{1}{1/S_U + S_U - 1} \le 1 = \frac{1}{\sqrt{S_L S_U}}.$$

[2] k_2 need be less than S_U for a bounded price of anarchy. A simple construction to show this is Pigou's network in Fig. 5.2 with $s = S_U$.

Thus the price of anarchy for T^+ is dictated by (5.41). Substituting $S'_{\mathrm{L}} = \frac{S_{\mathrm{L}}(1-k)}{1-kS_{\mathrm{L}}}$ and $\beta = k/(1-k)$ for $k \in [0, \frac{1}{S_{\mathrm{L}}+S_{\mathrm{U}}}]$ gives the price of anarchy in Proposition 5.4.2. □

B.2 Omitted proofs of Chapter 6

B.2.1 Proofs for general congestion networks

Proof of Lemma 4: To prove the first claim, consider the prior μ_0 on α and the signalling policy $\pi : A \to \Delta(S)$. If the signal $s \in S$ is sent to users, they update their belief via Bayesian inference to $\mu_s(\alpha) = \frac{\pi(s|\alpha)\cdot\mu_0(\alpha)}{\psi(s)}$. In a flow f, user $x \in [0, r_i]$ taking path $P_x \in \mathcal{P}_i$ experiences an expected cost of

$$\begin{aligned} J_x(P_x; f, \mu_s) &= \mathop{\mathbb{E}}_{\alpha\sim\mu_s}\left[\sum_{e\in P_x}\sum_{k=1}^{|\mathcal{K}|}\alpha_{e,k}\cdot\ell_k(f_e)\right] \\ &= \sum_{e\in P_x}\sum_{k=1}^{|\mathcal{K}|}\mathbb{E}[\alpha_{e,k}|s]\ell_k(f_e). \end{aligned}$$

Note that if f were not a Nash flow in the congestion game with coefficients $\mathbb{E}[\alpha|s]$, then by (6.3) there exists a user x who would be able to deviate their strategy $\sigma_x(s)$ and experience lower cost. Therefore, the only Bayes-Nash flows occur when $f(s)$ is a Nash flow with respect to $\mathbb{E}[\alpha|s]$ for all $s \in S$. Further, because the total latency in a Nash flow is unique, so too is the expected total latency in a Bayes-Nash flow.

To prove the second claim, consider the distribution μ_0 on α, and let f be some flow.

The expected total latency

$$\mathcal{L}(f;\mu) = \mathop{\mathbb{E}}_{\alpha\sim\mu}\left[\sum_{e\in E} f_e \cdot \sum_{k=1}^{|\mathcal{K}|} \alpha_{e,k}\ell_k(f_e)\right]$$
$$= \sum_{e\in E} f_e \cdot \sum_{k=1}^{|\mathcal{K}|} \mathop{\mathbb{E}}_{\alpha\sim\mu}\left[\alpha_{e,k}\right]\ell_k(f_e) = \mathcal{L}(f; \mathop{\mathbb{E}}_{\alpha\sim\mu}[\alpha]),$$

which follows from the linearity of expected value. □

Proof of Proposition 6.3.1: Consider that the users receive signal s from signalling policy π and prior μ_0 (forming posterior μ_s), and reach a flow of f. From Lemma 4 the expected total latency in this flow is equal to the total latency of this flow in the expected network, i.e., $\mathcal{L}(f;\mu_s) = \mathcal{L}(f;\overline{\alpha}_s)$, where $\overline{\alpha}_s = \mathbb{E}_{\alpha\sim\mu_s}[\alpha]$. Thus, an optimal flow at the reception of signal s is one that satisfies $f^{\text{opt}}(s) \in \arg\min_f \mathcal{L}(f;\overline{\alpha}_s)$.

We now look for an incentive that will influence users such that $f^{\text{opt}}(s)$ becomes a Nash flow in a congestion game with latency coefficients $\overline{\alpha}_s$. To do so, we note that G with flow-varying incentive functions $\tau_e(f_e)$ is a potential game [152] with potential function $\Phi(f;\alpha) = \sum_{e\in E}\int_0^{f_e} \ell_e(z) + \tau_e(z)dz$. As such, the flow in $\arg\min_f \Phi(f;\alpha)$ is a Nash equilibrium. For the polynomial latency functions considered in this work, let $\tau_e(z) = \sum_{k=1}^{|\mathcal{K}|} z\alpha_{e,k}\ell_k'(z)$. Now, the potential function becomes

$$\Phi(f;\alpha) = \sum_{e\in E}\int_0^{f_e} \sum_{k=1}^{|\mathcal{K}|} \alpha_{e,k}\ell_k(z) + z\alpha_{e,k}\ell_k(z)dz$$
$$= \sum_{e\in E}\sum_{k=1}^{|\mathcal{K}|} f_e\alpha_{e,k}\ell_k(f_e) = \mathcal{L}(f;\alpha),$$

and the Nash flow that minimizes Φ also minimizes $\mathcal{L}$; as such, $f^{\text{opt}}(s)$ becomes a Nash equilibrium in the game with the coefficients $\overline{\alpha}_s$.

Finally, notice that by selecting the fixed incentive $\tau_e^\star(s) = \tau_e(f_e^{\text{opt}}(s))$, the equilibrium conditions do not change and $f^{\text{opt}}(s)$ remains a Nash flow. Nash flows retain the same uniqueness properties under fixed incentives, and thus assigning $\tau^\star(s)$ minimizes the expected total latency when s is sent. If this is done for each signal, the total latency with each signal will be minimal, and so too will the overall expected total latency, making $T^\star$ an optimal incentive mechanism. □

B.2.2 Proofs for parallel network polynomial latency

Remark 2. *In parallel, polynomial Bayesian congestion games, without loss of generality, we can assume a unit traffic rate, $r = 1$, even when $r \sim \nu$ is a random variable.*

Proof: Consider a congestion game G with demand r and latency functions from the basis set of polynomials $\mathcal{D}$. Define a mapping $Q(G,\gamma)$ that outputs a new congestion game $\hat{G}$ with latency functions $\hat{\ell}_e(f_e) = \sum_{d\in\mathcal{D}} \frac{\alpha_{e,d}}{\gamma^{d+1}}(f_e)^d$. Let f be a flow in G with total traffic r. Now, consider the flow $\gamma f = \{\gamma \cdot f_e\}_{e\in E}$ in $\hat{G}$. Each edge $e \in E$ will have latency $\hat{\ell}_e(\gamma f_e) = \sum_{d\in\mathcal{D}} \frac{\alpha_{e,d}}{\gamma^{d+1}}(\gamma f_e)^d = \frac{1}{\gamma}\ell_e(f_e)$. Notice that latency on each edge is scaled by $1/\gamma$, and the preference structure is preserved; therefore, if f is a Nash flow in G, then γf is a Nash flow in $\hat{G}$. Further, $\mathcal{L}(\gamma f; \hat{G}) = \sum_{e\in E} \gamma f_e \hat{\ell}_e(\gamma f_e) = \sum_{e\in E} f_e \ell_e(f_e) = \mathcal{L}(f; G)$, and the two networks will have the same total latency.

If $(\alpha, r) \sim \mu_0$, e.g., $\mu_0(x,y) = \mathbb{P}[\alpha = x, r = y]$, then consider that $\hat{\alpha} \sim \hat{\mu}_0$, where $\hat{\mu}_0(z) = \sum_{x,y|Q(x,1/y)=z} \mu_0(x,y)$. Now $\hat{\alpha}$ has the same distribution over total latency. □

Proof of Lemma 5 We assume $r = 1$, which is without loss of generality from Remark 2. We note that $\mathcal{L}^{\text{Ne}}(\alpha)$ is continuous but need not be continuously differentiable; as such, we look for the largest gradient in the differentiable regions of the support. Let $f^{\text{Ne}}(\alpha)$ be the Nash flow in the parallel congestion game with polynomial coefficients α, i.e., $\mathcal{L}^{\text{Ne}}(\alpha) = \mathcal{L}(\alpha, f^{\text{Ne}}(\alpha)) = \sum_{e\in E}\sum_{d\in\mathcal{D}} \alpha_{e,d}(f_e^{\text{Ne}})^{d+1}$. First, we seek to bound the partial derivative of $\mathcal{L}^{\text{Ne}}(\alpha)$ with respect to some parameter $\alpha_{e,d}$. Clearly, $\frac{\partial}{\partial \alpha_{e,d}}\mathcal{L}^{\text{Ne}}(\alpha) \geq 0$ in parallel networks as no Braess's paradox type example can exist [195]. To upper-bound this partial derivative, we will consider a case where by increasing $\alpha_{e,d}$ any mass of traffic that chooses to leave edge e will all choose the edge e'; in general, this may not occur with every change in $\alpha_{e,d}$, as users may disperse over multiple edges, however, if we consider that users do all move to the same edge, and we pick edge e' as the one that increases the total latency most rapidly, then the following upper-bound will hold. With this, note $\frac{\partial}{\partial \alpha_{e,d}} f_e^{\text{Ne}} = -\frac{\partial}{\partial \alpha_{e,d}} f_{e'}^{\text{Ne}}$, and the partial derivative is

$$\frac{\partial}{\partial \alpha_{e,d}} \mathcal{L}^{\text{Ne}}(\alpha) = (f_e^{\text{Ne}})^{d+1} + \left(\sum_{d'' \in \mathcal{D}} \alpha_{e',d''} (d''+1) (f_{e'}^{\text{Ne}})^{d''+1} - \sum_{d' \in \mathcal{D}} \alpha_{e,d'} (d'+1) (f_e^{\text{Ne}})^{d'} \right) \frac{\partial}{\partial \alpha_{e,d}} f_{e'}^{\text{Ne}} \quad \text{(B.29)}$$

Now, we note that latency on edges e and e' must be the same in a Nash flow, thus $\ell_e(f_e^{\text{Ne}}) = \ell_{e'}(f_{e'}^{\text{Ne}})$ and $\frac{\partial}{\partial \alpha_{e,d}} \ell_e(f_e^{\text{Ne}}) = \frac{\partial}{\partial \alpha_{e,d}} \ell_{e'}(f_{e'}^{\text{Ne}})$. Using this equality, and the fact that $\frac{\partial}{\partial \alpha_{e,d}} f_e^{\text{Ne}} = -\frac{\partial}{\partial \alpha_{e,d}} f_{e'}^{\text{Ne}}$, we can evaluate the derivative and rearrange to get

$$\frac{\partial}{\partial \alpha_{e,d}} f_{e'}^{\text{Ne}} = \frac{(f_e^{\text{Ne}})^d}{\ell'_e(f_e^{\text{Ne}}) + \ell'_{e'}(f_{e'}^{\text{Ne}})} \leq \frac{(f_e^{\text{Ne}})^d}{2\rho_1^-}, \quad \text{(B.30)}$$

where $\rho_1^- = \min_{e \in E} \check{\alpha}_{e,1}$. Substituting (B.30) into (B.29) gives us

$$\frac{\partial}{\partial \alpha_{e,d}} \mathcal{L}^{\text{Ne}}(\alpha) \leq (f_e^{\text{Ne}})^{d+1} + \frac{\rho^+ - \rho_0^-}{2\rho_1^-} (f_e^{\text{Ne}})^d, \quad \text{(B.31)}$$

where $\rho_0^- = \min_{e \in E} \check{\alpha}_{e,0}$ and $\rho^+ = \max_{e \in E} \sum_{d \in \mathcal{D}} (d+1) \hat{\alpha}_{e,d}$. Now, the gradient of $\mathcal{L}^{\text{Ne}}(\alpha)$

must satisfy

$$
\begin{aligned}
\|\nabla \mathcal{L}^{\text{Ne}}(\alpha)\|_2 &\leq \sqrt{\sum_{e\in E}\sum_{d\in\mathcal{D}}\left((f_e^{\text{Ne}})^{d+1} + \frac{\rho^+ - \rho_0^-}{2\rho_1^-}(f_e^{\text{Ne}})^d\right)^2} \\
&\leq \sqrt{\left(\sum_{e\in E}\sum_{d\in\mathcal{D}}(f_e^{\text{Ne}})^{d+1} + \frac{\rho^+ - \rho_0^-}{2\rho_1^-}(f_e^{\text{Ne}})^d\right)^2} \\
&\leq \sum_{d\in\mathcal{D}}\left(\sum_{e\in E} f_e^{\text{Ne}}\right)^{d+1} + \frac{\rho^+ - \rho_0^-}{2\rho_1^-}\cdot \\
&\qquad \left(\sum_{e\in E}(f_e^{\text{Ne}})^0 \sum_{d\in\mathcal{D}\setminus\{0\}}\left(\sum_{e\in E} f_e^{\text{Ne}}\right)^d\right) \\
&= |\mathcal{D}| + \frac{\rho^+ - \rho_0^-}{2\rho_1^-}\left(|E| + |\mathcal{D}| - 1\right),
\end{aligned}
$$

where the first inequality holds from (B.31), the second and third hold from the super-additivity of convex monomials of positive terms, and the final equality holds from the assumption that $r = 1$. Finally, consider two sets of coefficients $a, b \in A$.

$$
\begin{aligned}
&\frac{\mathcal{L}^{\text{Ne}}(a) - \mathcal{L}^{\text{Ne}}(b)}{\|a-b\|_2} \\
&= \frac{1}{\|a-b\|_2}\int_{\lambda=0}^{1}(a-b)^T\nabla\mathcal{L}^{\text{Ne}}(\lambda a + (1-\lambda)b)d\lambda \\
&\leq \frac{1}{\|a-b\|_2}\int_{\lambda=0}^{1}\|a-b\|_2\cdot\|\nabla\mathcal{L}^{\text{Ne}}(\lambda a + (1-\lambda)b)\|_2 d\lambda \\
&\leq \int_{\lambda=0}^{1}|\mathcal{D}| + \frac{\rho^+ - \rho_0^-}{2\rho_1^-}\left(|E| + |\mathcal{D}| - 1\right)d\lambda \\
&= |\mathcal{D}| + \frac{\rho^+ - \rho_0^-}{2\rho_1^-}\left(|E| + |\mathcal{D}| - 1\right).
\end{aligned}
$$

Where the first inequality holds from Cauchy-Schwarz, and the second holds from our observation above on the norm of the gradient of $\mathcal{L}^{\text{Ne}}(\alpha)$. □

Proof of Lemma 6: This proof follows very similarly to the proof of Lemma 5, but now, in an optimal flow $f^\star(\alpha)$, the latency on each edge is not equal, however, the marginal-cost on each edge is [33]. Let $\nu_e(f_e) = \sum_{d\in\mathcal{D}}(d+1)\alpha_{e,d}(f_e)^d$ be the marginal cost on edge e with flow f_e. Now, in the optimal flow $f^\star$ (which emerges from using the tolls $T^\star$), $\nu_e(f_e^\star) = \nu_{e'}(f_{e'}^\star)$ and $\frac{\partial}{\partial\alpha_{e,d}}\nu_e(f_e^\star)4 = \frac{\partial}{\partial\alpha_{e,d}}\nu_{e'}(f_{e'}^\star)$. Evaluating and rearranging these derivatives gives

$$\frac{\partial}{\partial\alpha_{e,d}}f_{e'}^\star = \frac{(d+1)(f_e^\star)^d}{\nu_e'(f_e^\star)+\nu_{e'}'(f_{e'}^\star)} \leq \frac{(d+1)(f_e^\star)^d}{4\rho_1^-}. \tag{B.32}$$

With this, we can upper bound the partial derivative of $\mathcal{L}^\star$ as

$$\frac{\partial}{\partial\alpha_{e,d}}\mathcal{L}^\star(\alpha) \leq (f_e^\star)^{d+1} + \frac{\rho^+ - \rho_0^-}{4\rho_1^-}(d+1)(f_e^\star)^d. \tag{B.33}$$

Following the same steps as in the proof of Lemma 5, the gradient of $\mathcal{L}^\star$ must satisfy

$$\|\nabla\mathcal{L}^\star(\alpha)\|_2 \leq |\mathcal{D}| + \frac{\rho^+ - \rho_0^-}{4\rho_1^-}\left(|E| + \sum_{d\in\mathcal{D}\setminus\{0\}}(d+1)^d\right).$$

Finally, we can use this bound as in the proof of Lemma 5 to complete the proof. □

B.3 Omitted proofs of Chapter 7

Proof of Lemma 9: Consider a network $G \in \mathcal{G}$ with affine latency functions on each link $\ell_i(f) = a_i f + b_i$. Let $\hat{G}$ have cost functions $\hat{\ell}_i(f) = \hat{a}_i f + \hat{b}_i$ with $\hat{a}_i \geq 0$ and $\hat{b}_i \geq 0$. We first show that simply removing the constant latency term on the first edge b_1 strictly increases the price of anarchy under any scaled marginal-cost toll. Using the optimal and Nash flow in (B.34), if $\hat{b}_2 = b_2 - b_1$ and $\hat{b}_1 = 0$ then G and $\hat{G}$ will have the same optimal flow and Nash flow for a distribution s. From (7.2), we observe that

$\mathcal{L}^{\text{opt}}(G) = \mathcal{L}^{\text{opt}}(\hat{G}) + b_1$ as well as $\mathcal{L}^{\text{nf}}(G, s, k) = \mathcal{L}^{\text{nf}}(\hat{G}, s, k) + b_1$; therefore,

$$\text{PoA}\left(G, \mathcal{S}(\overline{s}), T(k)\right) = \frac{\mathcal{L}^{\text{nf}}\left(\hat{G}, s, T(k)\right) + b_1}{\mathcal{L}^{\text{opt}}(\hat{G}) + b_1}$$

$$\leq \frac{\mathcal{L}^{\text{nf}}\left(\hat{G}, s, T(k)\right)}{\mathcal{L}^{\text{opt}}(\hat{G})} = \text{PoA}(\hat{G}, \mathcal{S}(\overline{s}), T(k)).$$

Thus, for any network $G \in \mathcal{G}$, there exists a network $\hat{G}$ with a linear latency function on an edge with higher price of anarchy.

Next, we show a network $G \in \mathcal{G}$ will have the same price of anarchy as a network $\hat{G} \in \mathcal{G}$ under the same linear toll if the latency functions of $\hat{G}$ equal the latency functions of G times a scaling factor c. Under a distribution $s \in \mathcal{S}$, G and $\hat{G}$ will have the same Nash flow. Using the indifferent sensitivity S_{ind} that is the solution to (7.14), the Nash flow and optimal flow on the first edge are

$$f_1^{\text{opt}} = \frac{2a_2 + b_2 - b_1}{2(a_1 + a_2)}, \quad f_1^{\text{Ne}} = \frac{(1 + S_{\text{ind}}k)a_2 + b_2 - b_1}{(1 + S_{\text{ind}}k)(a_1 + a_2)}. \tag{B.34}$$

Under the same distribution s, S_{ind} will satisfy

$$(1 + S_{\text{ind}}k)ca_1 f_1^{\text{Ne}} + cb_1 = (1 + S_{\text{ind}}k)ca_2 f_2^{\text{Ne}} + cb_2, \tag{B.35}$$

which are the latency functions for the network $\hat{G}$. It is now clear that G and $\hat{G}$ will have the same Nash and optimal flows. From the definition of total latency in (7.2), the latency in $\hat{G}$ will be c times the latency in G under the same flow. The price of anarchy, which is the ratio of two total latencies, will be identical in G and G'.

Lastly, we show that by decreasing a_2 in a network, the price of anarchy will increase. In Lemma 7, it was shown that any feasible Nash flow can be induced by a bimodal sensitivity distribution in which users are segregated on either link by their sensitivity. The price of anarchy for the network G with a Nash flow caused by s will therefore be,

$$\mathrm{PoA}(G, s, T(k)) = \frac{\ell_1(f_1^{\mathrm{Ne}})f_1^{\mathrm{Ne}} + \ell_2(f_2^{\mathrm{Ne}})f_2^{\mathrm{Ne}}}{\ell_1(f_1^{\mathrm{opt}})f_1^{\mathrm{opt}} + \ell_2(f_2^{\mathrm{opt}})f_2^{\mathrm{opt}}}. \tag{B.36}$$

Let us consider the case where $f_2^{\mathrm{Ne}} > f_2^{\mathrm{opt}}$. Now, consider a new network, $\hat{G}$ which replaces latency function $\ell_2(f) = a_2 f + b_2$ in G with $\hat{\ell}_2(f) = a_2 f + \hat{b}_2$ where $\hat{b}_2 = b_2 + \delta$ such that $\delta > 0$. Because the users are segregated on the links, the Nash flow will not change. Note that because $f_2^{\mathrm{Ne}} > f_2^{\mathrm{opt}}$

$$\frac{\ell_2(f_2^{\mathrm{Ne}})}{\ell_2(f_2^{\mathrm{opt}})} = \frac{a_2 f_2^{\mathrm{Ne}} + b_2}{a_2 f_2^{\mathrm{opt}} + b_2} < \frac{f_2^{\mathrm{Ne}}}{f_2^{\mathrm{opt}}}.$$

It can now be shown that

$$\begin{aligned}\frac{\mathcal{L}^{\mathrm{nf}}(G, s, T(k))}{\mathcal{L}^{\mathrm{opt}}(G)} &< \frac{\mathcal{L}^{\mathrm{nf}}(G, s, T(k)) + \delta f_2^{\mathrm{Ne}}}{\mathcal{L}^{\mathrm{opt}}(G) + \delta f_2^{\mathrm{opt}}} \\ &= \frac{\mathcal{L}^{\mathrm{nf}}(\hat{G}, s, T(k))}{\mathcal{L}^{\mathrm{opt}}(\hat{G})}.\end{aligned}$$

Thus the price of anarchy has increased in the new network $\hat{G}$, under the same sensitivity distribution and toll, when b_2 was increased, which has the same effect as decreasing the other terms and holding b_2 constant. A very similar argument can be followed for when $f_2^{\mathrm{Ne}} < f_2^{\mathrm{opt}}$ by picking $\hat{a_2} = a_2 - \delta$, and the price of anarchy then again increases. □

Proof of Proposition 7.3.2: The k that solves (7.17) equates the price of anarchy for G_β and G_α. It is shown in Lemma 10 that these networks realize the worst-case inefficiency and in Lemma 8 it is shown the worst-case distribution will be $s_\ell^{(\bar{s},G,k)}$ and $s_u^{(\bar{s},G,k)}$ respectively, both defined as the bimodal distribution $(S_{\mathrm{L}}, S_{\mathrm{U}})$ with mean $\bar{s}$ and mass R of the lower sensitivity type. To show this is optimal, it is sufficient to show that the price of anarchy for the network G_β is decreasing with k while the price of anarchy for the network G_α is increasing with k. If the networks have this relation with k, then the k that minimizes the price of anarchy must equalize them.

Consider a network $G \in \mathcal{G}^{\text{lc}}$ characterized by $\gamma = b_2/a_1$. If this satisfies that $s_\ell^{(\bar{s},G,k)} = (S_{\text{L}}, S_{\text{U}})$ or $s_u^{(\bar{s},G,k)} = (S_{\text{L}}, S_{\text{U}})$, then the price of anarchy for this network will be

$$\text{PoA}(G, \mathcal{S}(\bar{s}), T) = \begin{cases} \frac{R^2 - \gamma R + \gamma}{\gamma - \gamma^2/4}, & \gamma < 2 \\ R^2 - \gamma R + \gamma, & \gamma \geq 2. \end{cases} \tag{B.37}$$

This piecewise-continuous expression is locally minimized by $\gamma = 2R$, further, by differentiation, it can be observed that it is monotone decreasing for $0 < \gamma < 2R$ and monotone increasing for $\gamma > 2R$.

For the previously defined network G_β, under the bimodal distribution $(S_{\text{L}}, S_{\text{U}})$,

$$\beta = (1 + S_{\text{L}} k) R = R/z_L, \tag{B.38}$$

where β is dependent on the scaling factor k. From [194], the optimal scaling factor k will be in $(1/S_{\text{U}}, 1/S_{\text{L}})$. Therefore, for any k, $\beta < 2R$. The price of anarchy for this network is therefore monotone decreasing with β, and from (B.38), β is clearly increasing with k. The price of anarchy of the network is therefore decreasing with k. Similarly for G_α, under the distribution $(S_{\text{L}}, S_{\text{U}})$, the worst network is found by maximizing (B.37) over $\gamma = \alpha$, giving

$$\frac{1}{2(z_U - z_U^2)} = \alpha > 2R, \tag{B.39}$$

when $\frac{1}{2(z_U - z_U^2)} \geq R$, and $R/z_U = \alpha > 2R$, and the price of anarchy will be increasing with k. □

Bibliography

[1] B. L. Ferguson, D. Paccagnan, B. S. R. Pradelski, and J. R. Marden, *Collaborative Decision-Making and the k-Strong Price of Anarchy in Common Interest Games*, Nov., 2023. arXiv:2311.01379 [cs, eess].

[2] B. L. Ferguson, P. N. Brown, and J. R. Marden, *Information Signaling With Concurrent Monetary Incentives in Bayesian Congestion Games*, *IEEE Transactions on Intelligent Transportation Systems* (2024) 1–14. Conference Name: IEEE Transactions on Intelligent Transportation Systems.

[3] B. L. Ferguson, D. Paccagnan, and J. R. Marden, *The cost of informing decision-makers in multi-agent maximum coverage problems with random resource values*, *IEEE Control Systems Letters* **7** (2023) 2928–2933.

[4] B. L. Ferguson and J. R. Marden, *Robust Utility Design in Distributed Resource Allocation Problems with Defective Agents*, *Dynamic Games and Applications* (Aug., 2022) 1–23. Publisher: Birkhauser.

[5] B. L. Ferguson, P. N. Brown, and J. R. Marden, *The Effectiveness of Subsidies and Tolls in Congestion Games*, *IEEE Transactions on Automatic Control* **67** (2022), no. 6 2729–2742. arXiv: cs.GT/2102.09655.

[6] B. L. Ferguson, P. N. Brown, and J. R. Marden, *The Effectiveness of Subsidies and Taxes in Atomic Congestion Games*, *IEEE Control Systems Letters* **6** (2022) 614–619. Conference Name: IEEE Control Systems Letters.

[7] D. Paccagnan, R. Chandan, B. L. Ferguson, and J. R. Marden, *Optimal Taxes in Atomic Congestion Games*, *ACM Trans. Econ. Comput.* **9** (Aug., 2021). Publisher: Association for Computing Machinery Place: New York, NY, USA.

[8] B. L. Ferguson, P. N. Brown, and J. R. Marden, *Value of Information in Incentive Design: A Case Study in Simple Congestion Networks*, *IEEE Transactions on Computational Social Systems* **10** (Dec., 2023) 3077–3088. Conference Name: IEEE Transactions on Computational Social Systems.

[9] B. L. Ferguson, D. Paccagnan, B. S. R. Pradelski, and J. R. Marden, *Collaborative coalitions in multi-agent systems: Quantifying the strong price of*

anarchy for resource allocation games, in *2023 62nd IEEE conference on decision and control (CDC)*, pp. 3238–3243, 2023.

[10] R. Zhang, Y. Zhang, R. Konda, B. Ferguson, J. Marden, and N. Li, *Markov Games with Decoupled Dynamics: Price of Anarchy and Sample Complexity*, in *2023 62nd IEEE Conference on Decision and Control (CDC)*, pp. 8100–8107, Dec., 2023. ISSN: 2576-2370.

[11] A. K. Chen, B. L. Ferguson, D. Shishika, M. Dorothy, J. R. Marden, G. J. Pappas, and V. Kumar, *Path Defense in Dynamic Defender-Attacker Blotto Games (dDAB) with Limited Information*, in *2023 American Control Conference (ACC)*, pp. 447–453, May, 2023. ISSN: 2378-5861.

[12] B. L. Ferguson, D. Shishika, and J. R. Marden, *Ensuring the Defense of Paths and Perimeters in Dynamic Defender-Attacker Blotto Games (dDAB) on Graphs*, in *2022 58th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 1–7, Sept., 2022.

[13] B. L. Ferguson, P. N. Brown, and J. R. Marden, *Avoiding Unintended Consequences: How Incentives Aid Information Provisioning in Bayesian Congestion Games*, in *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 3781–3786, Dec., 2022. ISSN: 2576-2370.

[14] B. L. Ferguson and J. R. Marden, *Robust Utility Design in Distributed Resource Allocation Problems with Defective Agents*, in *2021 60th IEEE Conference on Decision and Control (CDC)*, pp. 1650–1655, IEEE, Dec., 2021.

[15] Y. Yue, B. L. Ferguson, and J. R. Marden, *Incentive Design for Congestion Games with Unincentivizable Users*, in *2021 60th IEEE Conference on Decision and Control (CDC)*, pp. 4515–4520, Dec., 2021. ISSN: 2576-2370.

[16] B. L. Ferguson and J. R. Marden, *The Impact of Fairness on Performance in Congestion Networks*, in *2021 American Control Conference (ACC)*, pp. 4521–4526, 2021.

[17] B. L. Ferguson, P. N. Brown, and J. R. Marden, *Carrots or Sticks? The Effectiveness of Subsidies and Tolls in Congestion Games*, in *2020 American Control Conference (ACC)*, pp. 1853–1858, July, 2020. ISSN: 2378-5861.

[18] R. Chandan, D. Paccagnan, B. L. Ferguson, and J. R. Marden, *Computing optimal taxes in atomic congestion games*, in *dl.acm.org*, pp. 1–1, 2019.

[19] B. L. Ferguson, P. N. Brown, and J. R. Marden, *Utilizing Information Optimally to Influence Distributed Network Routing*, in *Proceedings of the IEEE Conference on Decision and Control*, vol. 2019-Decem, pp. 5008–5013, Institute of Electrical and Electronics Engineers Inc., Dec., 2019. arXiv: 1907.10172 ISSN: 07431546.

[20] N. Tucker, B. Ferguson, and M. Alizadeh, *An Online Pricing Mechanism for Electric Vehicle Parking Assignment and Charge Scheduling*, in *2019 American Control Conference (ACC)*, pp. 5755–5760, July, 2019. ISSN: 2378-5861.

[21] B. Ferguson, V. Nagaraj, E. C. Kara, and M. Alizadeh, *Optimal Planning of Workplace Electric Vehicle Charging Infrastructure with Smart Charging Opportunities*, in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1149–1154, Nov., 2018. ISSN: 2153-0017.

[22] J. v. Neumann, *On the Theory of Games of Strategy*, in *1. On the Theory of Games of Strategy*, pp. 13–42. Princeton University Press, Mar., 2016.

[23] J. F. Nash Jr, *Equilibrium points in n-person games*, *Proceedings of the national academy of sciences* **36** (1950), no. 1 48–49. Publisher: National Acad Sciences.

[24] D. Monderer and L. S. Shapley, *Potential Games*, *Games and Economic Behavior* **14** (May, 1996) 124–143.

[25] J.-F. Mertens, *Repeated Games*, in *The New Palgrave Dictionary of Economics*, pp. 1–4. Palgrave Macmillan UK, London, 2016.

[26] R. Selten, *Spieltheoretische Behandlung Eines Oligopolmodells Mit Nachfrageträgheit: Teil I: Bestimmung Des Dynamischen Preisgleichgewichts*, *Zeitschrift für die gesamte Staatswissenschaft / Journal of Institutional and Theoretical Economics* **121** (1965), no. 2 301–324. Publisher: Mohr Siebeck GmbH & Co. KG.

[27] H. P. Young, *The Evolution of Conventions*, *Econometrica* **61** (1993), no. 1 57–84. ISBN: 0792335198.

[28] W. Vickrey, *Counterspeculation, auctions, and competitive sealed tenders*, *The Journal of Finance* **16** (1961) 8–37. ISBN: 00221082.

[29] D. Fudenberg and D. K. Levine, *The Theory of Learning in Games*. MIT Press, 1998. Google-Books-ID: G6vTQFluxuEC.

[30] J. D. Hartline, *Mechanism design and approximation*, *Book draft. October* (2013). ISBN: 9781612844596.

[31] A. E. Roth, *The Origins, History, and Design of the Resident Match*, *JAMA* **289** (Feb., 2003) 909–912.

[32] A. A. Cournot, *Researches Into the Mathematical Principles of the Theory of Wealth*. Macmillan, 1897. Google-Books-ID: eGgPAAAAYAAJ.

[33] A. C. Pigou, *The Economics of Welfare*. Macmillan, New York, 1920.

[34] E. Koutsoupias and C. Papadimitriou, *Worst-Case Equilibria*, *Computer Science Review* **3** (May, 2009) 65–69. Publisher: Elsevier Inc.

[35] T. Roughgarden, *How Bad Is Selfish Routing?*, *J. ACM* **49** (2002), no. 2 236–259.

[36] M. Gairing, *Covering games: Approximation through non-cooperation*, in *Internet and Economics*, pp. 184–195, 2009. ISSN: 03029743.

[37] S. Li and T. Başar, *Distributed algorithms for the computation of noncooperative equilibria*, *Automatica* **23** (July, 1987) 523–533.

[38] T. Başar, *Relaxation techniques and asynchronous algorithms for on-line computation of non-cooperative equilibria*, *Journal of Economic Dynamics and Control* **11** (Dec., 1987) 531–549.

[39] J. Shamma and G. Arslan, *Dynamic fictitious play, dynamic gradient play, and distributed convergence to Nash equilibria*, *IEEE Transactions on Automatic Control* **50** (2005), no. 3 312–327.

[40] S. Wollenstein-Betech, A. Houshmand, M. Salazar, M. Pavone, C. G. Cassandras, and I. C. Paschalidis, *Congestion-aware Routing and Rebalancing of Autonomous Mobility-on-Demand Systems in Mixed Traffic*, in *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1–7, 2020.

[41] A. Khamis, A. Hussein, and A. Elmogy, *Multi-robot task allocation: A review of the state-of-the-art*, *Cooperative robots and sensor networks 2015* (2015) 31–51. Publisher: Springer.

[42] V. Ranganathan, P. Kumar, U. Kaur, S. H. Li, T. Chakraborty, and R. Chandra, *Re-Inventing the Food Supply Chain with IoT: A Data-Driven Solution to Reduce Food Loss*, *IEEE Internet of Things Magazine* **5** (Mar., 2022) 41–47. Conference Name: IEEE Internet of Things Magazine.

[43] K. Tsakalozos, H. Kllapi, E. Sitaridi, M. Roussopoulos, D. Paparas, and A. Delis, *Flexible use of cloud resources through profit maximization and price discrimination*, in *Proc. International Conference on Data Engineering*, pp. 75–86, 2011. ISSN: 10844627.

[44] F. G. Filip, *Decision support and control for large-scale complex systems*, *Annual Reviews in Control* **32** (Apr., 2008) 61–70.

[45] C. Daini, P. Goatin, M. L. D. Monache, and A. Ferrara, *Centralized Traffic Control via Small Fleets of Connected and Automated Vehicles*, in *2022 European Control Conference (ECC)*, pp. 371–376, July, 2022.

[46] L. Fang and H. Li, *Centralized resource allocation based on the cost–revenue analysis*, *Computers & Industrial Engineering* **85** (July, 2015) 395–401.

[47] G. Antonelli, *Interconnected dynamic systems: An overview on distributed control*, *IEEE Control Systems Magazine* **33** (2013), no. 1 76–88. Publisher: IEEE.

[48] J. R. Marden, G. Arslan, and J. S. Shamma, *Cooperative Control and Potential Games*, *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* **39** (2009), no. 6 1393–1407.

[49] R. M. Murray, *Recent research in cooperative control of multivehicle systems*, *Journal of Dynamic Systems, Measurement, and Control* **129** (May, 2007) 571–583. tex.eprint: https://asmedigitalcollection.asme.org/dynamicsystems/article-pdf/129/5/571/5538275/571_1.pdf.

[50] A. Das, T. Gervet, J. Romoff, D. Batra, D. Parikh, M. Rabbat, and J. Pineau, *Tarmac: Targeted multi-agent communication*, in *International conference on machine learning*, pp. 1538–1546, 2019. tex.organization: PMLR.

[51] D. D. Šiljak and A. I. Zečević, *Control of large-scale systems: Beyond decentralized feedback*, *Annual Reviews in Control* **29** (Jan., 2005) 169–179.

[52] Z. Xu and V. Tzoumas, *Resource-Aware Distributed Submodular Maximization: A Paradigm for Multi-Robot Decision-Making*, in *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 5959–5966, Dec., 2022. ISSN: 2576-2370.

[53] G. Orosz, *Connected cruise control: modelling, delay effects, and nonlinear behaviour*, *Vehicle System Dynamics* **54** (2016), no. 8 1147–1176. ISBN: 0042-3114 1744-5159.

[54] H. Nawaz, H. M. Ali, and A. A. Laghari, *UAV Communication Networks Issues: A Review*, *Archives of Computational Methods in Engineering* **28** (May, 2021) 1349–1369.

[55] A. Lazaridou and M. Baroni, *Emergent multi-agent communication in the deep learning era*, *arXiv preprint arXiv:2006.02419* (2020).

[56] R. J. Aumann, *Acceptable points in general cooperative n-person games*, *Contributions to the Theory of Games* **4** (1959) 287–324.

[57] R. Nessah and G. Tian, *On the existence of strong Nash equilibria*, *Journal of Mathematical Analysis and Applications* **414** (2014), no. 2 871–885. Publisher: Elsevier.

[58] N. Gatti, M. Rocco, and T. Sandholm, *On the verification and computation of strong Nash equilibrium*, *arXiv preprint arXiv:1711.06318* (2017).

[59] R. Holzman and N. Law-Yone, *Strong equilibrium in congestion games*, *Games and economic behavior* **21** (1997), no. 1-2 85–101. Publisher: Elsevier.

[60] T. Harks, M. Klimm, and R. H. Möhring, *Strong Nash equilibria in games with the lexicographical improvement property*, in *Internet and network economics: 5th international workshop, WINE 2009, rome, italy, december 14-18, 2009. Proceedings 5*, pp. 463–470, 2009. tex.organization: Springer.

[61] J. B. Clempner and A. S. Poznyak, *Finding the strong nash equilibrium: Computation, existence and characterization for markov games*, *Journal of Optimization Theory and Applications* **186** (2020) 1029–1052. Publisher: Springer.

[62] A. Epstein, M. Feldman, and Y. Mansour, *Strong equilibrium in cost sharing connection games*, in *Proceedings of the 8th ACM conference on Electronic commerce*, EC '07, (New York, NY, USA), pp. 84–92, Association for Computing Machinery, June, 2007.

[63] N. Andelman, M. Feldman, and Y. Mansour, *Strong price of anarchy*, *Games and Economic Behavior* **65** (Mar., 2009) 289–317.

[64] J. Barreiro-Gomez, G. Obando, and N. Quijano, *Distributed Population Dynamics: Optimization and Control Applications*, *IEEE Transactions on Systems, Man, and Cybernetics: Systems* (2016) 1–11.

[65] A. Fiat, H. Kaplan, M. Levy, and S. Olonetsky, *Strong price of anarchy for machine load balancing*, in *ICALP*, vol. 4596, pp. 583–594, 2007. tex.organization: Springer.

[66] L. Epstein and R. van Stee, *The price of anarchy on uniformly related machines revisited*, *Information and Computation* **212** (2012) 37–54. Publisher: Elsevier.

[67] S. Chien and A. Sinclair, *Strong and pareto price of anarchy in congestion games.*, in *ICALP (1)*, pp. 279–291, 2009. tex.organization: Citeseer.

[68] Y. Bachrach, V. Syrgkanis, E. Tardos, and M. Vojnović, *Strong Price of Anarchy, Utility Games and Coalitional Dynamics*, in *Algorithmic Game Theory* (R. Lavi, ed.), Lecture Notes in Computer Science, (Berlin, Heidelberg), pp. 218–230, Springer, 2014.

[69] M. Feldman and O. Friedler, *A unified framework for strong price of anarchy in clustering games*, in *Automata, languages, and programming: 42nd international colloquium, ICALP 2015, kyoto, japan, july 6-10, 2015, proceedings, part II*, pp. 601–613, 2015. tex.organization: Springer.

[70] T. Roughgarden, *Intrinsic robustness of the price of anarchy*, *Communications of the ACM* **55** (2012), no. 7 116–123. ISBN: 9781605585062.

[71] V. Bilò and C. Vinci, *Dynamic taxes for polynomial congestion games*, in *EC 2016 - Proceedings of the 2016 ACM Conference on Economics and Computation*, (New York, New York, USA), pp. 839–856, ACM Press, 2016.

[72] 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** (Nov., 2020) 4616–4631. arXiv: 1807.01333.

[73] R. Zhang, Y. Zhang, R. Konda, B. Ferguson, J. Marden, and N. Li, *Markov Games with Decoupled Dynamics: Price of Anarchy and Sample Complexity*, Apr., 2023. arXiv:2304.03840 [cs, math].

[74] R. Konda, R. Chandan, D. Grimsman, and J. R. Marden, *Optimal Design of Best Response Dynamics in Resource Allocation Games*, Apr., 2022. arXiv:2204.10364 [cs, eess].

[75] W. Saad, Z. Han, M. Debbah, A. Hjorungnes, and T. Basar, *Coalitional game theory for communication networks*, *Ieee signal processing magazine* **26** (2009), no. 5 77–97. Publisher: IEEE.

[76] H. Bayram and H. I. Bozma, *Multirobot communication network topology via centralized pairwise games*, in *2013 IEEE international conference on robotics and automation*, pp. 2521–2526, 2013. tex.organization: IEEE.

[77] D. Cappello and T. Mylvaganam, *Distributed differential games for control of multi-agent systems*, *IEEE Transactions on Control of Network Systems* **9** (2021), no. 2 635–646. Publisher: IEEE.

[78] A. Vetta, *Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions*, *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.* (2002) 416–425. ISBN: 0-7695-1822-2.

[79] S. Durand and B. Gaujal, *Complexity and Optimality of the Best Response Algorithm in Random Potential Games*, in *Algorithmic Game Theory* (M. Gairing and R. Savani, eds.), (Berlin, Heidelberg), pp. 40–51, Springer, 2016.

[80] Y. Giannakopoulos, *A Smoothed FPTAS for Equilibria in Congestion Games*, July, 2023. arXiv:2306.10600 [cs].

[81] G. A. Korsah, A. Stentz, and M. B. Dias, *A comprehensive taxonomy for multi-robot task allocation*, *The International Journal of Robotics Research* **32** (2013), no. 12 1495–1512. Publisher: SAGE Publications Sage UK: London, England.

[82] N. Mathew, S. L. Smith, and S. L. Waslander, *Planning paths for package delivery in heterogeneous multirobot teams*, *IEEE Transactions on Automation Science and Engineering* **12** (2015), no. 4 1298–1308. Publisher: IEEE.

[83] J. R. Marden and J. S. Shamma, *Game Theory and Distributed Control*, in *Handbook of Game Theory Vol. 4* (H. P. Young and S. Zamir, eds.). Elsevier Science, 2014.

[84] B. Swenson, R. Murray, and S. Kar, *On best-response dynamics in potential games*, *SIAM Journal on Control and Optimization* **56** (July, 2018) 2734–2767. arXiv: 1707.06465 Publisher: Society for Industrial and Applied Mathematics Publications.

[85] S. Rahili, J. Lu, W. Ren, and U. Al-Saggaf, *Distributed Coverage Control of Mobile Sensor Networks in Unknown Environment Using Game Theory: Algorithms and Experiments*, *IEEE Transactions on Mobile Computing* **17** (2018), no. 6 1303–1313.

[86] V. Ramaswamy, D. Paccagnan, and J. R. Marden, *Multiagent Maximum Coverage Problems: The Tradeoff Between Anarchy and Stability*, *IEEE Transactions on Automatic Control* **67** (Apr., 2022) 1698–1712. Publisher: Institute of Electrical and Electronics Engineers Inc.

[87] D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar, *Informational Braess' Paradox: The Effect of Information on Traffic Congestion*, *Operations Research* **66** (2018), no. 4.

[88] E. Kamenica and M. Gentzkow, *Bayesian Persuasion*, *American Economic Review* **101** (Oct., 2011) 2590–2615.

[89] M. Castiglioni, A. Celli, A. Marchesi, and N. Gatti, *Signaling in Bayesian Network Congestion Games: the Subtle Power of Symmetry*, *Proceedings of the AAAI Conference on Artificial Intelligence* **35** (May, 2021) 5252–5259. Number: 6.

[90] H. Tavafoghi and D. Teneketzis, *Informational incentives for congestion games*, in *55th Annual Allerton Conference on Communication, Control, and Computing, Allerton 2017*, vol. 2018-Janua, pp. 1285–1292, Institute of Electrical and Electronics Engineers Inc., Jan., 2018.

[91] Y. Chen and J. Zhang, *Signalling by Bayesian Persuasion and Pricing Strategy*, *The Economic Journal* **130** (May, 2020) 976–1007. Publisher: Oxford Academic.

[92] S. Shah, S. Amin, and P. Jaillet, *Optimal Information Provision for Strategic Hybrid Workers*, in *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 3807–3814, Dec., 2022. ISSN: 2576-2370.

[93] F. Sezer, H. Khazaei, and C. Eksin, *Social Welfare Maximization and Conformism via Information Design in Linear-Quadratic-Gaussian Games*, *arXiv preprint arXiv:2102.13047* (2021).

[94] Y. Zhu and K. Savla, *Information Design in Nonatomic Routing Games With Partial Participation: Computation and Properties*, *IEEE Transactions on Control of Network Systems* **9** (June, 2022) 613–624. Publisher: Institute of Electrical and Electronics Engineers Inc.

[95] M. Wu and S. Amin, *Information design for regulating traffic flows under uncertain network state*, in *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 671–678, IEEE, 2019.

[96] D. Paccagnan, R. Chandan, and J. R. Marden, *Utility and mechanism design in multi-agent systems: An overview*, *Annual Reviews in Control* **53** (Jan., 2022) 315–328.

[97] H. Kitano, S. Tadokoro, I. Noda, H. Matsubara, T. Takahashi, A. Shinjou, and S. Shimada, *RoboCup Rescue: search and rescue in large-scale disasters as a domain for autonomous agents research*, in *IEEE SMC'99 Conference Proceedings. 1999 IEEE International Conference on Systems, Man, and Cybernetics (Cat. No.99CH37028)*, vol. 6, pp. 739–743 vol.6, Oct., 1999. ISSN: 1062-922X.

[98] C. E. Nugraheni and L. Abednego, *Multi Agent Hyper-Heuristics based framework for production scheduling problem*, in *2016 International Conference on Informatics and Computing (ICIC)*, pp. 309–313, Oct., 2016.

[99] R. Brown, F. Rossi, K. Solovey, M. Tsao, M. T. Wolf, and M. Pavone, *On Local Computation for Network-Structured Convex Optimization in Multi-Agent Systems*, *IEEE Transactions on Control of Network Systems* (2021) 1.

[100] J. R. Marden, *The Role of Information in Distributed Resource Allocation*, *IEEE Transactions on Control of Network Systems* **4** (2017), no. 3 654–664.

[101] F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S. L. Smith, *Dynamic Vehicle Routing for Robotic Systems*, *Proceedings of the IEEE* **99** (2011), no. 9 1482–1504.

[102] L. Bertizzolo, S. D'Oro, L. Ferranti, L. Bonati, E. Demirors, Z. Guan, T. Melodia, and S. Pudlewski, *SwarmControl: An Automated Distributed Control Framework for Self-Optimizing Drone Networks*, in *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*, pp. 1768–1777, July, 2020. ISSN: 2641-9874.

[103] J. S. Shamma, *Cooperative control of distributed multi-agent systems*. Wiley Online Library, 2007.

[104] N. Li and J. R. Marden, *Designing games for distributed optimization*, *IEEE Journal on Selected Topics in Signal Processing* **7** (2013), no. 2 230–242. ISBN: 9781612848006.

[105] R. Chandan, D. Paccagnan, and J. R. Marden, *When Smoothness is Not Enough: Toward Exact Quantification and Optimization of the Price of Anarchy*, *Proceedings of the IEEE Conference on Decision and Control* **2019-Decem** (Nov., 2019) 4041–4046. arXiv: 1911.07823 Publisher: Institute of Electrical and Electronics Engineers Inc.

[106] D. Grimsman, M. S. Ali, J. P. Hespanha, and J. R. Marden, *Impact of information in greedy submodular maximization*, *2017 IEEE 56th Annual Conference on Decision and Control, CDC 2017* **2018-Janua** (2018), no. Cdc 2900–2905. arXiv: 1807.10639 ISBN: 9781509028733.

[107] H. Jaleel, W. Abbas, and J. S. Shamma, *Robustness Of Stochastic Learning Dynamics To Player Heterogeneity In Games*, in *2019 IEEE 58th Conference on Decision and Control (CDC)*, pp. 5002–5007, Dec., 2019. ISSN: 2576-2370.

[108] E. Yildiz, D. Acemoglu, A. Ozdaglar, A. Saberi, and A. Scaglione, *Discrete Opinion Dynamics with Stubborn Agents*, *Public Choice: Analysis of Collective Decision-Making eJournal* (2011).

[109] G. W. Brown, *Iterative Solution of Games by Fictitious Play*, in *Activity Analysis of Production and Allocation, T. C. Koopmans (Ed.)*, pp. 374–376. Wiley, New York, 1951.

[110] J. Hannan, *Approximation to Bayes Risk in Repeated Plays*, in *Contribution to the Theory of Games*, pp. 97–139. Princeton University Press, Princeton, NJ, 1957.

[111] D. Paccagnan and M. Gairing, *In Congestion Games, Taxes Achieve Optimal Approximation*, *ArXiV* (May, 2021). arXiv: 2105.07480.

[112] M. Goemans, E. Li, V. Mirrokni, and M. Thottan, *Market sharing games applied to content distribution in ad-hoc networks*, pp. 55–66, 2004.

[113] M. Salazar, M. Tsao, I. Aguiar, M. Schiffer, and M. Pavone, *A congestion-aware routing scheme for autonomous mobility-on-demand systems*, in *European Control Conference*, 2019.

[114] W. Mei, N. E. Friedkin, K. Lewis, and F. Bullo, *Dynamic models of appraisal networks explaining collective learning*, in *IEEE Conference on Decision and Control, CDC*, pp. 3554–3559, 2016. arXiv: 1609.09546 ISSN: 00189286.

[115] R. Johari, G. Y. Weintraub, and B. Van Roy, *Investment and Market Structure in Industries with Congestion*, *Operations Research* **58** (2010), no. 5 1303–1317. ISBN: 0030-364X.

[116] P. Palensky and D. Dietrich, *Demand side management: Demand response, intelligent energy systems, and smart loads*, *IEEE Transactions on Industrial Informatics* **7** (Aug., 2011) 381–388.

[117] C. Narasimhan, *A Price Discrimination Theory of Coupons*, *Marketing Science* **3** (1984), no. 2 128–147. ISBN: 07322399.

[118] M. Esmaeili, M. B. Aryanezhad, and P. Zeephongsekul, *A game theory approach in seller-buyer supply chain*, *European Journal of Operational Research* **195** (June, 2009) 442–448.

[119] C. Papadimitriou, *Algorithms, games, and the internet*, in *STOC '01: Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pp. 749–753, 2001. ISSN: 03029743.

[120] P. Kleer and G. Schäfer, *Path deviations outperform approximate stability in heterogeneous congestion games*, in *International Symposium on Algorithmic Game Theory*, pp. 212–224, 2017. arXiv: 1707.01278.

[121] L. J. Ratliff, R. Dong, S. Sekar, and T. Fiez, *A Perspective on Incentive Design: Challenges and Opportunities*, *Annual Review of Control, Robotics, and Autonomous Systems* **2** (2018), no. 1 1–34.

[122] R. Cole, Y. Dodis, and T. Roughgarden, *Pricing network edges for heterogeneous selfish users*, in *Proc. of the 35th ACM symp. on Theory of computing*, (New York, New York, USA), pp. 521–530, 2003. ISSN: 07349025.

[123] D. Fotakis, G. Karakostas, and S. G. Kolliopoulos, *On the existence of optimal taxes for network congestion games with heterogeneous users*, *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **6386 LNCS** (2010), no. M4D 162–173. ISBN: 3642161693.

[124] I. Caragiannis, C. Kaklamanis, and P. Kanellopoulos, *Taxes for Linear Atomic Congestion Games*, *ACM Trans. Algorithms* **7** (Dec., 2010) 13:1–13:31. Publisher: ACM Place: New York, NY, USA.

[125] L. Fleischer, K. Jain, and M. Mahdian, *Tolls for Heterogeneous Selfish Users in Multicommodity Networks and Generalized Congestion Games*, in *Proc. 45th IEEE Symp. on Foundations of Computer Science*, (Rome, Italy), pp. 277–285, 2004.

[126] W. Zou, M. Yu, and S. MIZOKAMI, *Mechanism Design for an Incentive Subsidy Scheme for Bus Transport*, *Sustainability* **11** (Mar., 2019) 1740. Publisher: MDPI AG.

[127] T. A. Taylor, *Supply Chain Coordination Under Channel Rebates with Sales Effort Effects*, tech. rep., 2002.

[128] S. A. Bagloee and M. Sarvi, *A modern congestion pricing policy for urban traffic: subsidy plus toll*, *Journal of Modern Transportation* **25** (Sept., 2017) 133–149. Publisher: Springer Berlin Heidelberg.

[129] L. Chen and H. Yang, *Managing congestion and emissions in road networks with tolls and rebates*, *Transportation Research Part B: Methodological* **46** (Sept., 2012) 933–948. Publisher: Elsevier Ltd.

[130] E. Gerstner, J. D. Hess, E. Gerstner, and J. D. Hess, *A Theory of Channel Price Promotions*, *American Economic Review* **81** (1991), no. 4 872–86. Publisher: American Economic Association.

[131] A. Ali, M. A. Jolson, and R. Y. Darmon, *A model for optimizing the refund value in rebate promotions*, *Journal of Business Research* **29** (1994), no. 3 239–245.

[132] T. Roughgarden, *Selfish Routing and the Price of Anarchy.* MIT Press, 2005.

[133] V. Bonifaci, M. Salek, and G. Schäfer, *Efficiency of restricted tolls in non-atomic network routing games*, in *Algorithmic Game Theory*, vol. 6982 LNCS, pp. 302–313, 2011. arXiv: dx.doi.org/10.1007/978-3-642-24829-0_27 ISSN: 03029743.

[134] D. Fotakis and P. G. Spirakis, *Cost-Balancing Tolls for Atomic Network Congestion Games*, in *Internet and Network Economics*, pp. 179–190. Springer Berlin Heidelberg, Berlin, Heidelberg, Dec., 2007.

[135] G. Karakostas and S. G. Kolliopoulos, *Edge pricing of multicommodity networks for heterogeneous selfish users*, in *Proc. 45th IEEE Symp. on Foundations of Computer Science*, (Rome, Italy), pp. 268–276, 2004.

[136] P. N. Brown and J. R. Marden, *Optimal Mechanisms for Robust Coordination in Congestion Games*, *IEEE Transactions on Automatic Control* **63** (2018), no. 8 2437–2448.

[137] I. Arieli, *Transfer Implementation in Congestion Games*, *Dynamic Games and Applications* **5** (2015), no. 2 228–238.

[138] P. Maillé and N. E. Stier-Moses, *Eliciting Coordination with Rebates*, *Transportation Science* **43** (2009), no. 4 473–492.

[139] W. H. Sandholm, *Evolutionary Implementation and Congestion Pricing*, *The Review of Economic Studies* **69** (2002), no. 3 667–689.

[140] A. Mas-Colell, *On a Theorem of Schmeidler*, *Mathematical Economics* **13** (1984) 201–206.

[141] T. Roughgarden, *The price of anarchy is independent of the network topology*, *J. Computer and System Sciences* **67** (Sept., 2003) 341–364.

[142] P. N. Brown and J. R. Marden, *The Robustness of Marginal-Cost Taxes in Affine Congestion Games*, *Transactions on Automatic Control* **62** (2017), no. 8 3999–4004. ISBN: 0001409107.

[143] G. Hardin, *The Tragedy of the Commons*, *Science* **162** (1968), no. 3859 1243–1248.

[144] J. Liu, S. Amin, and G. Schwartz, *Effects of information heterogeneity in Bayesian routing games*, *arXiv preprint arXiv:1603.08853* (2016).

[145] D. Bergemann and S. Morris, *Information design: A unified perspective*, *Journal of Economic Literature* **57** (2019), no. 1 44–95.

[146] R. Lindsey, T. Daniel, E. Gisches, and A. Rapoport, *Pre-trip information and route-choice decisions with stochastic travel conditions: Theory*, *Transportation Research Part B: Methodological* **67** (2014), no. C 187–207.

[147] A. Fleischer, A. Tchetchik, and T. Toledo, *Does it pay to reveal safety information? The effect of safety information on flight choice*, *Transportation Research Part C: Emerging Technologies* **56** (July, 2015) 210–220.

[148] O. Rotem-Mindali and I. Salomon, *The impacts of E-retail on the choice of shopping trips and delivery: Some preliminary findings*, *Transportation Research Part A: Policy and Practice* **41** (Feb., 2007) 176–189.

[149] T. Rambha, S. D. Boyles, A. Unnikrishnan, and P. Stone, *Marginal cost pricing for system optimal traffic assignment with recourse under supply-side uncertainty*, *Transportation Research Part B: Methodological* **110** (2018) 104–121.

[150] Y. Liu and Z. Yang, *Information provision and congestion pricing in a risky two-route network with heterogeneous travelers*, *Transportation Research Part C: Emerging Technologies* **128** (July, 2021) 103083.

[151] J. Wardrop, *Some Theoretical Aspects of Road Traffic Research*, in *Proc. of the Institution of Civil Engineers, Part II, Vol. 1, No. 36*, pp. 352–362, 1952.

[152] R. W. Rosenthal, *A class of games possessing pure-strategy Nash equilibria*, *International Journal of Game Theory* **2** (Dec., 1973) 65–67. Publisher: Physica-Verlag.

[153] C. S. Fisk, *Game theory and transportation systems modelling*, *Transportation Research Part B: Methodological* **18** (Aug., 1984) 301–313. Publisher: Pergamon.

[154] D. Braess, A. Nagurney, and T. Wakolbinger, *On a Paradox of Traffic Planning*, *Transportation Science* **39** (Nov., 2005) 446–450.

[155] D. Levinson, *Micro-foundations of congestion and pricing: A game theory perspective*, *Transportation Research Part A: Policy and Practice* **39** (Aug., 2005) 691–704. Publisher: Pergamon.

[156] F. Farokhi and K. H. Johansson, *A Study of Truck Platooning Incentives Using a Congestion Game*, *IEEE Transactions on Intelligent Transportation Systems* **16** (2015), no. 2 581–595. arXiv: 1310.5534v2 ISBN: 1524-9050 VO - 16.

[157] A. Yekkehkhany and R. Nagi, *Risk-Averse Equilibria for Vehicle Navigation in Stochastic Congestion Games*, *IEEE Transactions on Intelligent Transportation Systems* (2022) 1–17.

[158] S. Spana, L. Du, and Y. Yin, *Strategic Information Perturbation for an Online In-Vehicle Coordinated Routing Mechanism for Connected Vehicles Under Mixed-Strategy Congestion Game*, *IEEE Transactions on Intelligent Transportation Systems* **23** (2022), no. 5 4541–4555.

[159] Y. Chen and V. Gupta, *Toll Design for Routing Games With Stochastic Demands*, *IEEE Control Systems Letters* **6** (2022) 3445–3450.

[160] Y. Zhu and K. Savla, *A Semidefinite Approach to Information Design in Non-atomic Routing Games*, *CoRR* **abs/2005.0** (2020).

[161] M. Wu, S. Amin, and A. E. Ozdaglar, *Value of Information in Bayesian Routing Games*, *Operations Research* **69** (Jan., 2021) 148–163. arXiv: 1808.10590.

[162] J. Verbree and A. Cherukuri, *Inferring the prior in routing games using public signalling*, 2021.

[163] M. H. Lotfi, R. J. La, and N. C. Martins, *Bayesian Congestion Game with Traffic Manager: Binary Signal Case*, in *2018 IEEE Conference on Decision and Control (CDC)*, pp. 327–333, 2018.

[164] O. Massicot and C. Langbort, *Competitive comparisons of strategic information provision policies in network routing games*, *IEEE Transactions on Control of Network Systems* **9** (2022), no. 4 1589–1599.

[165] O. Massicot and C. Langbort, *Public signals and persuasion for road network congestion games under vagaries*, *IFAC-PapersOnLine* **51** (2019), no. 34 124–130. Publisher: Elsevier.

[166] E. T. Verhoef, R. H. M. Emmerink, P. Nijkamp, and P. Rietveld, *Information provision, flat and fine congestion tolling and the efficiency of road usage*, *Regional Science and Urban Economics* **26** (Aug., 1996) 505–529.

[167] J. Enrique Fernández L., J. de Cea Ch, and G. Germán Valverde, *Effect of advanced traveler information systems and road pricing in a network with non-recurrent congestion*, *Transportation Research Part A: Policy and Practice* **43** (June, 2009) 481–499.

[168] L. M. Gardner, S. D. Boyles, and S. T. Waller, *Quantifying the benefit of responsive pricing and travel information in the stochastic congestion pricing problem*, *Transportation Research Part A: Policy and Practice* **45** (Mar., 2011) 204–218.

[169] R. Zhang and E. T. Verhoef, *A monopolistic market for advanced traveller information systems and road use efficiency*, *Transportation Research Part A: Policy and Practice* **40** (June, 2006) 424–443.

[170] A. de Palma, R. Lindsey, and N. Picard, *Risk Aversion, the Value of Information, and Traffic Equilibrium*, *Transportation Science* **46** (2012).

[171] R. H. Emmerink, E. T. Verhoef, P. Nijkamp, and P. Rietveld, *Endogenising demand for information in road transport*, *The Annals of Regional Science* **30** (1996) 201–222. Publisher: Springer.

[172] S. Gubins, E. T. Verhoef, and T. de Graaff, *Welfare effects of road pricing and traffic information under alternative ownership regimes*, *Transportation Research Part A: Policy and Practice* **46** (Oct., 2012) 1304–1317.

[173] H. Yang, *Evaluating the benefits of a combined route guidance and road pricing system in a traffic network with recurrent congestion*, *Transportation* **26** (1999) 299–322. Publisher: Springer.

[174] C. Chen, G. Zhang, H. Wang, J. Yang, P. J. Jin, and C. Michael Walton, *Bayesian network-based formulation and analysis for toll road utilization supported by traffic information provision*, *Transportation Research Part C: Emerging Technologies* **60** (Nov., 2015) 339–359.

[175] W. H. Sandholm, *Potential Games with Continuous Player Sets*, *Journal of Economic Theory* **97** (2001) 81–108.

[176] F. Farokhi and K. H. Johansson, *A piecewise-constant congestion taxing policy for repeated routing games*, *Transportation Research Part B: Methodological* **78** (2015) 123–143. Publisher: Elsevier Ltd ISBN: 9781510811539.

[177] R. Singh and R. Dowling, *Improved Speed-Flow Relationships: Application To Transportation Planning Models*, *Seventh TRB Conference on the Application of Transportation Planning Methods* (2002), no. 4 340–349.

[178] S. P. Boyd and L. Vandenberghe, *Convex optimization.* Cambridge university press, 2004.

[179] U. Bhaskar, Y. Cheng, Y. K. Ko, and C. Swamy, *Hardness results for signaling in bayesian zero-sum and network routing games*, in *Proceedings of the 2016 ACM conference on economics and computation*, pp. 479–496, 2016.

[180] S. M. Griesbach, M. Hoefer, M. Klimm, and T. Koglin, *Public Signals in Network Congestion Games*, 2022.

[181] J. B. Lasserre, *A semidefinite programming approach to the generalized problem of moments*, *Mathematical Programming 2007 112:1* **112** (Feb., 2007) 65–92. Publisher: Springer.

[182] D. Henrion, J. B. Lasserre, and J. Lofberg, *GloptiPoly 3: Moments, optimization and semidefinite programming*, *Optimization Methods and Software* **24** (Aug., 2009) 761–779. arXiv: 0709.2559.

[183] R. Johari and J. N. Tsitsiklis, *Efficiency Loss in a Network Resource Allocation Game*, *Mathematics of Operations Research* **29** (Aug., 2004) 407–435.

[184] G. Piliouras, E. Nikolova, and J. S. Shamma, *Risk sensitivity of price of anarchy under uncertainty*, in *Proceedings of the 14th ACM Conference on Electronic Commerce*, vol. 9, (New York, New York, USA), pp. 715–732, 2013. Issue: 4.

[185] J. M. Hoque, G. D. Erhardt, D. Schmitt, M. Chen, and M. Wachs, *Estimating the uncertainty of traffic forecasts from their historical accuracy*, *Transportation Research Part A: Policy and Practice* **147** (May, 2021) 339–349.

[186] R. Tang, S. Wang, and H. Li, *Game theory based interactive demand side management responding to dynamic pricing in price-based demand response of smart grids*, *Applied Energy* **250** (Sept., 2019) 118–130.

[187] R. Schroll and R. Grohs, *Uncertainty in Prerelease Advertising*, *Journal of Advertising* **48** (Mar., 2019) 167–180. Publisher: Routledge _eprint: https://doi.org/10.1080/00913367.2018.1532853.

[188] J. Verbree and A. Cherukuri, *Inferring the prior in routing games using public signalling*, Feb., 2023. arXiv:2109.05895 [cs, eess, math].

[189] G. Kreindler, *Peak-Hour Road Congestion Pricing Experimental Evidence and Equilibrium Implications*, *Harvard University* (2020).

[190] P. A. Akbar, V. Couture, G. Duranton, and A. Storeygard, *Mobility and Congestion in Urban India*, Nov., 2018.

[191] N. Buchholz, L. Doval, J. Kastl, F. Matějka, and T. Salz, *The Value of Time: Evidence From Auctioned Cab Rides*, May, 2020.

[192] P.-A. Chen, B. D. Keijzer, D. Kempe, and G. Schäfer, *Altruism and Its Impact on the Price of Anarchy*, *ACM Transactions on Economics and Computation* **2** (Oct., 2014) 17:1–17:45.

[193] P. N. Brown and J. R. Marden, *Can Taxes Improve Congestion on All Networks?*, *IEEE Transactions on Control of Network Systems* **7** (Dec., 2020) 1643–1653. Conference Name: IEEE Transactions on Control of Network Systems.

[194] P. N. Brown and J. R. Marden, *The Benefit of Perversity in Distributed Network Routing*, in *56th IEEE Conf. on Decision and Control*, (Melbourne, Australia), pp. 6229–6234, 2017.

[195] I. Milchtaich, *Network topology and the efficiency of equilibrium*, *Games and Economic Behavior* **57** (2006), no. 2 321–346.